

The London School of Economics and Political Science

BRINGING ABOUT EQUALITY

Christopher John Marshall

A thesis submitted to the Department of Philosophy, Logic, and Scientific Method
for the degree of Doctor of Philosophy, September 2017

DECLARATION

I certify that the thesis I have presented for examination for the PhD degree of the London School of Economics and Political Science is solely my own work. The copyright of this thesis rests with the author. Quotation from it is permitted, provided that full acknowledgement is made. The thesis may not be reproduced without my prior written consent. I warrant that this authorisation does not, to the best of my belief, infringe the rights of any third party.

I declare that my thesis consists of 50,371 words.

ABSTRACT

Sometimes people ought to do something for the sake of equality but it would be wrong to force them to do it. Contrariwise, sometimes it would be permissible to force people to do something for the sake of equality but it is unclear whether they ought to do it without coercion.

This gives rise to moral obstacles to redistributing benefits and burdens in unequal situations of different degrees and kinds. In these situations, what are individuals required to do, permitted to do, and forbidden from doing for the sake of equality? In this thesis, I address four aspects of this problem. I defend the following central claims.

When we fall short of realising equality using the coercive power of the state, individuals have an egalitarian reason to take up the moral slack and voluntarily redistribute their own wealth, even when this makes them worse off than they would be if society were equal.

It is sometimes permissible for individuals to flout property laws in the name of equality by dispossessing people of things that they are legally, but not morally, entitled to. This is often true when property laws are severely unjust and is sometimes true when property laws are only partially unjust.

Egalitarian justice requires people to benefit each other by making a productive contribution to society in order to realise a Pareto optimal level of equality. This is neither implausibly morally demanding, nor commits one to the view that justice licences legally coercive job allocation.

Sometimes, due to brute bad luck, indivisible harm will occur and cannot be fairly divided. Some views say that it is wrong to harm innocent people in the course of defending yourself from an equivalent harm, but permissible to “duck” harm even though you foresee that doing so will lead to an innocent person being harmed as a result. I argue that, by the lights of such views, the mode of agency used to impose or redistribute harm makes no difference to permissibility.

ACKNOWLEDGEMENTS

Thanks to Mike Otsuka, Alex Voorhoeve, Gabriel Wollner, Susanne Burri, Tom Rowe, Adam White, Goreti Faria, Silvia Milano, Nicolas Wuethrich, Todd Karhu, Philippe van Basshuysen, Bastian Steuwer (especially for suggesting the title to Chapter Two), the Arts and Humanities Research Council, the LSE Philosophy Department, and audiences at Warwick, Berkeley, Manchester, Pavia, Braga, Belfast, Boulder, and Reading. I've been lucky to meet so many great people during my time at the LSE. But most of all, I'm lucky to know Mariam Jalloq, whom I would like to thank for being so supportive and understanding over the past four years.

CONTENTS

INTRODUCTION

- I Bringing About Equality 7
- II Remarks on Methodology 9
- III Chapter Summaries 11

1 FAIR SHARES AND DEGREES OF INEQUALITY

- 1.1 Introduction 15
- 1.2 Preliminary Remarks 16
- 1.3 The Institutional Thesis 18
- 1.4 The Fair Shares Thesis 24
 - 1.4.1 The Fair Shares Thesis and Inequality Measurement 26
 - 1.4.2 The Fair Shares Thesis and Partial Compliance 28
- 1.5 The Maximisation Thesis 31
 - 1.5.1 Distributive Patterns and Egalitarian Complaints 32
 - 1.5.2 Competing Claims and Justifiability to Each 32
- 1.6 Two Problems 36
- 1.7 Conclusion 40

2 IF YOU'RE AN EGALITARIAN, WHY NOT ROB A BANK?

- 2.1 Introduction 42
- 2.2 Preliminary Remarks 45
- 2.3 Dimensions of Vigilantism 47
 - 2.3.1 Violence and Coercion 47
 - 2.3.2 Varieties of Property 47
 - 2.3.3 Direct and Indirect Losses 48
 - 2.3.4 Consolidated and Widely Dispersed Losses 49
 - 2.3.5 Individuals and Group Entities 50
 - 2.3.6 Redistributive and Expressive 51
 - 2.3.7 Summary 52
- 2.4 Degrees of Injustice 53
 - 2.4.1 Just Enough 54
 - 2.4.2 Severe Injustice 55
 - 2.4.3 Partial Injustice 56

2.5	Two Moral Obstacles to Vigilantism	57
2.5.1	The Democracy Argument	58
2.5.2	The Kantian Argument	65
2.6	Conclusion	66
3	OPTIMAL EQUALITY AND FORCED LABOUR	
3.1	Introduction	67
3.2	Equality, Pareto Optimality, Forced Labour	68
3.3	The Efficiency Dilemma	69
3.3.1	Reply to the First Horn	73
3.3.2	Reply to the Second Horn	76
3.3.3	Summary	79
3.4	The Forced Labour Objection	79
3.4.1	The First Formulation	80
3.4.2	The Second Formulation	82
3.4.3	Summary	90
3.5	Conclusion	90
3.6	Appendix: The Equality Objection	91
4	KILLING AND DUCKING: SELF-DEFENCE AND MORAL EQUIVALENCE	
4.1	Introduction	100
4.2	Threats, Bystanders, Killing, Ducking	102
4.3	The Moral Immunity Thesis	103
4.3.1	The Right of Self-Ownership	106
4.3.2	The Side-Effect Principle	109
4.3.3	The Redistribution Principle	112
4.3.4	Killing and Letting Die	114
4.3.5	Property Rights and Fair Chances	117
4.4	Two Objections	119
4.5	Conclusion	121
	BIBLIOGRAPHY	123

INTRODUCTION

Sometimes people ought to do something for the sake of equality but it would be wrong to force them to do it. Contrariwise, sometimes it would be permissible to force people to do something for the sake of equality but it is unclear whether they ought to do it without coercion.

For example, it is permissible to tax the rich in order to bring about a more egalitarian society but it is unclear whether the rich ought to voluntarily redistribute their wealth. To give another example, one might believe that people have an egalitarian obligation to make a productive contribution to society but think it is impermissible for the state to legally conscript people into certain occupations in the name of equality.

This gives rise to moral obstacles to redistributing benefits and burdens in unequal situations of different degrees and kinds. What are individuals required to do, permitted to do, and forbidden from doing for the sake of equality in these situations? In this thesis, I address four aspects of this problem.

In this introductory chapter, I explain why I think this central problem is of interest and set out the four central questions I will answer. Then, I provide some remarks on the philosophical methodology that I use throughout. Finally, I state the central claims of each chapter and briefly summarise my supporting arguments.

I Bringing About Equality

Most philosophical attention on inequality has focused on the following questions: does inequality really matter, or is it a surrogate for a concern to give priority to benefitting the worse off? Is our concern for equality about distributive fairness, or ensuring that we are socially equal? If the former, what should we equally distribute? Between whom is inequality bad? What is the relationship between equality and justice?

For the most part, I set aside these problems in order to focus on a different one which has received comparatively little attention.¹ When faced with unfair inequality, what should individuals do insofar as they are concerned to reduce it? And, given that

¹ A notable exception is G. A. Cohen's *If You're An Egalitarian, How Come You're So Rich?* (Cambridge, MA: Harvard University Press, 2002) Chapter 10.

equality is not the only thing that matters, what are individuals required to do, permitted to do, and forbidden from doing about inequality, all things considered?

I think that these questions are interesting because while we want to know what an ideally egalitarian society would look like, we also want to know what we ought to do when faced with inequality of different degrees and kinds. And as individuals, most of us cannot determine what our social institutions are like or what other people do. But we can determine what *we* will do. These considerations give rise to an abundance of interesting subsidiary questions and I can only begin to answer a handful of them. In this thesis, I provide answers to the following four.

First, when we fall short of realising equality using the coercive power of the state, do individuals have an egalitarian reason to take up the moral slack and voluntarily redistribute their own wealth? This question is of special interest to people with egalitarian convictions. Is it okay to be a rich egalitarian in an unequal society, or a “champagne socialist”? Or does a belief in egalitarianism commit you to making certain choices in an unequal society? I think that everyone, other than the worst-off person(s) in society, has an egalitarian reason to redistribute benefits until doing so further would worsen inequality, relative to the factual status quo.

Second, is it ever permissible for individuals to reduce unfair inequality by unilaterally dispossessing people of things to which they are legally, but not morally, entitled? Such flouting of property laws can take many forms, including banditry, embezzlement, fraud, looting, occupation, and trespass. I think that sometimes it is. This question is of interest because although much has been said about the justifiability of civil disobedience in response to injustice, less has been said about the justifiability of uncivil disobedience which neither is intended as a political protest, nor involves the willingness to submit to punishment.

Third, does egalitarian justice require individuals to benefit others by making a productive contribution to society? And, given that egalitarians are typically willing to coercively enforce equality through taxation, are they also committed to endorsing legally coercive job allocation in the name of equality? I think that the answers are “yes” to the former question and “no” to the latter. These questions are of interest because many attempts to change society in the name of equality have come at the price of economic efficiency or legal freedom of occupational choice, with devastating human costs.

Finally, sometimes, due to brute bad luck, harm will inevitably occur and cannot be fairly divided. What are the moral differences between the modes of agency that

individuals can employ in order to avoid being the one who suffers it? At first glance, this question may not seem closely related to the problems described above. But I think it is closely related because it concerns the moral obstacles to imposing and redistributing the burden of brute bad luck when fair distribution is not possible. I think the mode of agency used to distribute indivisible harm makes less of a difference than is often supposed.

II Remarks on Methodology

Because this thesis concerns how to respond to unfair inequality, I have been asked if I am working on “non-ideal theory”. Before proceeding further, I briefly summarise my understanding of the difference between “ideal” and “non-ideal” theory. Then I will explain where my approach fits in, insofar as it does. Finally, I clarify the philosophical methodology I will use throughout the thesis.

As I understand it, the debate about “ideal” and “non-ideal” theory is a set of methodological questions about the right or best way to do political philosophy, including, at least, the following ones.²

First, should we work out what our rights and duties are in a world in which everybody is able and willing to comply with them? Or should we deduce what our rights and duties are under the realistic assumption that some people will transgress them? We can call this the difference between deducing our rights and duties under the assumption of *full compliance* and under the assumption of *partial compliance*.

Second, does what we ought to do depend on facts about what we can do, under realistic conditions? Or is the truth about moral values and principles independent of the facts about whether or not we can put them into practice? This is sometimes called the difference between *realistic* views and *utopian* views.

Third, should we expend our philosophical efforts on trying to work out what a perfect society should look like, or should we concentrate on working out what would make the world better, in comparison to the status quo? We can call this the difference between

² See Zofia Stemplowska, “Non-ideal Theory” in *A Companion to Applied Philosophy*, ed. by Kasper Lippert-Rasmussen, Kimberley Brownlee and David Coady (Chichester: Wiley Blackwell, 2017), 284-96, and Laura Valentini, “Ideal and Non-Ideal Theory: A Conceptual Map”, *Philosophy Compass*, 7.9 (2012), 654-64.

thinking about the *ends* of an ideal society and the *transition* from the status quo to something imperfect but better.

I think these questions are of special interest for methods of moral reasoning which construct principles by asking what rules people could or would accept under hypothetical conditions in view of certain facts. When conducting such forms of reasoning, we need to know what facts to include and exclude from the hypothetical conditions. For example, should we include basic facts of economics and human behaviour? And should we go further and include facts about what is practically possible, in light of the status quo? The answers to these questions will partly depend on whether we think the primary purpose of moral and political philosophy is to work out what we ought to *do*, or to work out what we ought to *believe*.

They are also of special interest insofar as we want to determine what rules our social institutions should enforce. Perhaps we shouldn't make our institutions enforce rules which are based on false or unrealistic assumptions. Instead, perhaps we should ensure that our institutions enforce rules which take into account the fact that sometimes people won't or can't comply with them. And perhaps we should ensure that our institutions enforce rules which take into account what the factual status quo is like.

But the questions are less significant for other forms of moral reasoning which do not derive principles in this way, I think. And insofar as what we ought to do as individuals goes, these questions are less important. I think that what we ought to do as individuals depends upon a wide spectrum of considerations, not all of which can be reduced to rules that we could or would accept under hypothetical conditions. I do not try to formulate a theory which tells us how individuals should respond to all kinds of inequality. Instead, I try to illuminate the moral considerations that are at stake in different unequal situations in order to work out what this tells us about the moral requirements, permissions, and prohibitions that apply to us.

For what it's worth, I do not think that what we ought to do depends on what we can do, or that we should exclusively expend our efforts on trying to judge what would make the world better, relative to the factual status quo. On the contrary, I think moral and political philosophy is mostly about what we ought to believe, and that it is important to abstract away from the actual world in order to reason about values and principles independently what would happen if we tried to realise or implement them.

That being said, I also think that the complexity of the actual world gives rise to an interesting and largely neglected set of questions that are of both philosophical interest and practical importance. This thesis *is* concerned with "non-ideal" situations, insofar

as that is understood as thinking about what is morally at stake in the face of injustice and inequality, including the fact that some people do not, or will not, do as they ought, but *not* insofar as it is understood as employing a particular kind of reasoning about moral and political philosophy.

III Chapter Summaries

In what remains of this introduction, I summarise the main claims and supporting arguments of each chapter. The first three chapters of the thesis address questions in the domain of distributive justice. The fourth chapter steps outside the domain of distributive justice and considers a related problem in the domain of normative ethics concerning the imposition and redistribution of harm.

1 Fair Shares & Degrees of Inequality

When faced with unfair inequality, what should individuals do insofar as they are concerned to reduce it? I argue that there are at least three appealing answers to this question, and, given the diversity of reasons to object to inequality, the logic of different egalitarian views will imply different answers.

The first answer, which I call *the institutional thesis*, says that individuals in unequal societies have no egalitarian reason, or only a very trivial one, to try to reduce inequality unilaterally because isolated redistribution will only make a negligible difference, and, under realistic conditions, only institutions can ensure mass conformity with the demands of equality. I argue that the institutional thesis is supported by some varieties of instrumental egalitarianism, deontological egalitarianism, procedural egalitarianism and contractualist egalitarianism. But those views do not capture the full spectrum of egalitarian concern.

The second answer, which I call *the fair shares thesis*, says that individuals should give away everything above their “fair share”, where one’s fair share is defined as the amount of goods that one would have in a counterfactual world in which that which is captured by the metric of equality is equally distributed. I argue that the fair shares thesis is supported by left-libertarianism and luck egalitarianism combined with rule consequentialism. It is also supported by combining luck egalitarianism with some implausible measures of inequality. But, I argue, the view depends on the idea that the amount of that to which the metric of equality applies is fixed. We should not make that assumption.

The third answer, which I call *the maximisation thesis*, says that everyone, other than the very worst-off person in society, has an egalitarian reason to redistribute benefits to those who are worse off until doing so any further would worsen inequality, relative to the factual status quo. An implication of this is that there is still an egalitarian reason to redistribute benefits even if this makes a person worse off than they would be in the counterfactual world in which all existing goods were equally distributed. I argue that this view is supported by luck egalitarianism and the competing claims view, at least. All told, I believe the maximisation thesis to be the most thoroughgoing answer.

2 If You're An Egalitarian, Why Not Rob a Bank?

It is sometimes permissible to reduce unfair inequality by dispossessing people of things that they are legally, but not morally, entitled to, without the informed consent of the legal owner and without providing compensation. I call this egalitarian vigilantism. I start by considering the dimensions across which acts of egalitarian vigilantism are easier or harder to justify, independently of the degree of injustice in society as a whole. I conclude that the most justifiable kinds of vigilantism are those which dispossess group entities of money, food, and disused land and buildings, absent special attachment claims, where any losses suffered by individuals are indirect and widely dispersed.

I then turn to consider how the extent to which property laws are unjust affects the justifiability of vigilantism. Unsurprisingly, I conclude that vigilantism is almost impossible to justify when property laws are just or just enough, and much easier to justify when they are severely unjust. The most interesting and difficult questions about vigilantism emerge when property laws, and societies as a whole, are *partially* unjust.

I claim that there are three additional moral presumptions against redistributive vigilantism in partially unjust situations. First, law-abiding citizens have a claim-right that others respect legal property entitlements that were allocated in a reasonably democratic way. Second, vigilantism is a disproportionate response to injustice when democratic or other legal courses of action are available. Third, widespread acts of vigilantism would often make things worse rather than better.

I argue that, taken together, these arguments do not rule out the expropriation of property from groups or people who culpably perpetrate injustice, when democratic or other legal forms of recourse are highly unlikely to be effective, provided that such acts do not become too widespread.

3 Optimal Equality & Forced Labour

G. A. Cohen claimed that we could bring about an egalitarian society without sacrificing economic efficiency or legal freedom of occupational choice if most people, motivated by an egalitarian ethos, freely and efficiently allocated their labour in exchange for an egalitarian wage. In the third chapter, I defend this claim from two counterarguments.

The first counterargument, *the efficiency dilemma*, says that on the one hand, a merely egalitarian ethos would not suffice to induce people to efficiently allocate their labour, and, on the other hand, a productive ethos, which induced people to efficiently allocate their labour, is beyond what justice could plausibly require.

In reply, I argue that egalitarian justice requires the realisation of a Pareto optimal level of equality and, accordingly, the free realisation of egalitarian justice requires a productive ethos. I argue that a productive ethos is not an implausibly demanding requirement of justice because when the currency of justice is access to advantage, it is, generally speaking, in the rational self-interest of each person to allocate their labour efficiently.

The second counterargument, *the forced labour objection*, says that Cohen cannot consistently advocate the use of coercive taxation for the sake of equality but oppose legally coercive job allocation for the sake of equality. In reply, I argue that we can distinguish between the coercive enforcement of a just distribution of income and the enforcement of a just distribution of labour in a number of different ways.

4 Killing & Ducking: Self-Defence & Moral Equivalence

I start by contrasting two cases in which a person who is endangered by an innocent threat can defend themselves using different modes of agency in each case: by killing the innocent threat in the first case and by ducking the threat in the second case. According to *the moral immunity thesis*, it is wrong to kill a person who innocently threatens your life, on the grounds that doing so is morally indistinguishable from killing an innocent bystander when defending your life. By contrast, proponents of the account plausibly hold that you may permissibly evade, or “duck”, harm when doing so will foreseeably lead to the death of an innocent person.

I then set out my central claim: if killing an innocent threat is morally equivalent, in terms of permissibility, to killing an innocent bystander, then both are equivalent to ducking a threat. As a result, if the moral immunity thesis is correct, then ducking the

threat is impermissible. Since this is highly counterintuitive, I suggest that we should instead reject (or revise) the moral immunity thesis.

I defend this claim in two stages. First, I offer an intermediate case which falls between killing and ducking. I claim that we can proceed, by a transitive relation of moral equivalence, from the view that killing an innocent threat is wrong to the claim that ducking is wrong. Second, I consider five ways in which a proponent of the moral immunity thesis might try to resist my claim: by appealing to the right of self-ownership, the side-effect principle, the redistribution principle, the distinction between killing and letting die, and the moral significance of fair chances and property rights. Using counterexamples, I argue that all five attempts are unsuccessful.

1 FAIR SHARES AND DEGREES OF INEQUALITY

1.1 Introduction

The distribution of benefits and burdens in the actual world is unjustly unequal and will continue to be unjustly unequal for the foreseeable future. For the fortunate, life is better than it would be if our unequal world were transformed into a justly egalitarian one. By contrast, many people are worse off than they would be if inequality were reduced through redistribution, and some of us find it morally repugnant that people enjoy extravagant riches in a world where others struggle to make ends meet.

Some of us believe that we ought to reduce inequality by redistributing benefits and burdens using the coercive power of the state. But when, as is actually the case, we fall short of realising equality through redistributive public policies, it is unclear whether and how individuals ought to take up the moral slack and do what they can to reduce inequality on their own. What does egalitarianism demand of individuals in unequal societies? In this chapter, I consider three answers to this question and examine which of the answers are supported by different varieties of egalitarianism, widely understood.

The first answer, which I will call *the institutional thesis*, is that individuals in unequal societies have either *no* egalitarian reason to unilaterally redistribute benefits, or only a very *trivial* one, because isolated redistribution will, at best, only negligibly reduce inequality, and, under realistic conditions, only institutions can ensure that everybody contributes to widespread redistribution. I argue that this view is supported by some instrumental, deontological, procedural, and contractualist varieties of egalitarianism, but that these views do not capture the full spectrum of egalitarian concern.

I will call the second answer *the fair shares thesis*. This is the view that individuals *do* have an egalitarian reason to reduce inequality, but this only extends to redistributing or foregoing goods which are in excess of one's fair share, where one's fair share is defined as the amount of goods that one *would have* in a counterfactual world in which the metric of equality is equally distributed. The fair shares thesis has intuitive appeal and enjoys argumentative precedent when it comes to our moral duties to help the needy. However, I argue that this view is unsupported by most strands of egalitarian thought, contrary to what has been assumed by some.

I call the third answer *the maximisation thesis*. This is the view that everyone, other than the very worst-off person in society, has an egalitarian reason to redistribute benefits to those who are worse off until further doing so would worsen inequality, *relative to the factual status quo*. The maximisation thesis appears to lack common sense support; it does not seem widely believed that everybody, including people of average or below average means, has an egalitarian reason to redistribute benefits to those who are worse off. Yet somewhat surprisingly, this answer is implied, I believe, by at least two of our best egalitarian views: luck egalitarianism and the competing claims view.

The chapter is divided into four main parts. In what remains of the first part, I provide some preliminary remarks, explain how the problem differs from some other questions, and then explain why it matters. In the second part, I set out the institutional thesis, and explain which varieties of egalitarianism support it, and point out that these views do not capture everything that egalitarians care about. In the third part, I set out the fair shares thesis, explain which varieties of egalitarianism support it, and then point out that, contrary to what some have assumed, the view is not supported by most varieties of egalitarianism. In the fourth part, I set out the maximisation thesis and defend the claim that the maximisation thesis gains support from combining some of our best egalitarian views with some of our best ways of measuring the badness of inequality. Following that, I address some problems for the maximisation thesis before concluding.

1.2 Preliminary Remarks

I focus on the question of what individuals in unequal societies have *egalitarian* reason to do, or what people should do insofar as they want to act upon egalitarian beliefs. I set aside the question of what people in unequal societies are morally *required* to do, and the question of what people have most reason to do, all things considered. Other reasons, values, constraints, and agent-centred options will sometimes permit, or forbid people to do what is best from an egalitarian perspective.¹

That being said, many of us take equality to be a conception, or aspect, of distributive justice, and reasons of justice are often taken to carry special weight, even if they may sometimes be overridden by other concerns. So, if individuals in unequal societies do

¹I set aside the difficult but important question of how individuals should decide what to do in light of empirical uncertainty about the distribution of benefits and burdens in society and moral uncertainty about which choices people are morally responsible for, insofar as that affects their distributive shares. Individuals in the actual world would need a decision procedure to help them decide how to best carry out egalitarian redistribution, given these epistemic and normative obstacles.

have an egalitarian reason to redistribute or forego goods, this is not to be taken lightly and there will sometimes be a decisive reason to act upon it when the costs of doing so are not too onerous and no other good reasons tell against doing so.

Surprisingly, this question has not received much attention from philosophers, despite its practical significance, and despite its significance for a range of further questions. For example, the answer has a bearing on whether it is permissible to be extremely rich in an unequal society and whether it is hypocritical to be a rich egalitarian, although complete answers to such questions will require further premises about the correct metric of equality and what one has most reason to do, all things considered.²

Now, individual egalitarian acts will typically take the form of redistributing benefits from oneself to the worse off. By “redistribution”, I mean any actions in which a person imposes losses in the metric of equality on themselves in order to bestow benefits in the metric of equality on another person or persons.³ However, egalitarian acts could also take the form of *foregoing* benefits which will worsen inequality. The reason in favour of foregoing benefits may sometimes be stronger than the reason in favour of redistributing benefits because it is plausibly a greater sacrifice to give up what one has than it is to forego things one doesn't have (although I will consider a problem with this claim towards the end of the chapter). I will set aside acts which reduce inequality by imposing losses on *others*.

This question is distinct from some other important questions which have been more widely addressed. For instance, sometimes the culpable and innocent beneficiaries of injustice ought to redistribute things in order to rectify wrongdoing, including wrongdoing by omission.⁴ But on *some* views, an unfairly unequal distribution can arise without anybody acting wrongly, and without anybody's rights being violated. So, we might sometimes have an egalitarian reason to redistribute benefits to others even when

² The question of whether it is hypocritical to be a rich egalitarian is examined in Cohen, *If You're an Egalitarian, How Come You're So Rich?*, Chapter Ten, and Saul Smilansky “On Practicing What We Preach”, *American Philosophical Quarterly*, 31.1 (1994), 73-79.

³ Apart from where otherwise stated, I do not take a stand on the metric of equality, whether it is welfare, resources, opportunity sets, income, primary goods, capabilities, and so on. Since some of those currencies cannot be directly transferred, any egalitarian reason to redistribute benefits will typically require the redistribution of the sources of those goods. I also set aside the question of the scope of egalitarianism, whether that is the set of individuals within a nation state or the set of all human beings.

⁴ For a defence of this view, see Daniel Butt, “On Benefitting From Injustice”, *Canadian Journal of Philosophy*, 37.1 (2007), 129-52.

rectification is not required because nobody has acted wrongly and nobody's rights have been violated.

Moreover, we sometimes have reasons to aid the needy (those who are *badly off*, in absolute terms) when we have no egalitarian reason to do so. We would, for example, have reasons to aid everyone even when everyone, including the aiders, is equally needy. And we sometimes have an egalitarian reason to redistribute benefits even when we have no humanitarian reason to do so; for example, when needs have been satisfied, hardship has been alleviated and everyone's life is sufficiently good.⁵ That being said, duties to rectify wrongdoing and to aid the needy may sometimes overlap with reasons to reduce inequality such that our reasons to redistribute benefits are overdetermined.

Aside from being practically significant, this problem is significant for normative ethics because the soundness of some egalitarian views has been questioned on the grounds that the views do not directly tell us anything about how we ought to act in the actual world. These problems are sometimes taken to be shortcomings which undermine the plausibility of those accounts. But that conclusion is premature, or so I shall suggest.

1.3 The Institutional Thesis

The institutional thesis, as I shall call it, consists of three claims: (1) there is *no* egalitarian reason for individuals to unilaterally try to reduce inequality, or only a very trivial one, because (2) unlike widespread redistribution, isolated egalitarian acts, at most, only negligibly reduce inequality and, (3) under realistic conditions, only coercive institutions can ensure mass conformity with the demands of equality.⁶

⁵ The question of what we are morally required to do to help the needy is widely addressed. For examples, see Garrett Cullity, *The Moral Demands of Affluence* (Oxford: Clarendon Press, 2004), and Peter Singer, "Famine, Affluence and Morality", *Philosophy and Public Affairs*, 1.1 (1972), 229-43.

Arguments for a moral requirement to help the needy often proceed by parity of reasoning from cases in which one ought to save a life at little cost to oneself. No comparable argument can be made on the part of equality. Helping people who are badly off is morally urgent. By contrast, sometimes reducing inequality, important though it may be, is not always morally urgent. For example, unfair inequality between two people who are very well off, in absolute terms, seems morally significant but it does not seem morally urgent to redress it.

⁶ By "negligible", I mean "unimportant", rather than numerically negligible. This distinguishes the institutional thesis from what we can call the "drop in the ocean" claim, which says: "Yes, I could reduce inequality through my own choices, but by doing so, I would only make a tiny difference in the grand scheme of things."

To illustrate, suppose, for the sake of argument, that the metric of equality is wealth, and suppose further that the timespan of egalitarian concern is whole lives. If the institutional thesis is correct, then we would have, at most, a very trivial egalitarian reason, over the course of our lives in an unequal society, to redistribute wealth to people who are worse off than us (though we may have other reasons to do so).

This is practically significant because if egalitarian reasons in favour of independent redistribution are very trivial, then these reasons would rarely, if ever, cross the threshold of moral requirement. I take claim (3) to be true, so I will explore which varieties of egalitarianism make the conjunction of claims (1) and (2) true.

Now, on distributive egalitarian views, the institutional thesis is false. For example, on views which hold that a state of affairs containing distributive inequality is intrinsically bad, in one respect, when and because some are unfairly or arbitrarily worse off than others,⁷ it is clearly not generally true that only widespread egalitarian action would reduce inequality in a significant way, since there are many circumstances in which independent redistribution would do a great deal to reduce inequality.

However, people object to inequality on a variety of grounds and the different objects of egalitarian concern provide different answers to the question of whether reducing inequality in small degrees makes a moral difference or whether the badness or injustice of inequality can only be removed by widespread action. For example, one might believe that equality would be good for society as a whole, but not believe that reducing inequality in small ways between individuals would do anything to promote that ideal. In fact, I think the institutional thesis is supported by at least four egalitarian views.

First, while people think distributive inequality is intrinsically bad, others think it is only *instrumentally* bad insofar as it leads to bad effects like divided communities, hierarchies of social status and power, or morally troubling attitudes which undermine solidarity. For example, Martin O'Neill has proposed a variety of egalitarianism which he calls *non-intrinsic egalitarianism*.⁸ He claims that distributive inequality is bad, regardless of the effect of inequality on individual well-being, but this badness does not reside in the

⁷ G. A. Cohen, *Rescuing Justice and Equality*, (Cambridge, MA: Harvard University Press, 2008), Shlomi Segall, *Why Inequality Matters: Luck Egalitarianism, Its Meaning and Value* (Cambridge: Cambridge University Press, 2016), and Larry Temkin, *Inequality* (Oxford: Oxford University Press, 2003).

⁸ In Martin O'Neill, "Constructing a Contractualist Egalitarianism: Equality After Scanlon", *The Journal of Moral Philosophy*, 10 (2013), 429-61, and "What Should Egalitarians Believe?" *Philosophy and Public Affairs*, 36.2 (2008), 119-56.

relational property of some being worse off than others in some metric of equality. Rather, distributive inequality is bad when and because it causes these kinds of social and political inequality.

On this view, we should redistribute from better to worse off *as a means* to remove these bad effects. So *if* redistributing or foregoing benefits would serve to reduce social or political inequality, then there would be an egalitarian reason for individuals to do so. But typically, isolated acts of redistribution will do nothing to remedy social and political inequality, since that is a systemic feature of societies. Voluntary redistribution may even be counterproductive when and because it humiliates or undermines the dignity of the recipients.

Although it certainly makes sense to talk about social and political equality in degrees (we can say that one society is more socially and politically equal than another), individual contributions to reducing social and political equality will typically take the form of holding certain attitudes towards others or refraining from treating people in certain ways which undermine egalitarian relationships.

A worry about this view is that it becomes difficult to see why we should favour widespread redistribution of things from the better to worse off in degrees, where doing so will not also reduce social or political inequality. Of course, these egalitarians may favour redistribution for non-egalitarian reasons as well (to alleviate hardship for example). But on this view it seems that unless redistribution reduces social inequality, there is little egalitarian reason to favour it.

Second, the institutional thesis gains some plausibility from deontological egalitarian views which object not to distributive inequality itself, but rather to the way it was brought about. For example, some believe that inequality matters when and because it results from *treating people unequally*, and we should aim to reduce distributive inequality when and because that is what treating people equally requires.

To give one example of such a view, A.J. Julius has argued, roughly speaking, that distributive equality is morally relevant only insofar as we need to justify treating each other fairly through our shared institutions. He says that we have an egalitarian reason to refrain from inducing each other to act in ways which unfairly benefit us, including promoting or perpetuating institutions which unfairly advance our own interests.⁹ On this view, we should aim for distributive equality in order to refrain from treating each

⁹ A. J. Julius, "Basic Structure and the Value of Equality", *Philosophy and Public Affairs*, 31.4 (2003), 321-55 (p.347).

other unfairly, and there is no egalitarian reason to redistribute benefits when the failure to do so does not constitute unfair treatment.

This view does, however, hold that we have egalitarian reasons to refrain from promoting or supporting coercive institutions which work to our unfair advantage. For example, we have an egalitarian reason to refrain from lobbying for changes to taxation policies which unfairly benefit us. So, while we may have no egalitarian reason to redistribute things we already possess, *if* we are responsible for distributing some new benefits and burdens, then we should distribute them equally.

A third way in which the institutional thesis could be true is if equality only requires ensuring the metric of equality is distributed through *fair procedures or rules*. The following metaphor captures such a view. “I deplore the fact that the cards are stacked against some people in life, and in favour of others. We should establish rules and procedures which deal everyone a fair hand and even out everyone’s life prospects. If everyone had a fair hand and played by fair rules, it wouldn’t matter morally if it led to unequal outcomes.” On this view, trying to promote a more equal outcome will not reduce inequality in a morally relevant respect. Instead, we should aim to make our procedures and rules fairer.

A deontological formulation of this claim could say that we should aim for fair rules and procedures because that’s what treating people equally requires. This coheres well with the hypothesis that we ought to reduce inequality through the coercive power of the state but have no reason to do so in our own lives. When acting together through the state, we ought to treat each other impartially by playing by the same rules, but people need not treat each other impartially in their own lives.¹⁰

We can distinguish between two ways of understanding procedural views of this variety. On one formulation, the unfairness of procedures is grounded in the fact that they produced an unfair outcome. One might not object to the same procedure if it produced a fair outcome accidentally, for example. On this view, we should aim to change our

¹⁰ Ronald Dworkin locates the justice of equality in the way the state treats its citizens, in *Sovereign Virtue* (Cambridge, MA: Harvard University Press, 2000). One might grant that that reducing inequality through redistributive personal choices would make an outcome better, but not more just. One might maintain that the kind of inequality that is brought about through social institutions is a distinctive kind of injustice, which the state has reasons to reduce through its tax and transfer policies, but individuals have no equivalent reason to reduce through their voluntary choices. The distinction is significant because we typically have stronger reasons to remove injustice than we do to remove (mere) disvalue.

institutional procedures *because* they tend to produce maldistributions. But if we take that view, then it doesn't matter whether a distribution is brought about by a change in procedures or voluntary redistributive acts.

On a second, more distinctive, formulation, the badness of unequal outcomes is grounded in the unfairness of procedures; inequality is bad when and because it was produced by unfair procedures. On this view, there is no fact of the matter about what a fair distribution would look like in the absence of fair procedures and there is no independent standard for a fair outcome.¹¹

All that being said, we can ask what *would* have happened had a fair procedure taken place. Even if there is no *unique* fact of the matter about the distribution that would have resulted from a fair procedure, there is a set of possible distributions that might have been produced by fair procedures. Of those possible distributions, we could either aim for the one which was most likely to occur, or minimise expected injustice or minimise the maximum possible injustice done to any one individual by bringing about a more equal, rather than less equal, distribution.¹²

Finally, one might defend the institutional thesis by offering a *contractualist* defence of the claim that equality is concerned with fair rules and procedures. On a contractualist formulation of this claim, we could say that we should aim for fair rules and procedures because that's what we would agree to in a fair hypothetical situation in which our individual interests are set aside; when distributing benefits and burdens across individuals, we do so in accordance with rules and principles that can be justified to each.

Consider John Rawls's claim that the basic structure of society is the primary (though not only) agent of justice within his free-standing contractualist theory of justice.¹³ For several reasons which are specific to Rawls's contractualist theory, a distribution is fully just, if, and only if, and because it is brought about by full compliance to principles, rules, and institutions that would be in everybody's rational self-interest to agree to

¹¹ These two ways of formulating a deontological concern with how inequality is produced are distinguished in Derek Parfit, "Equality or Priority?" in *The Ideal of Equality*, eds. Matthew Clayton and Andrew Williams (Basingstoke: Palgrave, 2002), p.89.

¹² David Miller defends this justification for bringing about equality under uncertainty in "Equality and Justice", *Ratio*, 10.3 (1997), 222-37 (pp.227-228), although he argues against the claim that distributive equality under certainty is morally significant.

¹³ John Rawls, *A Theory of Justice: Revised Edition* (Cambridge, MA: Harvard University Press, 1999).

behind a non-probabilistic veil of ignorance, on the condition that everybody else would comply as well.

It would not be fair, nor would it be in everybody's common rational interest, to agree to principles if they could not be assured that others would reciprocate in terms of compliance with those principles. That is one reason why the basic structure of society is the primary agent of justice in Rawls's theory; the coercive power of the state can ensure that most people will comply with the principles of justice and ensure the fairness and rationality of agreeing to adhere to those principles behind the veil of ignorance. On this view, one might think that, as a matter of *distributive* justice, individuals are duty-bound to comply with legally enforced obligations to institutions that *are* just, to refrain from supporting and sustaining institutions that realise unjust inequality, and perhaps play a part in the reform or abolition of those institutions, but no more.¹⁴

Whether Rawls is right about the significance of the basic structure of society within the scope of his own theory has been widely discussed.¹⁵ But *that* discussion concerns the question of whether individuals can do more to promote justice through their personal choices when social institutions are fully just (or, put another way, whether a society can be less than fully just even when the basic structure of society is fully just).

The present question is of interest regardless of one's stance on the significance of the basic structure within Rawls's theory of justice which is restricted in scope to the choice of principles for institutions in sufficiently developed societies. The claim that that individuals could do nothing more to promote equality through their personal choices in a society with a just basic structure differs from, and is consistent with, the claim that individuals faced with unjust inequality can *reduce inequality* through voluntary egalitarian acts, even if individual action alone cannot bring about a fully just distribution.

In order to deny that individuals faced with inequality have an egalitarian reason to carry out direct redistribution, one must affirm a different view. One must hold that

¹⁴ Rawls does of course also say that individuals can have further natural duties which fall outside the sphere of interaction with the basic structure of society. *A Theory of Justice* (pp.98-101).

¹⁵ For example, see Michael G. Titelbaum, "What Would a Rawlsian Ethos of Justice Look Like?", *Philosophy and Public Affairs*, 36.3 (2008), 289-322, David Estlund, "Liberalism, Equality, and Fraternity in Cohen's Critique of Rawls", *The Journal of Political Philosophy*, 6.1 (2008), 99-112, Kok Chor Tan, "Justice and Personal Pursuits", *Journal of Philosophy*, 101.7 (2004), 331-62, and Liam Murphy, "Institutions and the Demands of Justice", *Philosophy and Public Affairs*, 27.4 (1999), 251-91.

only changes to social institutions and procedures can reduce unjust inequality and individuals are *incapable* of reducing distributive unfair inequality through redistributive personal choices. That is implausible, and it is not a claim that Rawls explicitly defends.¹⁶

In summary, the institutional thesis is supported by at least four egalitarian views. Whether these views should count as genuinely egalitarian ones, and whether they are plausible in their own right, are further questions. I think that some are plausible, even if different views ultimately turn out to be addressing different aspects or kinds of inequality. And even if social and political equality, treating people equally, fair procedures, and governing our interactions on egalitarian terms which can be justified to each are important aspects of egalitarian concern, they are not, in my view, the whole truth about egalitarian justice. In any case, these views certainly do not exhaust the spectrum of egalitarian views, and I now turn to examine a different answer to the question with which I began: the fair shares thesis.

1.4 The Fair Shares Thesis

The fair shares thesis holds that individuals faced with inequality have an egalitarian reason to redistribute or forego benefits, regardless of what others do, but only goods which are in excess of their *fair share*. For example, suppose, for the sake of argument, that the metric of equality is resources and suppose further that the timespan of egalitarian concern is whole lives. If the fair shares thesis is correct, then we have an egalitarian reason to redistribute resources in excess of our fair share to those with less than their fair share, over the course of our lives.

But a problem with this view is that there are several ways in which one might define what one's fair share is and no immediate way of determining which is the most plausible. For example, each person has a fair share of the metric of equality relative to the factual status quo. That is, by the amount of goods one would have in the counterfactual world in which *existing* benefits and burdens were equally distributed.

But this assumes that amount of the metric of equality is fixed. We should not make this assumption. For counterfactual worlds may contain different amounts of the metric of

¹⁶ For an argument in support of this conclusion, see Thomas Porter, "The Division of Moral Labour and the Basic Structure Restriction", *Politics, Philosophy and Economics*, 8 (2009), 173-99.

equality.¹⁷ For example, a counterfactually egalitarian society may contain fewer goods due to less being produced. Alternatively, it may contain more goods, due to talents being more efficiently harvested and labour being more efficiently distributed.¹⁸ Now, we shouldn't dismiss the fair shares thesis for this reason alone because we might be able to determine a plausible baseline for fair shares from the nature of different egalitarian views. There are at least four egalitarian rationales for the fair shares thesis, some more plausible than others.

First, according to the family of views called left-libertarianism, which combine a right of self-ownership with a luck egalitarian principle of distribution, each individual is rightfully entitled to an egalitarian share of the world's resources.¹⁹ Egalitarianism based on natural rights and entitlements could straightforwardly include a moral duty to redistribute benefits as part of a theory of rectification, since, on these views, each person has a right to a fair share of existing resources, and inequality violates the natural rights of those who are deprived of their rightful entitlements.²⁰ Fair shares, on this view, could be determined relative to the facts concerning the amount of resources in the world and each person's ability to transform resources into welfare.

Second, we might adopt the fair shares thesis if we were to combine a luck egalitarian view about the comparative unfairness of outcomes with a rule consequentialist account of the duty to follow rules which, if internalised and acted upon by everyone, would make things go best.²¹ On this view, fair shares would plausibly be determined by the

¹⁷ The assumption is relatively plausible if the metric of equality is natural resources or natural resources are the only source of the proper metric of equality. But even then, natural resources can be depleted and replenished.

¹⁸ A further peculiarity is that teleological views like luck egalitarianism are sometimes taken to be unlimited in scope across time and space. So a teleological fair shares view, without further auxiliary premises to limit the scope of egalitarianism, suggest the possibility that one's fair share is relative to the amount of the metric of equality that will exist in total throughout time and space.

¹⁹ See, for example, Michael Otsuka, *Libertarianism Without Inequality* (Oxford: Clarendon Press, 2003). An egalitarian share can be defined in several ways, including an equal distribution of non-human resources, or an equally advantageous distribution, that is, a distribution of resources which realises equal opportunity for welfare, which is what Otsuka defends.

²⁰ Note, however, that according to left-libertarian views, possessing an inegalitarian share of the world's resources would be a form of benefitting from injustice. So a duty to redistribute resources may be better described as a form of rectifying wrongdoing than an egalitarian reason to redistribute.

²¹ See, for example, Brad Hooker, *Ideal Code, Real World: A Rule-Consequentialist Theory of Morality* (Oxford: Clarendon Press, 2000). See pp.44-51 for a discussion of equality, fairness and justice.

shares one would have in a counterfactual world in which everybody followed rules which would bring about the optimal possible outcome (this would typically be a higher fair share than one's fair share relative to the factual status quo).

Both left-libertarianism and luck egalitarianism combined with rule consequentialism plausibly support the fair shares thesis. But some have also suggested luck egalitarianism might support the fair shares thesis *without* a rule consequentialist criterion of rightness. For example, Peter Vallentyne has proposed that a luck egalitarian account of the redistributive duties that we owe to each other could hold that

an agent with more than her fair share owes (at least a *pro tanto*) duty, to each of those with less than his fair share [...] to make the distribution as good as possible without the agent leaving herself less than her fair share.²²

Similarly, G.A. Cohen asked why rich egalitarians faced with unjust inequality do not “pursue equality by donating the extra that [they] would lack in a just society to poor people” and assumes that “no one is obliged to sacrifice so much that she drops to a level *worse* than she would be at in an egalitarian society”.²³

I suspect that some people are attracted to this view because it is intuitive that, from an egalitarian perspective, people are owed equal shares of that which is already unequally distributed. This view also seems to naturally cohere with the belief that an ideally egalitarian government ought to redistribute everything above their fair share through taxation. This view might be explained in two ways.

1.4.1 The Fair Shares Thesis and Inequality Measurement

One could try to defend the fair shares thesis by combining luck egalitarianism with some ways of measuring the badness of inequality. If the amount of inequality in a distribution is determined wholly by (1) the number of individuals who are worse off than average (relative to the amount of the metric of equality in the factual status quo), (2) the deviation from a state in which all existing goods are equally distributed, or (3) the aggregation of the differences between people's counterfactual fair share and their

²² Peter Vallentyne, “Justice, Interpersonal Morality, and Luck Egalitarianism” in *Distributive Justice and Access to Advantage* ed. by Alexander Kaufman (Cambridge: Cambridge University Press, 2014), p.9.

²³ Cohen, *If You're an Egalitarian, How Come You're So Rich?*, p.161 and p.176.

actual shares, then a transfer which makes someone worse off than their fair share might *worsen* inequality.

However, we should be sceptical about whether these measures of the badness of inequality are plausible. All three contradict the Pigou-Dalton principle, which states that all even, non-rank-switching transfers from a better off individual to a worse off individual, which leave others unaffected in an absolute sense, reduce inequality.²⁴ Many transfers which satisfy the Pigou-Dalton principle would not reduce inequality on each of these measures because they would *not* reduce the number of people who are worse off than average or alter the aggregate deviation from a state in which the metric of equality is equally distributed.²⁵

For example, redistribution from the best-off person to someone who is just above the average, or redistribution from a person who is just below the average to the worst-off person would neither alter the amount of people who are worse off than average nor the deviation from a state in which existing goods are equally distributed, but such transfers *would* satisfy the Pigou-Dalton principle.

The first two measures of inequality also counterintuitively imply that it is better, from an egalitarian perspective, to raise someone who is just below their fair share up to their fair share rather than benefit someone who is *much* worse off whom we cannot raise to the level of their fair share, because only the *number* of people with their fair share is what determines the amount of inequality in society as a whole.

None of our better ways of comparing unequal outcomes imply that individuals have no reason to redistribute benefits any further once they reach or fall below their fair share. More plausible ways of comparing unequal outcomes turn out to support the maximisation thesis (I will postpone defending this point for now).

²⁴ See Hugh Dalton, "The Measurement of the Inequality of Incomes", *The Economic Journal*, 30 (1920), 348-61, revisited in Anthony B. Atkinson and Andrea Bandolini, "Unveiling the Ethics Behind Inequality Measurement", *The Economic Journal*, 125 (2015), 209-34.

²⁵ A similar discussion, in relation to distributive justice and the Pigou-Dalton Principle can be found in Matthew Adler, "The Pigou Dalton Principle and the Structure of Distributive Justice" (Working Paper, 2013).

1.4.2 The Fair Shares Thesis and Partial Compliance

Luck egalitarians might be tempted to adopt the fair shares thesis by drawing on arguments about duties of beneficence (requirements to promote the well-being of others) and duties to avert harm in situations of *partial compliance*. Situations of partial compliance are situations in which members of a group of people have a moral duty to act to bring about a good outcome or avoid a bad outcome, where responsibility is fairly divided between each member of the group, and at least one person does not carry out their fair share of this duty. I will set out two ways of thinking about such situations in more detail and then explain why I don't think they work in the case of inequality.

In *Moral Demands in Non-Ideal Theory*, Liam Murphy argues, against act consequentialism, that the moral requirements of beneficence are constrained by fairness.²⁶ He points out that when some people do not bring about the best outcome, an individual acting on the requirements of beneficence will be doing more than their fair share, since they would be making sacrifices that they would not be asked to make if everybody did as they ought. He concludes that each person is only required to make the same sacrifices for the sake of promoting the well-being of others that they *would* make in the counterfactual world in which everybody did as they ought in order to bring about an optimal outcome.

Similarly, in "Taking Up the Slack", David Miller considers the requirements of justice in situations in which several agents share a duty to avert an unjust harm at moderate cost and some agents do not do not carry out their fair share of that responsibility.²⁷ Examples of these kinds of situations in the actual world mentioned by Miller include the alleviation of world poverty, the conservation of natural resources, and the aversion of environmental catastrophe. Miller concludes that in situations of partial compliance, individuals are generally only duty-bound, as a matter of justice, to do what they would be required to do under full compliance because that that fully serves *their* responsibility; people are not responsible, as a matter of justice, for what others do or

²⁶ Liam Murphy, *Moral Demands in Non-Ideal Theory* (Oxford: Oxford University Press, 2000).

²⁷ David Miller, "Taking Up The Slack? Responsibility and Justice in Situations of Partial Compliance" in *Responsibility and Distributive Justice* ed. by in Carl Knight and Zofia Stemplowska (Oxford: Oxford University Press, 2011). Miller discusses what he calls duties of justice, but these are duties of justice in the sense of the duties that we owe to each other as a matter of right, rather than distributive justice.

not do. (Miller grants that there might be a humanitarian obligation to take up the slack in morally urgent cases, for example, when the lives of innocents are at stake.)²⁸

However, there are at least three respects in which reducing inequality differs from duties of beneficence and duties to avert harm, and these differences provide good grounds to deny that these arguments support *the fair shares view* when it comes to reducing inequality.

First, while the objection to taking up the slack in cases of beneficence or averting harm is made on the grounds of fairness, fairness is an aspect of the very thing we are trying to promote in the case of reducing inequality. Fairness seems to pull us in different directions. Which situation is more unfair: asking someone to take up the slack in order to ameliorate the unfairness or leaving some worse off than they could be in the name of fairly distributing the responsibility for reducing inequality? In many cases, the latter answer will be more unfair. Consider cases in which the costs of doing more than one's fair share are modest and the benefits to the worse off are large.²⁹ More plausibly, the requirements of fair redistribution are set by ratio of prospective costs to the distributor and benefits to recipients.

Second, in the case of averting harm or promoting the well-being of others, there is typically one optimal outcome to aim for, even if we can move closer to the outcome in degrees. For example, when it comes to harm aversion, the optimal outcome is one in which harm is completely averted, although in some cases we may be able to reduce harm in degrees (for example, if there are multiple lives at stake). And when it comes to promoting the well-being of others, there is one optimal state (or multiple optimal states which are equally good), although we will typically be able to promote the well-being of others in degrees. In these cases, a person's fair contribution can be identified with reference to what they would have to do if everybody fairly contributed to the realisation of the optimal outcome. (There seems to be no reason for an agent to do their fair share in situations of partial compliance in which a threshold needs to be passed, which could not be passed by an agent doing only their fair share).

By contrast, there are many different outcomes which could reduce inequality. When it comes to teleological views, it is mistaken to attach moral significance to the *absolute*

²⁸ Miller also observes that “[e]veryone can agree that in these cases compliers have a strong reason to take up the slack” in “Taking Up The Slack”, p.242.

²⁹ Keith Horton raises this objection to Murphy in “Fairness and Fair Shares”, *Utilitas*, 23.1 (2011), 88-93 (p.92).

level of advantage people would have in any counterfactually egalitarian world since equality is concerned with how individuals fare relative to others, rather than how they fare absolutely. Inequality can be reduced in degrees, or by increasing or decreasing the amount of goods to be distributed (by levelling up or down), in addition to redistributing the goods that already exist, and so there is no uniquely egalitarian counterfactual alternative.

One might say that there is something morally special about the “closest state of equality” with the implication being that is the state in which existing goods are equally distributed. But one could also think of the closest egalitarian state as being the one in which the least amount of the metric of equality has to be added or removed from the distribution in order to realise equality.³⁰

The third difference is related to the second. The idea that one ought to do one’s “fair share” derives its force from the idea that some moral obligations are *collective* or *shared*, in a way which is irreducible to individual moral obligations.³¹ There are situations where it is clear that people *ought* to act together to avert harm or promote a good outcome. This seems plausible in cases in which the harm to be averted, or the good outcome to be brought about, could not be achieved by individuals acting alone. But it is not clear that inequality, as a feature of states of affairs, is something that we are collectively responsible for removing.

To say that inequality is a bad feature of an outcome is to imply that there is reason to reduce or remove the inequality, but it does not imply, by itself, that people are collectively responsible for doing so. Whether one ought to reduce inequality will typically depend on the balance of reasons for and against doing so. To be sure, one *could* think of the realisation of equality as a collective responsibility. For instance, one might think that inequality is bad only when it *harms* people by making some worse off than they could otherwise be. Then the duty to reduce inequality becomes an example of

³⁰ Larry Temkin, “Exploring the Roots of Egalitarian Concerns”, *Theoria*, 69 (2003), 125-51 (p.138).

³¹ William Edmundson has argued that some group agents, like democratic states, are bound by moral requirements which fairly divide amongst each member of the group owing to the fact that the group agent roughly expresses the will of the group (provided that sufficient standards of democratic legitimacy are met), and each member is morally required to fulfil their share and only their share, regardless of what others do, in “Distributive Justice and Distributed Obligations”, *The Journal of Moral Philosophy* (2017) and “Ought We Do What We Ought to Be Made to Do?” in *Reasons and Intentions in Law and Practical Agency* ed. by George Pavlakos and Veronica Rodriguez-Blanco (Cambridge: Cambridge University Press, 2015).

harm aversion, in the sense that Miller discusses. But that is *not* a feature of teleological views like luck egalitarianism.

In my view, the arguments in defence of doing one's fair share in terms of beneficence and averting harm do not readily export to the problem of reducing inequality in situations of partial compliance with its requirements. Aside from left-libertarian views and teleological views combined with rule consequentialism, there is little egalitarian reason to support the fair shares thesis. Some of the reasons for rejecting the fair shares thesis point us towards a different, and, in my view, better answer: *the maximisation thesis*.

1.5 The Maximisation Thesis

The maximisation thesis says that everyone, other than the worst-off person in society, has an egalitarian reason to maximise their contribution to reducing inequality by redistributing or foregoing benefits until doing so any further would worsen inequality, relative to the factual status quo. On this view, there may well be an egalitarian reason to make oneself worse off than one would be in counterfactual worlds in which the metric of equality is equally distributed, and there may sometimes be an egalitarian reason to redistribute benefits until one is no better off than the worst-off person in society.

For example, suppose, for the sake of argument, that the metric of equality is welfare and suppose further that the timespan of egalitarian concern is whole lives. If the maximisation thesis is correct, then we have an egalitarian reason, over the course of our whole lives, to sacrifice our own welfare in order to improve the lives of the worse off than us until any further sacrifices would worsen inequality.

The maximisation thesis is supported, I believe, by luck egalitarianism (or *equality as comparative fairness*) and *equality as competing claims*. These two views compare how people fare relative to each other, and, on these views, there is no uniquely egalitarian counterfactual distribution which we can use to determine fair shares; we can either increase or reduce inequality by increasing or decreasing the extent to which people's shares of the metric of equality differ. To illustrate this point, consider two of our best ways of comparing the relative injustice or badness of different unequal distributions, neither of which support the conclusion that there ceases to be an egalitarian reason for people to redistribute benefits once they reach, or fall below, their fair share.

1.5.1 Distributive Patterns & Egalitarian Complaints

The first way of comparing unequal distributions is to adopt Larry Temkin's account of *equality as comparative fairness*. On this view, the badness of inequality resides in the impersonal disvalue of the pattern of welfare that obtains across a population. Temkin proposes that we measure the badness of inequality in a situation by assigning each individual an egalitarian *complaint*, the size of which is determined by how they fare relative to others.

Temkin suggests that the size of an individual complaint may be determined by the size of the gap between each individual and (1) the best-off person within the population, or (2) each of the individuals who are better off than them, or (3) the average level of advantage. On the first two views, everybody, other than the best-off person, has an egalitarian complaint. On the third view, everybody below the average level of advantage will have an egalitarian complaint and everyone above the average level will not.³² The size of the inequality in any distribution is a function of these egalitarian complaints.

Temkin also proposes three ways of determining the *overall* amount of inequality in a distribution: (1) an additive view, in which one distribution is more equal than another if the sum of complaints is smaller, (2) a weighted additive view in which one distribution is more equal than another if the sum of *weighted* complaints is smaller, and (3) a maximum view, in which one distribution is more equal than another if the *largest* complaint is smaller.

An implication of the additive and weighted additive views is that if a complaint can be reduced or removed without giving rise to an equivalent or larger complaint elsewhere in the distribution, then inequality has been reduced (on the maximum view, only reducing the size of the largest complaint would reduce inequality). There is nothing about this view which precludes inequality from being reduced by redistributions which cause a person's goods to fall below their fair share, as I will shortly illustrate.

1.5.2 Competing Claims & Justifiability to Each

Another way of measuring the badness of injustice of inequality is to adopt Thomas Nagel's *competing claims* view.³³ In contrast to Temkin's view, on this view injustice

³² Temkin, *Inequality*, Chapter 2.

³³ Thomas Nagel, "Equality" in *Mortal Questions* (Cambridge: Cambridge University Press, 1979).

does not reside in distributive patterns, or outcomes, but rather in the extent to which individuals' competing claims are fairly satisfied.

The size of each individual claim is determined by the following features: (1) each individual has a claim to an outcome, or benefit, if it makes them better off, (2) the larger the possible increase in advantage available to them, the stronger their claim, and (3) the lower the absolute level of advantage they are at, the stronger their claim. One distribution is more just than another if, and only if, it more fairly satisfies everyone's competing claims in proportion to their strength.

An implication of *this* view is that if a claim can be satisfied without giving rise to an equivalent or larger claim, then the outcome will more fairly satisfy people's competing claims. But again, there is nothing about this view which precludes unfair inequality from being reduced by redistributions which cause a person's benefits to fall below their fair share, providing a claim is satisfied without giving rise to an equivalent or larger claim.

So, on Temkin's complaint view, redistribution from better to worse off reduces inequality if it reduces the aggregate magnitude of egalitarian complaints, and on an adaptation of Nagel's competing claims view, redistribution from a better off to a worse off person reduces inequality if it results in a fairer satisfaction of claims.³⁴ Neither of these views imply that individuals can do nothing more to reduce injustice once they have done their "fair share" relative to the status quo or any counterfactual state of equality.³⁵

Consider the following illustration, making the simplifying assumptions that that each person has an equal claim to a share of the goods to be distributed and (perhaps contrary to fact) that people do not have special claims or entitlements to things they have already legally earned or been given under the distribution which constitutes the status quo.

³⁴ There is some evidence that these conclusions are not widely supported by people's intuitive judgements. Yoram Amiel and Frank Cowell's *Thinking About Inequality* (Cambridge: Cambridge University Press, 2009), (pp. 39-46), reports that respondents to questions about the Pigou-Dalton principle disagreed over whether all even transfers from a person who has more of a good to a person who has less, where the absolute positions of others remained unchanged, reduced inequality.

³⁵ A third way of comparing unequal distributions is to ask how an agent acting in their rational self-interest would rank the distributions behind a non-probabilistic veil of ignorance which strips them of all knowledge of who they are. On this view, one distribution is more just than another if it would be rationally preferred behind the veil of ignorance, but I shall set aside this measure.

	Person A	Person B	Person C
D1	1	3	5
D2	2	2	5

Each of the following is true. (1) The “fair share” of each individual within these distributions, relative to the status quo, is 3. (2) D2 could be brought about from D1 by a transfer from Person B to Person A. (3) When comparing the distributive patterns, D2 is more equal than D1 on the “relative to everyone” view of complaints combined with the additive and weighted additive views, and on the maximin view combined with any of the ways of determining the size of an egalitarian complaint.³⁶

D2 also more fairly satisfies competing claims than D1, due to the fact that, being worse off in absolute terms in D1, A has a stronger claim to benefits than B. What this example shows is that on each of these ways of measuring the badness or injustice of inequality, there is still an egalitarian reason for a person to redistribute benefits to the worse off even when they reach, or fall below, their fair share.

We can also go further and say that D2 is more *just* than D1, since it is plausibly true that, other things being equal, for any two Pareto incomparable distributions concerning the same population and roughly the same overall amount of goods, the distribution which is more equal is more just.³⁷

³⁶ It is true that D1 and D2 do not differ in terms of inequality on the “relative to average” and “relative to best-off” view of complaints combined with the non-weighted additive view. These measures support the fair shares thesis, although, as I have argued, they have counterintuitive implications which should lead us to be suspicious about their plausibility.

³⁷ For example, a pluralist egalitarian would consider an equal distribution more just than an unequal distribution containing the same amount of goods. An equal distribution would be more justifiable to each person than an unequal distribution containing the same amount of goods, other things being equal.

This definition of unjust inequality is consistent with the claims that (a) equality is only one aspect of distributive justice to be balanced alongside other, distinct values, (b) inequality is only unjust when, and because it makes some worse off than they could otherwise be, and (c) an equal distribution is less just than an unequal, Pareto superior distribution.

Moreover, prioritarianism would also favour a more equal distribution over a less equal one when the amount of goods to be distributed remains fixed. It would also be in everyone’s rational self-interest to prefer a more equal distribution to a less equal distribution containing the same amount of goods behind a non-probabilistic veil of ignorance.

Of course, people in unequal societies are differently located with respect to inequality and it is always best, from an egalitarian perspective, for transfers to go from the best-off person possible to the worst-off person possible. Furthermore, it is plausible that absolute gaps between people's shares of the metric of equality are morally worse at low levels than high levels. For example, the inequality in (10, 20) is worse than the inequality in (100, 110), even though the absolute size of the gap between the two individuals is the same in each. Accordingly, the egalitarian reason is scalar: one's egalitarian reason to redistribute benefits is stronger the better off you are in comparison to others for the simple reason that the more one has to redistribute, the greater the size and amount of egalitarian complaints that one can reduce or remove, and the harder it is to justify retaining one's holdings in light of the greater claims of others.

I should acknowledge some of the further complexities raised by these measures of inequality. On the complaint view, not *all* transfers from a better off person to a worse off person will reduce inequality for the following reason: any transfer will alter the *relative* levels of individuals whose absolute share of goods is unaffected by the transfer. That is, a transfer from one person to another could increase or decrease the size of the egalitarian complaints of otherwise unaffected individuals, and, indeed, the overall number of egalitarian complaints within the distribution. By contrast, on the competing claims view, this makes no difference. For that reason, the two accounts are not extensionally equivalent. All transfers that reduce inequality on the complaints account, however, will form a subset of the transfers which reduce injustice on the competing claims account.

Finally, many possible redistributive acts do not meet the simplifying assumptions I introduced in the illustrative example. The example shows an *even* redistribution in which the value of the transferred good does not increase or decrease during the transfer and, consequently, the total amount of goods to be distributed remains fixed.

Redistributive acts can also be efficient or inefficient (leaky), rather than even, and thus increase or decrease the total amount of the metric of equality. For example, if money has diminishing marginal utility, and the metric of equality is welfare, then transfers of money from a better off person to a worse off person will be efficient and increase the total amount of the metric of equality. Redistributive transfers can also result in the parties to the transfer changing location within the distribution, that is, a person who has more of a good than the recipient ends up with less than the recipient.

These are all interesting complexities that require additional examination. But as far as demonstrating the plausibility of the maximisation thesis as far as luck egalitarianism

and the competing claims view goes (at least when it comes to even transfers), my argument concludes here.

1.6 Two Problems

In this final section, I confront some problems for the maximisation thesis which are of practical importance and offer an interesting route for further enquiry. If either problem is insurmountable, then some of the above conclusions are practically insignificant in a range of realistic situations. Since this is a result I wish to avoid, I explore some ways of overcoming them.

The first problem concerns the egalitarian reason to redistribute benefits for the sake of equality. When some patterns of inequality obtain across large populations, the maximisation thesis faces a practical trilemma. To repeat, the maximisation thesis says that everyone, other than the worst-off person(s), has an egalitarian reason to reduce inequality, relative to the factual status quo.

To illustrate the trilemma, consider the following situation in which there is a small group of very well-off people and a much larger group of worse off people. The illustration depicts a population of 100 people. Suppose for the sake of illustration that the metric of equality is wealth. In D1, which represents the status quo, two individuals *A* and *B*, are equally well-off and 98 individuals are equally badly-off. Suppose *B* is unwilling to redistribute anything for the sake of equality and *A* must choose whether to unilaterally redistribute. *A* faces at least three options.

	Person 1	Person 2	...	Person 98	Person A	Person B
D1	1	1	1	1	100	100
D2	1	1	1	10	90	100
D3	1.1	1.1	1.1	1.1	90	100
D4	2	2	2	2	2	100

The first option, represented by D2, is to transfer a fixed amount of wealth which falls short of a maximal sacrifice (in this example, 10 units) to *one* of the worse off individuals (this could also be modified so that *A* disperses a fixed amount of wealth across a small subset of the worse off members population). But by doing so, inequality is *worsened* in some respects. Although the gap between *A* and all of the worse off 98 people has been reduced, there is now unfair inequality where none existed previously: between the people with 1 unit and the person with 10, and between *A* and *B*.

The second option, represented by D3, is to transfer a fixed amount of wealth which falls short of a maximal sacrifice (again, in this example, ten units) and equally disperse it across *all* 98 of the worse off people. But this only very *trivially* reduces the gap between the worse off and the better off *and* it introduces unfair inequality between A and B where there was no inequality before.

The third option, represented by D4, is to make a maximal sacrifice by transferring 98 units of wealth and equally dispersing it across all of the 98 worse off people. But again, this only trivially reduces inequality. Although it completely removes the gap between A and the rest of the population, it does not do much to close the gap between the worse off and the best-off, and it introduces a new large gap between A and B where none existed previously.

Although D4 is probably the best distribution from an egalitarian perspective (though this is by no means obvious) the cost to A of bringing it about is extremely high. And, other than act consequentialists, most people believe that people cannot be morally required to make great sacrifices in order to bring about trivial improvements in outcomes. Although there *is* an egalitarian reason to widely disperse benefits, this reason is extremely trivial, and is unlikely to ever be of sufficient weight to become a moral requirement.

What this shows is that in certain unequal situations containing some patterns of inequality, unilateral redistribution may worsen inequality, or only reduce it very trivially. This is a problem because one might assume that situations with a small wealthy elite and a much larger population of worse off people are paradigmatic examples of situations in which the rich have an egalitarian reason to unilaterally redistribute benefits. Yet the opposite seems to be sometimes true: the more unequal the society, the more defensible it is to retain benefits.

Now, depending on how the facts of the illustration are filled out, the benefits to the lives of the worse off people may be fairly substantial, and there would be clear reasons of *beneficence* to make the transfer. But from an *egalitarian* perspective, the improvement in terms of reducing unfair inequality is very small. The trilemma is magnified across much larger populations in which the ratio between the worse off and better off is roughly the same as depicted in the example. The example, recall, concerns a population of only 100 people. The implications of the trilemma are much more pronounced across populations of thousands or millions because the egalitarian reason to independently redistribute benefits would become increasingly trivial as the size of the population grows larger.

I think there are at least three responses to this problem. First, it is much more of a worry for some metrics of equality than others. If the metric of equality is wealth, then in circumstances of extreme inequality of wealth, this problem would be acute. But if the metric of equality is *welfare*, for example, then the problem is typically likely to be less acute. That is because, since welfare cannot be transferred directly, the egalitarian reason to redistribute benefits will take the form of transferring the *sources* of welfare, like money. And given its diminishing marginal utility, transfers of money will not be so trivial, from the perspective of reducing inequality of welfare.

Second, in these situations, we should conclude that direct transfers of the metric of equality are not the *best* way to serve egalitarian ends as individuals. Instead, the best way to reduce inequality may be to direct resources to schemes which help to eliminate the sources of unfair inequality by directing resources into better health-care, education, housing, and so on. Again, the best way of doing this will depend on the metric of equality.

Third, we should concede that there are only very trivial objective egalitarian reasons for the well-off to carry out direct transfers in these kinds of cases. But we might still think that people with egalitarian *beliefs* still have good reasons to redistribute benefits, despite the fact that doing so would only make the outcome trivially better from an egalitarian perspective.

To draw an analogy, suppose, for the sake of argument, that refusing to purchase and eat factory farmed meat made absolutely no causal difference to the amount of suffering endured by non-human animals. I think that people who believe that factory farming animals for meat is morally abhorrent still have a subjective reason to refrain from participating in, and financially contributing to, a practice that they regard as deeply wrong, even if so refraining makes no causal difference to the amount of suffering endured by non-human animals. But people who falsely believe that factory farming is permissible do not obviously have any reason to refrain from contributing to that suffering when doing so would make no causal difference (although there might be objective reasons to refrain from individual complicity in collective wrongdoing even when one's own contribution is neither necessary nor sufficient).

Similarly, people with egalitarian beliefs have a reason to refrain from benefitting from what *they* regard as an unfair distribution of the metric of equality. There may be reasons to avoid complicity in something that one considers unfair or unjust. There are of course salient differences between the factory farming example and inequality. For one, in the case of animal suffering, the belief is that it is *wrong* to factory farm animals.

Whereas in the case of unfair inequality, people with egalitarian beliefs may regard inequality as *bad*, but not think it wrong. Moreover, to purchase or consume factory farmed meat is to actively take steps to contribute to wrongdoing, whereas to benefit from unfair inequality takes the form of allowing or sustaining something that one objects to. Nevertheless, I would suggest that this gives rise to a reason to refrain from complicity, even if doing so only makes trivial improvements to the distribution of the metric of equality.

Finally, I think that the competing claims view is better placed than luck egalitarianism to deal with this problem. That is because the competing claims view doesn't, or at least doesn't merely say that we should reduce inequality because it improves states of affairs. What's essential to the view is that we should reduce inequality when and because it *cannot be justified* to those who are worse off. The inability to justify retaining benefits in the face of the greater claims of others seems to give rise to a more significant reason to redistribute benefits than negligibly improving outcomes.

The second problem for the maximisation thesis concerns *foregoing* benefits for the sake of equality. Foregoing benefits can occur in at least two ways. First, it can take the counter-Paretian form of refusing to accept a benefit which would worsen inequality but with the result that the benefit will not be bestowed upon *anybody*. This would be the case, for example, if someone refrained from using expensive private healthcare or first-class seats on planes and trains even though doing so would not benefit anybody. (One could, of course, use the money to pay for someone *else* who is much worse off to secure the benefit, but that would be redistributive.)

Second, it can take the form of refusing to accept a benefit when doing so will result in *another equally well-off* individual securing the benefit instead. This would be the case, for example, if someone refused to pay for an expensive education but this made no difference to the number or social composition of people who are unfairly well-educated.

From the perspective of *fairness*, it does not make a difference whether or not one foregoes benefits that will be enjoyed by other equally well-off people anyway. And from an *all things considered* perspective, there does not seem to be a decisive reason make counter-Paretian decisions. The problem implies that foregoing benefits is pointless from an egalitarian perspective and unnecessary from an all things considered perspective. Of course, in some cases, foregoing a benefit *will* reduce unfairness without acting contrary to the Pareto principle. Specifically, if one foregoes a benefit that will

instead be bestowed on a worse off person, there is an egalitarian reason to forego the benefit. But often, this will not be the case.

I think that luck egalitarianism and the competing claims view are ill-equipped to explain how there could be a sufficient or decisive reason to forego benefits for the sake of equality in the kinds of situations described above (although I would again suggest that people with egalitarian *beliefs* have subjective reasons to refrain from accepting benefits that they consider unfair).

However, other aspects of egalitarianism are better able to explain why there could be a sufficient or decisive egalitarian reason to forego benefits in these kinds of situations. For example, I think social egalitarianism, or “non-intrinsic egalitarianism”, is better equipped to explain why there is an egalitarian reason to forego benefits for the sake of equality. To accept unfair benefits is to make an exception of oneself, to enjoy things that others can’t, and in many cases, never will. Accepting a large amount of unfair benefits insulates oneself from some of the hardships that others face, whereas foregoing benefits avoids that moral taint.

Moreover, deontological egalitarian views which object to distributive inequality when and because it results from treating people unequally may imply that there is an egalitarian reason to forego benefits. One form of treating people unequally is to causally perpetuate institutions which work to the unfair advantage of some at the expense of others. Accordingly, when accepting unfair benefits contributes to the perpetuation of unfair institutions, there is an egalitarian reason to refrain from doing so.

Individual aspects of egalitarianism, considered in isolation, will, in various situations, imply that individuals have no egalitarian reason to unilaterally redistribute or forego benefits. However, if we take into account the full arsenal of egalitarian considerations, there is a wide spectrum of reasons to do so.

1.7 Conclusion

I started by asking what egalitarianism demands of individuals in unequal societies. I said that there are at least three appealing answers to the question, and, given the diversity of reasons to object to inequality, it is unsurprising that different egalitarian views would have different implications for this practical question.

The institutional thesis is supported by some varieties of instrumental egalitarianism, deontological egalitarianism, procedural egalitarianism and contractualist egalitarianism. If the entirety of egalitarian concern was captured in those views, as some believe, then there is only a very trivial egalitarian reason for individuals in unequal societies to redistribute or forego benefits for the sake of equality. But those views do not capture the full spectrum of egalitarian concern.

The fair shares thesis is supported by left-libertarianism and luck egalitarianism combined with rule consequentialism. It is also supported by combining luck egalitarianism with some implausible measures of inequality, but otherwise lacks support. The rationale for adopting the fair shares thesis in the case of beneficence and the duty to avert harm do not extend to the case of reducing inequality. This is a surprising result, since some have mistakenly assumed that the fair shares thesis is supported by luck egalitarianism.

The maximisation thesis is supported by equality as comparative fairness and the competing claims view, at least. On these views, there is an egalitarian reason for each individual in an unequal society to maximise their contribution to reducing inequality, relative to the factual status quo. An implication of this is that there is still an egalitarian reason to redistribute benefits even if this makes a person worse off than they would be in the counterfactual world in which all existing goods were equally distributed.

So what does egalitarianism demand of individuals in unequal societies? In my view, because the institutional thesis derives support only from views which do not capture the full spectrum of egalitarian concern, and the fair shares thesis is unsupported by the logic of most egalitarian views, we should regard the maximisation thesis as the most thoroughgoing answer to this question.

2 IF YOU'RE AN EGALITARIAN, WHY NOT ROB A BANK?

2.1 Introduction

There are many ways of reducing unfair inequality. Some are morally permissible. It is permissible for a legitimate government to impose coercive taxation in order redistribute money from the better to worse off. Others are impermissible. It is wrong to forcibly redistribute kidneys, eyeballs, or bone marrow for the sake of equality. Sometimes rights stand in the way of fairness.

In some cases, however, the moral status of egalitarian redistribution is less clear. One class of difficult cases are those in which individuals reduce unfair inequality by possessing or redistributing money and property which legally belongs to somebody else, without the informed consent of the legal owner, and without providing compensation.¹ I will call this *egalitarian vigilantism*. Such flouting of property laws can take many forms, including banditry, embezzlement, fraud, looting, occupation, and trespass. Here are two examples.

In April 1932, over 400 people, led by Benny Rothman, an activist in the Young Communist League of Manchester, conducted a deliberate mass trespass of Kinder Scout, a moorland plateau in the Peak District in the North of England, legally owned by the Duke of Devonshire. According to a report in the *Manchester Guardian*:

Triumphant, the marchers gathered on the peak before returning to Hayfield singing *The Red Flag* and *The Internationale* and shouting “Down with the landlords and ruling class and up with the workers!” and “Down with the bobbies!” By the time they reached Hayfield the police presence had grown considerably and five “ringleaders” were arrested.²

¹ Although typically such vigilante acts are illegal, some legal regimes permit the intentional infringement of legal property rights in the right circumstances. For example, the legal doctrines of adverse possession, acquiescence, and necessity all allow people to acquire the legal title to property through infringing the legal entitlements of others.

² Reported in the *Manchester Guardian*, 25th April, 1932. Trespass was not actually a criminal offence in the United Kingdom in 1932. The arrests were made because the trespassers were involved in a violent skirmish with gamekeepers armed with sticks.

The stated purpose of the trespass was to protest against the Enclosure Acts, a series of Acts of Parliament which confiscated common land and bestowed legal rights of ownership onto private landowners, as well as to draw attention to the inequality and poverty of the 1930s.³ The trespass contributed to the eventual enactment of legislation which curtailed the legal rights of landowners to prevent the general population from accessing land, including the National Parks and Access to the Countryside Act 1949, and the Countryside and Rights of Way Act 2000.

In 2010, the Greek economy was in crisis which was partly caused, some believe, by financial corruption in state-owned banks. In February of that year, an armed man walked into a branch of the Greek National Bank and stole 250,000 euros. In October, the same man carried out two additional bank raids and escaped with 240,000 euros. It is alleged that the culprit was Vassilis Paleokostas, also known as the “Greek Robin Hood”.

Paleokostas was known to the police because he was responsible for carrying out a crime spree over the course of three decades, during which he stole millions from state-owned banks, kidnapped rich tycoons in order to extort ransoms, and then redistributed the proceeds to impoverished rural farmers. According to his family and friends, “Vassilis suffered his bosses’ capitalist exploitation, working as a wage slave in a factory. So, he turned against those bosses. [He] may have been a thief, but never a criminal”. At the time of writing, he remains a fugitive.⁴

I take it that while many would regard the mass trespass of Kinder Scout as a paradigmatic example of permissible direct action against the unequal distribution of land, many would have some reservations about whether the bank robberies and extortions were permissible forms of retaliation to the unequal distribution of wealth. Do these judgements withstand scrutiny and, if so, how are they best explained?

³ Benny Rothman, *The 1932 Kinder Trespass: A Personal View of the Kinder Scout Mass Trespass* (Altrincham: Willow Publishing, 1982), p.11.

⁴ Reported in Jeff Maysh, ‘The Uncatchable’, *BBC News Magazine*, 25 September 2014.

It appears that Paleokostas took to heart the Greek proverb: “If you steal something small, you are a petty thief, but if you steal millions, you are a gentleman of society”. There are many other examples of “modern day Robin Hoods”. To give one, in the early and mid-20th century, some Spanish and Argentinian anarchists practiced a doctrine of “expropriative anarchism”, and employed robbery, scams, and fraud as a means of financing rebellious activities. See Eric J. Hobsbawm, *Bandits* (London: Abacus, 2001), Chapter 8 for that example and others.

In this chapter, I examine the variables governing the permissibility of egalitarian vigilantism and, although I do not pretend to reach a complete or precise set of conclusions, I try to illuminate what is morally at stake in order to justify judgements about clear cases and inform deliberations about more difficult ones. My central conclusion is that in partially unjust societies, individuals may sometimes permissibly expropriate money and disused property from groups (like banks, corporations, and the state) which culpably perpetrate injustice when democratic or other legal forms of recourse are highly unlikely to be effective, provided that any property losses experienced by individuals are indirect and widely dispersed, and such acts are not so widespread as to be counterproductive.

The chapter proceeds in two stages. In the first part, I consider features of egalitarian vigilantism which make such actions easier or harder to justify independently of the justice or injustice of property laws. I argue that there are at least six dimensions across which an act of vigilantism may be easier or harder to justify.

The second part of the chapter is divided into three smaller sections and explores how the permissibility of egalitarian vigilantism varies in accordance with the extent to which property laws are unjust.

In the first section I point out that, in some situations, flouting property laws in order to reduce inequality is clearly impermissible. It would be morally wrong, as well as illegal, to expropriate money and property when property laws protect all and only moral property rights, property laws are decided upon through sufficiently fair procedures, or the distribution of the currency of justice is fair *enough*.

We can, however, certainly rule out the view that egalitarian vigilantism is *always* wrong. In some situations, flouting property laws is clearly permissible. For example, Robin Hood's acts of brigandage, real or imagined, are typically believed to be not only justified, but admirable acts of resistance against arbitrary inequality, as well as the injustice of tyranny and oppression. In the second section, I examine the reasons for thinking that egalitarian vigilantism is easier to justify when property laws are *severely unjust*.

Intermediate cases between these two extremes are more difficult to judge. In the third section, I turn to the more interesting, complex and murky question of when it would be permissible to expropriate things which are legally owned when property laws are *partially unjust*. This question is more interesting because some legal property regimes in the actual world plausibly fall into this category. But it is also more complex and murky because our intuitive judgements about such situations are less clear.

2.2 Preliminary Remarks

Throughout, I make the following three assumptions.⁵ First, I take it for granted that reducing unfair inequality makes things better in one respect, and I assume that equality is one aspect of distributive justice. Other things being equal, reducing inequality reduces injustice.⁶ What distributive justice requires must be balanced against other moral considerations in order to determine what is morally required, permitted, and forbidden, all things considered. I assume that there are moral constraints which prohibit people from treating others in certain ways, without their consent. It is wrong to reduce inequality by transgressing these constraints.

Second, we are sometimes morally permitted to disobey unjust laws. There is no general duty or obligation to obey all laws *merely* because they are laws. However, we are sometimes morally required to obey laws for some other reason.⁷ I will qualify this assumption by noting that the state may permissibly enforce a law even though it is permissible to disobey it. The permissibility of vigilantism does not entail anything about whether or not the state may permissibly arrest, convict, and punish vigilantes.

Third, although people have moral property rights, under realistic conditions, people are sometimes not morally entitled to everything that they legally own.⁸ I understand

⁵ Given alternative assumptions, the permissibility of egalitarian vigilantism would be more straightforward. For example, if act consequentialism is true, then the permissibility of an act of vigilantism would simply be determined by whether the act brings about consequences which are at least as good as any alternative act. If rule consequentialism is true, then the permissibility of an act of vigilantism would be determined by whether it would be required or prohibited by rules which, if followed by everyone, would make things go impartially best (even if a particular act of vigilantism which broke those rules would make things go impartially best). And if we *do* have a general moral obligation to obey all laws, then vigilantism would always be wrong.

⁶ My arguments should apply whether inequality is bad when and because it is unfair, where the badness resides in the relation of some being worse off than others, or whether inequality is bad when and because it cannot be justified to each. My arguments could also be modified in accordance with the assumption that equality is not an aspect of distributive justice, but *sufficiency* is. With this modification, most of my arguments would still apply, but there would be a reason of justice to carry out vigilantism only in order to raise people to, or towards, a sufficiency threshold.

⁷ See Christopher Heath Wellman and A. John Simmons, *Is There a Duty to Obey the Law? For and Against* (New York: Cambridge University Press, 2005).

⁸ I will restrict my discussion to the dispossession and redistribution of tangible property and exclude intellectual property, although the permissibility of expropriating intellectual property is also of interest. Suppose, for example, that a vigilante could acquire information about how to manufacture a patented medicine that could improve the lives of many people.

property as a relation between persons regarding an object. If I have a cluster of legal rights to possess, use, derive income from, and transfer something, then you are legally forbidden from possessing, using, deriving income from, or transferring that thing. Similarly, if I have a set of *moral* rights to possess, use, derive income from, and transfer something, then you are morally prohibited from possessing, using, managing, deriving income from, or transferring that thing.

If I have a legal right, but not a moral right, to something, then you have a legal duty to refrain from taking it, but you may have *moral* permission to take it. Similarly, if I have a moral right to something, but that right is not legally recognised, then you have a moral duty to refrain from taking it, even if that moral duty is not legally enforced.

Taken together, these three independently plausible assumptions yield the possibility of circumstances in which a vigilante could reduce unjust inequality by breaking laws which are not morally binding, without violating moral property entitlements. If it would be wrong to do so, there must be some further factors at stake which act as overriding presumptions against egalitarian vigilantism. I now turn to consider what those presumptions may be. In order to set aside unnecessary complications throughout, I will assume that in each example, vigilantism reduces inequality both non-trivially, and in the most effective possible way.⁹

Furthermore, I take moral property rights to be independent of the moral justification of a *system* or *institution* of legal property rights. The establishment of an institution of legal property rights may be morally justified on the grounds that, compared to possible alternatives, property rights secure stability in society. Nevertheless, as a matter of logical possibility, a legal system of property rights may be morally justified even though moral entitlements and legal entitlements do not perfectly correlate and some individual acts of lawbreaking are morally permissible. And although stability and order are important, another desideratum of a system of property rights is that it is responsive to reform and change in response to justified civil disobedience and changing social attitudes.

⁹ A worry about my approach is that one might think that if you're going to resort to dispossessing people of what they legally own, you ought to use the money to do the most impartial good that you can rather than using it merely to reduce unfair inequality. I grant this worry, but would add that if the arguments presented in the preceding chapter are correct, the demands of equality would often correlate with what is impartially best. In any case, my arguments still apply if we assume that vigilante action is aimed at doing what is impartially best rather than aimed at reducing inequality.

2.3 Dimensions of Vigilantism

In this part of the chapter, I consider six dimensions across which the moral status of egalitarian vigilantism varies, independently of the degree to which property laws are unjust.

2.3.1 Violence and Coercion

It is easier to justify non-violent or non-coercive forms of vigilantism than violent or coercive ones (which is not to say that politically motivated violence and coercion is never justified). Recall the Greek bank robberies described at the outset. Although Vassilis Paleokostas did not physically harm anybody, he did terrorise innocent bank employees while armed with a weapon. He also kidnapped rich tycoons in order to exact ransoms.¹⁰ But not all forms of vigilantism take a violent or coercive form. Contrast Paleokostas's bank robberies with the following made up case.

Harmless Hacker

A vigilante computer hacker gains access to the records of an exclusive bank. The hacker transfers large amounts of money from the accounts of the bank's millionaire clients into another account, before withdrawing and redistributing it to the least well-off.

This act of direct redistribution is intuitively more justifiable than the Greek bank robberies. Violent coercion, and the threat of it, violates moral constraints against interference. Treating people in such a way does not properly respect them as a person. Accordingly, there is a very strong presumption against it, even if doing so reduces unjust inequality.

2.3.2 Varieties of Property

It is harder to justify the forcible dispossession of some kinds of goods than others. Other things being equal, it is much worse to redistribute personal possessions, or property with special *attachment* claims, than disused buildings and land, food, and money, absent special attachment claims.

¹⁰ He did not physically harm the people he kidnapped. One of his victims is reported to have said: "My kidnapers' behaviour was not bad at all. I was not scared for myself. Actually, I enjoyed some wide-ranging discussions with the kidnapers."

Some believe, and I agree, that personal possessions and property with special attachment claims are akin to an extension of one's personhood, and control rights over them are a part of one's personal sovereignty.¹¹ Like violence, the expropriation of personal possessions, or the intrusion of property with special attachment claims (like home invasion) fails to respect the victim as a person. It would be presumptively wrong for even a democratically legitimate government, let alone an outlaw, to forcibly redistribute personal possessions or property with special attachment claims. For similar reasons, it is very hard to justify taking anything which is in someone's current possession.

By contrast, sometimes the link between property and personal sovereignty is weak or non-existent. Taking things which are not in the actual possession of the legal owner and lack special attachment claims, like money in a bank account or land and buildings which have long been left disused, is not an equivalent violation of personal sovereignty. It is also easier to justify the expropriation of money or food than property of either kind. There can be no special attachment claims to food and money (with the exception of collector's items).

Accordingly, the presumption against redistributing money, food, and disused buildings and land, absent special attachment claims, is much weaker than the presumption against the redistribution of personal possessions. There is a clear moral difference between poaching pheasants and pilfering gooseberries from the estate of a wealthy landowner and sneaking into the landowner's home to steal a family heirloom of equivalent value.¹² The mass trespass of Kinder Scout cannot plausibly be described as a violation of the 9th Duke of Devonshire's personal sovereignty.

2.3.4 Direct and Indirect Losses

In my view, it is harder to justify *directly* imposing property losses on individuals, relative to the factual status quo, than to indirectly make people worse off than they *would have been* had vigilante action not occurred. The imposition of indirect losses takes at least two forms. First, it can take the form of *diverting* funds from intended

¹¹ For example, see Margaret Jane Radin, "Property and Personhood", *Stanford Law Review*, 34.5 (1982), 957-1015 (p. 986).

¹² According to James C. Scott in *Two Cheers for Anarchism* (Princeton: Princeton University Press, 2012), p.11: "For the two centuries from roughly 1650 to 1850, poaching [...] from Crown or private lands was the most popular [frequent and approved of] crime in England".

recipients. For example, consider a clerk who is able to divert bequests to the arts or charities for pets and instead donate the money to organisations which help the needy. Second, it can take the form of imposing losses as an unintended *side-effect*. Consider the following example.

Victimless Fraud

An egalitarian vigilante has a variety of fake identities. He uses the fake identities to request personal and commercial loans from many different banks but has no intention of repaying them. Instead, once he has fraudulently obtained the money, he redistributes it to the least well-off. The losses are covered by the bank's insurance, and are not directly passed on to individuals. The insurance company voluntarily chose to accept the risks of covering losses imposed by expropriation.

In this case, no losses are *directly* imposed on individuals by the vigilante's action. Other examples of "victimless" vigilantism are deliberate overpayment of welfare benefits by state employees, the embezzlement of funds by corporate employees, welfare fraud, minor tax evasion, and some transactions in the underground economy.¹³

That being said, such acts may make people worse off than they would otherwise have been as a side-effect. For example, fraud may raise the price of insurance premiums. Minor tax evasion or welfare fraud takes funds that could have been otherwise distributed, and state action to tackle these transgressions gives rise to administrative costs, diverting funds that could have been distributed otherwise.

However, imposing indirect losses as a side-effect is intuitively easier to justify than imposing direct losses. I would posit that this is explained by the fact that indirect losses are typically more causally distant from the point of action than direct losses and, in at least the examples described above, there is another agent in the causal chain that bears some responsibility for the outcome.

2.3.5 Consolidated and Widely Dispersed Losses

It is easier to justify egalitarian vigilantism when property losses are widely dispersed across people and time than when they are consolidated into a small number of people or imposed suddenly.

¹³ See Tommie Shelby, "Justice, Deviance, and the Dark Ghetto", *Philosophy and Public Affairs*, 35.2 (2007), 126-60 (p.152).

Plausibly, it is worse to impose a large harm on one person in order to benefit others than it is to widely disperse tiny harms on many people in order to benefit others.¹⁴ Likewise, other things being equal, it is worse to impose large property losses on few people in order to benefit others than to widely disperse equivalent aggregate property losses across many people for the same purpose.

Similarly, I think it is typically worse to impose a sudden property loss on someone in order to benefit others than it is to disperse the same loss gradually over a long period of time. Consider the following modification of the *Harmless Hacker* case.

Harmless Hacker II

A vigilante computer hacker gains access to the records of an exclusive bank. Gradually, over a period of first weeks, then months, and then years, the hacker transfers small amounts of funds from the accounts of each of the bank's millionaire clients into another account, before withdrawing the money and redistributing it to the least well-off. Each individual transfer is so small that it goes undetected, even when the gradual accumulation of transfers amounts to millions of pounds.

This act of vigilantism seems much easier to justify than the act described in the former *Harmless Hacker* case, which imposed sudden, non-trivial, property losses. This is plausibly because of the disruption to people's lives that occurs when they suffer a sudden, non-trivial loss of holdings.¹⁵ It is not always wrong to impose sudden property losses: such a strong view attaches too much moral significance to the status quo. But I think it counts *to some extent*.

2.3.4 Individuals & Group Entities

It is harder to justify imposing property losses on individuals than imposing them on entities like banks, corporations, or the state. Vigilantism which targets individuals is

¹⁴ Larry Temkin calls this the "Disperse Additional Burdens View", *Rethinking The Good: Moral Ideals and the Nature of Practical Reasoning* (Oxford: Oxford University Press, 2012), Chapter 3.

¹⁵ In "Prerogatives to Depart From Equality", *Royal Institute of Philosophy Supplement*, 58 (2006), 95-111 (p.110), Michael Otsuka writes: "If [...] the better off in the actual world really have a moral entitlement to that which they legally own, this would most plausibly be explained by the hypothesis that present legal entitlements have moral force because of the disruption to people's lives which would occur if the status quo were overturned and transformed into a justly egalitarian society. There may be limits to what one can, in the name of equality, force a person to give up which he already possesses."

especially egregious because it shows a disregard for their moral status as a person, especially when doing so violates widely accepted social norms and attracts widespread disapproval.

But this does not extend to vigilantism targeted at group entities. Group entities may sometimes be agents, but they are not persons, and we are not required to treat them as such, except when the effects of actions directed at group entities indirectly trickle down to the persons who are members of the group. Consider the following variation of the *Harmless Hacker* case, in which property losses are imposed on the *bank*, rather than on individual customers.

Harmless Hacker III

A vigilante computer hacker gains access to the records of a major retail bank. Gradually, over a period of first weeks, then months, and then years, the hacker erases small amounts of consumer debt from the bank's records. Each individual act of tampering is so small that it goes undetected, even when the gradual accumulation of the erasures wipes millions of pounds of consumer debt.

In this case, no losses are directly imposed on persons, although indirect losses could be widely dispersed side-effects. Intuitively, it is easier to justify than each of the former two *Harmless Hacker* cases.

2.3.6 Redistributive and Expressive

Finally, we can divide such acts of expropriation into those which are *expressive* and those which are *acquisitive* or *redistributive*.¹⁶ *Expressive* acts of egalitarian vigilantism are those which are intended to draw public attention to unfair inequality in order to bring about a change in people's attitudes towards the legal status quo, and to ultimately bring about changes in the law itself. They may also be intended to spark widespread emulation.

Typically, expressive acts of defiance are also acts of civil disobedience, although they need not always be, depending on how the parameters of civil disobedience are defined (for example, some believe that a necessary condition of civil disobedience is a willingness on the part of the perpetrators to submit to punishment, but one may

¹⁶ These terms are used in Eduardo Moises Penalver, and Sonia K. Katyal, *Property Outlaws: How Squatters, Pirates and Protestors Improve the Law of Ownership* (New Haven: Yale University Press, 2010).

expressively flout the law without being so willing to submit). The mass trespass of Kinder Scout was an expressive act of vigilantism.

Acquisitive, or *redistributive*, acts of expropriation are not intended to influence public opinion or bring about a change in the law, Although those things might be regarded as welcome side-effects, the primary purpose of expropriation is to reduce inequality by transferring things from the better to worse off. The Greek bank robberies were a redistributive act of vigilantism.

Given that they are intended to help to bring about widespread and lasting change, expressive acts seem more easily justifiable than acts which are merely redistributive. I myself believe that this difference is relatively trivial, or, insofar as it matters, it matters when and because expressive acts do more good, in the long term.

2.3.7 Summary

Acts of vigilantism which contain all of the six factors identified on the left will be very difficult to justify in comparison to acts which contain all of the eight factors on the right.

<i>Harder to Justify</i>	<i>Easier to Justify</i>
Harmful or Coercive	Harmless & Non-Coercive
Possessions	Money & Disused Property
Direct	Indirect
Consolidated Losses	Dispersed Losses
Persons	Group Entities
Acquisitive	Expressive

In order to act with the strongest possible justification, egalitarian vigilantes should peacefully exact money or disused property from group entities, and ensure that any losses imposed on individuals are indirect and widely dispersed. That being said, acts containing all of the more difficult to justify features could be permissible in some situations, for example, when property laws are severely unjust. On the other hand, even acts containing all of the more easily justifiable features may nevertheless be wrong in some situations. They would typically be wrong, for example, when property laws are sufficiently just.

This concludes my examination of the variables governing the permissibility and justifiability of different *acts* of egalitarian vigilantism, independently of the justice or injustice of the law. What does this reveal about the examples I described at the outset?

The mass trespass of Kinder Scout was mostly accomplished without violence or coercion (apart from the skirmish with game-keepers). It was also “out in the open”, both figuratively and literally. It was a transgression of land, rather than personal possessions and was an expressive act of civil disobedience. Although losses were eventually imposed on one individual, the Duke of Devonshire, the only losses imposed were the eventual curtailment of legal rights to prevent others from accessing land, rather than legal title. All of these features help to justify the judgement that the trespass was a clear instance of permissible, indeed, admirable, vigilantism.

Reconsidering Vassilis Paleokostas, it is easy to see why his vigilante acts were more difficult to justify than the trespass of Kinder Scout. Terrorising bank employees and kidnapping tycoons involved the violation of innocent people’s rights. On the other hand, he targeted a bank and dispersed losses widely through taking mere money. So was he justified in robbing the bank? I think to answer that question we need to think about how degrees of injustice make a moral difference.

2.4 Degrees of Injustice

The more unjust the law, the easier it is to justify vigilantism. But what more can be said? People mean different things by “injustice”. So let me delineate three respects in which a property law may be unjust.

First, justice is partly concerned with moral rights that are, in principle, permissibly enforceable (this excludes impersonal wrongs and rights that are not, in principle, enforceable). Second, justice is partly concerned with the fair distribution of the currency of justice (whatever that is). Third, justice is partly concerned with *procedural fairness*, or fairness in the way things are decided upon.

Property laws are unjust when they fail to legally prohibit people from violating rights, unjustifiably sustain an unfair distribution of the currency of justice, or were unfairly decided upon (for example, because not everyone had a fair say in the decision, the decision lacked actual or hypothetical agreement, or was not justified by appeal to reasons that all reasonable people can recognise. I will remain neutral about what

exactly counts as procedural fairness in order to set aside introducing further complexity).¹⁷

The extent to which property laws are unjust across each of these dimensions is relevant because the permissibility of vigilantism partly depends upon whether or not it is a *proportionate* form of retaliation to injustice and whether there are other courses of action available which could achieve the same objectives.

2.4.1 Just *Enough*

Let's say that property laws are just *enough* if and only if all of the following necessary and jointly sufficient conditions are satisfied: (1) moral rights, including property rights, roughly correlate with legal rights, (2) the distribution of the currency of justice is fair enough, (3) property laws were decided upon by sufficiently fair procedures.

I think that these conditions are individually necessary because plausibly, property laws imposed by a benevolent dictator are unjust even if they fairly distribute the currency of justice and otherwise perfectly reflect everyone's rights. And property laws decided upon with fair procedures can be unjust if they distribute the currency of justice unfairly. My use of the words "enough" and "sufficient" are vague and imprecise. But a degree of imprecision and vagueness is inevitable when discussing non-ideal situations. I assume that, with due examination, a more precise specification could be provided, but I do not provide one here.

Many people think that moral property rights can arise, in various ways, independently of institutions to recognise and enforce them.¹⁸ Imagine a society in which the set of all legal property entitlements is roughly extensionally equivalent to the set of all moral entitlements. In that situation, by stipulation, each person has a moral right to most of the things they legally possess, and people are morally prohibited from taking the

¹⁷ Each of these considerations is independent of the question of whether or not the law in question is permissibly enforceable through the threat of punishment and whether or not people have a moral obligation to comply with it.

¹⁸ For example, some people believe that people acquire moral entitlements to unowned parts of the world that they appropriate, provided that they leave "enough and as good for others". Others believe that desert may ground moral property entitlements. It is also plausible that people are morally entitled to unequal shares of the currency of justice which arise from a state of fair equality through gifts, consensual, non-fraudulent transfers, and freely chosen optional gambles.

property of others, except under exceptional circumstances (for example, when doing so is the only way to save lives or avoid some other moral catastrophe).¹⁹

It is very hard to justify vigilantism when property laws ensure that the distribution of the currency of justice is fair *enough*. For one thing, if the degree of unfair inequality in society is very small, then there is only a relatively trivial reason to reduce it. For another, if property laws are sufficiently just, then there is a very strong presumption against vigilantism.

Finally, if property laws are decided upon by sufficiently fair procedures, then there is clear legal and procedural recourse for reducing unfair inequality. It would be disproportionate to resort to vigilantism when other courses are available which can achieve the same objectives. A society in which property laws satisfies these conditions could contain unfair inequality. But it would nevertheless be impermissible to reduce it through vigilantism.

2.4.2 Severe Injustice

I now turn to the opposite extreme. Let's say that property laws are *severely unjust* if any of the following sufficient conditions hold: (1) respecting legal property entitlements requires the violation of people's basic moral rights, like rights of bodily integrity and other basic freedoms, (2) respecting moral property entitlements require sustaining a morally intolerable degree of inequality, or (3) the procedures which determine legal property entitlements are severely unfair.

We are usually permitted to disobey property laws when respecting them requires tolerating the violation of people's moral rights. A slaveholder is not morally entitled to the body and labour of a person who they have legally but forcibly enslaved, even if their entire life revolves around the assumption that the legality of the enslavement would continue in perpetuity. It is permissible for an enslaved person to disobey the law by escaping, if they can, and a third party could permissibly assist their escape.

We are also usually permitted to flout property laws when the degree of distributive unfairness in society exceeds a morally tolerable degree. For example, we may flout property laws when that is the only way to satisfy people's basic needs. It is very

¹⁹ For example, it is permissible to steal someone's rightfully owned luxury yacht, even if compensation is not possible, in order to save a drowning person, as Peter Unger argues in *Living High and Letting Die: Our Illusion of Innocence* (Oxford: Oxford University Press, 1996), pp. 63-64.

plausible that the satisfaction of basic needs has moral priority over legal entitlements.²⁰ If stealing from the well-off is the only alternative to starvation, then it is permissible to steal. Although strictly speaking, these are not *egalitarian* reasons to flout the law, they are reasons of distributive justice.

Laws imposed by severely unfair procedures are not morally binding. For example, we are not morally required to obey laws imposed by a tyrannical or oppressive dictator (though we would still be required to act in accordance with tyrannically imposed laws which happen to correlate with what morality requires). Each of these considerations goes some way to explaining why social bandits in peasant or feudal societies, who, against a backdrop of famine or political oppression, steal from the rich in order to give to the needy, are upheld as admirable vigilantes.

2.4.3 Partial Injustice

I am sure that many would agree with me that it is sometimes permissible to break property laws when they are *severely unjust*, and typically impermissible to break them when they are *sufficiently just*. The interesting questions about vigilantism really arise when property laws are *partially unjust*.

We can say that property laws are *partially unjust* when (a) none of the sufficient conditions for *severe injustice* obtain and at least one, but not all, of the necessary conditions of *sufficient justice* obtain, or (b) when none of the sufficient conditions for either severe injustice or sufficient justice obtain (in other words, property laws exceeds neither the threshold for severity or sufficiency across each of the three dimensions).²¹ Is

²⁰ Aquinas thought that taking from another's "superabundance" in order to satisfy need was simply taking that to which the needy are morally entitled, in virtue of their need. Jeremy Waldron has said that "nobody should be permitted to use force to prevent another man from satisfying his very basic needs". Thomas Aquinas, *Summa Theologia*, 2.2, 66.7, Jeremy Waldron, *Liberal Rights: Collected Papers, 1981 – 1991* (Cambridge: Cambridge University Press, 1993), p.240 to 241.

²¹ More exactly, we can say that partial injustice obtains if and only if (1) property laws do not permit the violation of deontological constraints, but (2) some people are not morally entitled to everything that they legally possess, (3) the currency of justice is not distributed fairly enough, but nor is it unfairly distributed to a morally intolerable degree, (4) procedures are neither severely unfair nor sufficiently fair. To these, we should add: (5) the realisation of more just property laws is possible through people acting in ways in which they are capable of acting, and (6) the legal status quo with respect to property rights makes everybody better off than they would be if there were no laws.

egalitarian vigilantism permissible when property laws are merely partially unjust? A tempting response to this question is the following.

When property laws are partially unjust, the moral presumption *against* vigilantism is much stronger than when they are severely unjust but much weaker than when they are sufficiently just. And the case *for* vigilantism is weaker than when property laws are severely unjust, but stronger than when they are sufficiently just.

The best we can do is to try to judge whether vigilantism against partial injustice is justified or permissible on a case-by-case basis. We can compare the strength of the reasons for complying with the unjust law with strength of the reason for disobeying the law, and act in accordance with the stronger reason, other things being equal. But we should not expect to find any general underlying principles which tell us whether or not vigilantism is permissible in a given case.²²

I agree with the general spirit of this claim, but we are not straightforwardly entitled to it. There are at least two additional considerations in partially unjust situations which strengthen the moral presumption against vigilantism.

2.5 Two Moral Obstacles to Vigilantism

The first presumption is established by what I call *the argument from fair democratic procedures*, or *the democracy argument* for short. The argument has two aspects. First, it says that citizens who respect legal property rights are morally entitled to demand similar acquiescence on the part of others. Second, it says that that vigilantism is a *disproportionate* response to partial injustice when there are democratic alternatives which could serve the same aims.

The second presumption is established by what I call *The Kantian Argument*. This argument says that there is a decisive deontological reason to refrain from vigilantism because it could not be universalised: if everybody resorted to vigilantism, or believed vigilantism to be permissible, then egalitarian vigilantism would make things worse, rather than better.

²² A defence of a similar view applied to all forms of resistance to unjust social institutions can be found in Gabriel Wollner's *On Permissibly Resisting Injustice* (Working Paper).

I will present each of these arguments in turn, and argue that, although they succeed in ruling out some forms of vigilantism that would be morally permitted in the face of severe injustice, they do not comprehensively rule out *all* forms of vigilantism in the face of partial injustice. In the course of rejecting the arguments, I will argue that there are, in fact, some underlying considerations which govern the permissibility of vigilantism against the backdrop of partial injustice.

2.5.1 The Democracy Argument

Some people object to vigilante redistribution in a democratic society because it unjustly imposes one's own idea of distributive fairness onto others who reasonably disagree. They think that, when they conflict, the demands of *procedural* fairness override the demands of *distributive* fairness. Thomas Christiano writes

Citizens who skirt democratically made law act contrary to the equal right of all citizens to have a say in making laws when there is substantial and informed disagreement. Those who refuse to pay taxes or who refuse to respect property laws on the grounds that these are unjust are simply affirming a superior right to that of others in determining how the shared aspects of social life ought to be arranged [...] only by obeying the democratically made choices can citizens act justly [...] Each citizen has a right to one's obedience.²³

On this view, procedures gain overriding moral credentials from their ability to resolve reasonable disagreements in a publically egalitarian way, and vigilantes act with a moral abandon that is incompatible with each person's claim-right that others acquiesce to the outcomes of fair procedures. Law abiding citizens can raise the following complaint against vigilantes.

What right do *you* have to act above the law? Most of us respect legal property entitlements, even though they do not perfectly reflect what justice requires, and you owe it to us do likewise. The distribution of legal entitlements is not, after

²³ Thomas Christiano, *The Constitution of Equality* (Oxford: Oxford University Press, 2008), pp.250-252. Note that one can think that democratically made decisions have moral credentials which, provided that they meet other conditions, make them permissibly enforceable. But one can believe that it is permissible to coercively enforce democratically made decisions and also think that it is permissible to disobey them. Christiano's view is challenged by Gerhard Øverland and Christian Barry, "Do Democratic Societies Have a Right to Do Wrong?" in *The Journal of Social Philosophy*, 42.2 (2011), 111-31.

all, unjust to a severe extent. We all owe it to each other to respect the outcomes of our democratic procedures.

These claims are plausible when our procedures are reasonably but imperfectly just. But I doubt that anyone believes that *severe* injustice imposed by fair procedures gives rise to claim-rights which bind everyone to obey the law. So does this argument have any force when our procedures are partially unjust? I think there are two problems with it. The claim can be understood in at least two ways.

First, it can be understood as the claim that each individual has a claim-right against *everyone* that they do not interfere with legal property entitlements (and laws in general) that followed from partially fair democratic procedures. But that is too restrictive. It rules out some forms of vigilantism which are intuitively permissible, and seems to rule out civil disobedience in general. Consider again the mass trespass of Kinder Scout. Should we believe that the protesters impermissibly acted contrary to the claim-rights of other citizens? I think that is implausible.

Second, it may be understood as the claim that individuals *who themselves* comply with the outcomes of partially fair democratic procedures hold a claim-right that others do not interfere with their legal property entitlements. But this is insufficiently comprehensive. While it establishes a presumption against vigilantism which is carried out at the expense of innocent, law-abiding people (even if they are unfairly much better off than others) it does not rule out the egalitarian expropriation of things legally owned by individuals or groups who perpetuate injustice to their own advantage. To develop this idea, consider two ways in which individuals can be situated with respect to procedural injustice.²⁴

Let's say that *perpetrators* of injustice are those who perpetuate injustice in ways which are legally permissible but morally wrong. Consider groups which intentionally corrupt procedural fairness by using their wealth and power to exert an unfair influence on democratic decisions or intentionally avoid making a fair contribution to society, for example, by avoiding taxation. *Innocent bystanders* to injustice are people who are not culpably responsible for the perpetuation of injustice and do not have an unfair say over how laws are decided or democratic decisions are made.

²⁴ These categories do not exhaust the logical space of ways in which individuals can be situated with respect to injustice. For example, we could also distinguish between culpable and innocent beneficiaries of injustice, and victims of injustice.

The democracy argument, if successful, establishes that vigilantism which targets innocent bystanders would violate their claim-rights that others do not interfere with legal entitlements that followed from partially fair procedures. I do not, however, believe that vigilantes act contrary to the rights of culpable perpetrators of injustice. How could perpetrators of injustice who do not themselves respect the outcomes of fair procedures, or undermine the fairness of procedures themselves, hold a claim-right that others do so?

This idea derives further justification from its ability to explain the difference between the following two cases from the actual world. Both of the cases involve fraud, which one may regard as an additional moral taint that makes vigilantism much harder to justify, but since this factor is held equal between the cases, it plays no part in explaining the differences between them.

In 2010, an NHS bereavement services advisor was convicted of stealing over £750,000 from the estates of dead patients by forging documents in order to obtain bequests. Allegedly, £250,000 of the misbegotten funds was given to charitable causes, a further £10,000 was given to a homelessness project and, according to the defendant, some was directly handed out to prostitutes, asylum seekers, refugees, and the homeless.²⁵

Contrast that case with the following one; in 2003, Rev. Ozell Clifford Brazil of Los Angeles was sentenced to 41 months in a federal prison and ordered to pay \$716,179 in restitution on seven counts of federal student financial assistance fraud. He assisted hundreds of disadvantaged college applicants in filing fraudulent student aid forms involving understatement of family income or false claims to be orphans or wards of the court. Due to his crimes, many disadvantaged people were able to afford to go to college.²⁶

I am strongly inclined to see the latter act of fraud as more justifiable than the former (even if they are both ultimately impermissible due to the deception involved). In the first case, while it would have been permissible for a legitimate government to confiscate the estates through inheritance taxation and use the money to help the worse off, it was wrong to take the money fraudulently. Doing so betrayed the trust of (what we assume to be) innocent people (as well as the dereliction of professional duties).

²⁵ Some of the money was used for personal enrichment: the perpetrator bought two flats, a narrowboat, and a musical organ. The case was reported in *The Daily Telegraph*, 2nd May, 2010.

²⁶ Reported in *The Los Angeles Times*, 22nd February, 2003. As is usually the case in examples from the real world, the case is tainted by additional details, such as the fact that he charged prospective college applicants to attend seminars in which he counselled them to cheat the system.

By contrast, I think the latter case was much more justifiable. I would posit that this is *primarily* explained by the fact that the fraud was carried out at the expense of the state, rather than at the expense of individuals, combined with the thought that, in this example, the state is culpably responsible for the inequality which the fraud sought to ameliorate. It may also be partly explained by the fact that I am inclined to regard the crime as “victimless”, although it may have had the side-effect of making people worse off than they *would have been* if they missed out on state aid as a result.

I think that the democracy argument rules out vigilantism which targets innocent bystanders to injustice, who merely play by the rules. It does not, however, rule out vigilantism against perpetrators of injustice, because they lack the right to demand that others refrain from interfering with their legal property entitlements.

The second aspect of the democracy argument says that although people may sometimes justifiably resort to vigilantism as a last resort, it is a *disproportionate* response to partial injustice when one can turn to democratic or otherwise legal means instead. To illustrate, assume, for the sake of argument, that western democracies in the actual world are partially unjust, rather than severely unjust, and recall again the Greek bank robberies. One could say the following:

Yes, the distribution of wealth in Greece and other Western democracies is unjust. But we live in a democracy. Although our democratic procedures are less than fully fair, they are not severely unfair either. You should try to change property laws through them, rather than taking the law into your own hands. Vigilantism may be a last resort. But it’s a disproportionate form of resistance in a democratic society.

I think that this argument is insufficiently comprehensive. It is very plausible when democratic and legal forms of recourse are sufficiently fair or have a non-negligible probability of being effective. But it is less plausible when the degree of procedural unfairness is large enough that legal forms of have no chance, or very little chance, of being effective. Consider the following example as an illustration.²⁷

²⁷ Exploring these problems through simplified cases raises a dilemma. On the one hand it helps to strip away realistic details in order to cleanse our judgements of bias in favour of the status quo. But on the other hand, stripping away realistic details may hinder our capacity to form judgements which are practically relevant. In other circumstances, this would not matter, but when the purpose of the enquiry is to identify conclusions that are of practical relevance, abstracting away from realistic cases may be counterproductive. Nevertheless, I think that the example is sufficiently informative to be worthy of consideration.

Imagine that five people wash ashore on an unowned and uninhabited desert island which contains a fixed amount of resources. Assume, for the sake of argument, that distributive fairness requires that the island's resources be equally distributed. Suppose that the islanders all agree to hold a simple majority vote to determine how to distribute the island's resources.

All five islanders believe that justice requires an equal distribution, but three of the islanders collaborate and vote to unequally distribute the island's resources in their favour, granting them access to the large sandy beaches, pineapple plants and coconut trees (call them the rich islanders). The remaining islanders (the poor islanders) are confined to small, rocky beaches and subsist on seaweed from rock pools, but nevertheless have enough resources to live a decent life. Suppose that the unequal distribution is coercively enforced through threat of punishment and that, at periodic elections, the status quo is perpetually reinforced by the voting majority of the three rich islanders.

The aspect of justice concerned with procedural fairness is sufficiently realised on the grounds that it was decided by a majority vote that each islander expressly consented to and willingly participated in. The aspect of justice concerned with distributive fairness is not ideally realised but nor can it be described as severely unfair. In this case, it is plausible that the poor islanders are morally bound to the outcome of reasonably fair democratic procedures that they willingly agreed to and participated in. We can conclude that when procedural fairness is satisfied to this extent, there is a very strong presumption against vigilantism.

But to test the democracy argument in full, we need to identify a variation on the case in which the procedural aspect of justice is *partially* unfair. Let's add some additional details to the case to make it more closely resemble a realistic one which warrants the label of partial unfairness. Imagine that two of the islanders are replaced by two new islanders. The two new islanders do not expressly agree to the form of government and are unable to leave the island.

Imagine that there is a permanent majority in favour of unequally distributing the island's resources because it is in the rational self-interest of any three islanders to collude in favour of it. Although there are democratic procedures in place which offer the *possibility* of overturning the distributive status quo, egalitarian reform is, in practice, blocked. The form of government does not enjoy the express consent, or tacit consent by residence, of the whole population. Nevertheless, each islander willingly

participates in democratic votes on equal terms. I think we can say that the democratic procedures are *partially unfair*.

Suppose further that the identity of the rich and poor changes periodically. There are always three rich islanders and two poor but the composition of each group changes.²⁸ However, one of the new islanders is *always* among the poor. Call him the *unlucky islander*. Now consider the following case.

Harmless Poacher

At night, the unlucky islander sneaks onto the rich part of the island while the other islanders sleep. He trespasses on the sandy beaches and swipes fruit that will be partially, though not wholly, replenished. The covert raids go some way to reducing inequality of both resources and welfare between the islanders. The other islanders never find out about the trespass and never notice that the fruit has gone missing.

I am strongly inclined to believe that harmless poaching by the unlucky islander is permissible. Bear in mind that this is not equivalent to saying that the majority of the islanders had no right to unfairly distribute benefits and burdens in their own favour. It may have been wrong for them to do so, but they may have a right to do wrong which is grounded in the partial fairness of the voting procedure. Even if they had such a right, and even if they may permissibly enforce the outcome of their decision, it can still be permissible for the unlucky islander to resort to vigilantism. Those possibilities are all consistent.

The purpose of these examples is to illustrate that even in a reasonably “pure” case, it is not completely clear that vigilantism is prohibited when democratic and other legal forms of recourse are highly unlikely to be effective. Some might be unconvinced by these sparse examples so let me try to strengthen my claim by considering a real world example.

In August 2012, Juan Manuel Sánchez Gordillo, mayor of the small Andalusian town of Marinaleda, led farm labourers into supermarkets to expropriate food and then distributed the spoils to local food banks. This act of vigilantism ticks all of the boxes to make it more easily justified. It targeted a group entity without violence or coercion and it is unlikely that any individuals were made worse off as a result, or, if they were (for

²⁸ This allows us to set aside the problem of a permanently disadvantaged minority who did not expressly consent to the form of government, which can be described as an instance of *severe* procedural injustice.

example, due to supermarkets increasing their prices to recoup the costs of theft), any burdens were widely dispersed. It was also both a redistributive and an expressive protest. This makes it a good test case for the permissibility of vigilantism in partially unjust situations. Now consider the backdrop.

Unemployment [was] at 25% nationally (higher than Greece), 34% in Andalusia and 53% for 16-to-24-year-olds; [...] the collapse of the construction industry [...] left 800,000 empty homes, and, [in] May, the 8 million-strong *indignados* protest movement, a forerunner of Occupy, announced its total lack of faith in parliamentary democracy to solve any of these problems.²⁹

In a range of realistic situations like this one, there are obstacles to legal and procedural forms of addressing injustice which render them highly unlikely to be effective. In these situations, egalitarian vigilantism cannot be described as disproportionate on the grounds that better means to the same aims are available. While some varieties of vigilantism, for example, violent and coercive ones, may remain disproportionate response to inequality in a partially unjust society; they are disproportionate when and because peaceful and non-coercive forms of vigilantism are available.

One might respond to this claim by exploiting the ambiguity in my definition of “partial” injustice, and suggesting that the society in this case should count as severely unjust, rather than partially unjust. One might maintain that once unfair inequality exceeds a certain degree, or democratic procedures become inert and unable to bring about effective change, the procedures deserve the label of severe injustice. I would be happy to concede this point and accept a more expansive set of conditions for severe injustice, since my argument here is primarily intended to demonstrate that vigilantism is not disproportionate when democratic or other legal alternatives are highly unlikely to be effective. Nothing, in particular, depends upon the exact parameters of severe and partial injustice.

In summary, the democracy argument rules out vigilantism in the range of partially unjust situations in which there are legal alternatives to vigilantism which have a reasonable probability of being effective. I do not think, however, that it rules out vigilantism in the range of partially unjust situations in which there is little to no probability of legal and democratic means proving effective.

²⁹ Reported in *The Guardian*, 12th August 2012.

2.5.2 The Kantian Argument

One might maintain that even if my replies to the above argument succeed, there is still a decisive deontological reason to refrain from resorting to vigilantism against partial injustice, which can be roughly encapsulated in the slogan: what if everybody did it? We could say

It is wrong to resort to vigilantism because we could not rationally will it to be true that everyone did so (or, alternatively, because we could not rationally will it to be true that everyone believes such vigilantism to be morally permitted).

One might maintain that we could not rationally will it to be true that everybody resorted to vigilantism, or believed it was permissible to turn to vigilantism, because this would lead to widespread disorder which would make things worse, rather than better. But this claim, applied generally to all kinds of egalitarian vigilantism, is too restrictive. It would incorrectly render the mass trespass of Kinder Scout impermissible.

I agree that if extreme forms of vigilantism, like armed robbery, were pervasive, then it would be highly likely to lead to widespread disorder which would make things worse, rather than better. But if a handful of people covertly carried out acts of vigilantism like the *Harmless Hacker* and *Harmless Poacher* cases, the effects would be good. I think it therefore faces what Parfit calls the *Threshold Objection*.

In some cases [...] whether some act is wrong depends on how many people act in this way [...] if too many people acted on this maxim, these people's acts would have bad effects, but when fewer people act on this maxim the effects are neutral or good [...] [T]hough such acts would be wrong if too many people acted on this maxim, when fewer people act on this maxim such acts are permissible, and may even be morally required [...] [M]ost of us could not rationally will it to be true that everyone acts on these maxims.³⁰

Vigilantism is only justified when, in tandem with what everyone else does, it achieves the effect of reducing unfair inequality without making things worse, all things considered. At best, I think this argument establishes that it might be good if most people *believe* redistributive vigilantism to be impermissible, even though it is sometimes, in fact, permissible.

³⁰ Derek Parfit, *On What Matters* (Oxford: Oxford University Press, 2011), p.308.

2.6 Conclusion

I defined egalitarian vigilantism as possessing or redistributing money or property which legally belongs to somebody else, without their informed consent, and without providing compensation. I pointed out that flouting property laws in this way is one way of reducing unjust inequality. I said that although there are presumptions against such vigilantism, it is also true that, under realistic conditions, we are often morally permitted to break the law, and that people sometimes lack moral rights to what they legally possess.

I argued that the most justifiable acts of vigilantism are those which, without violence or coercion, expropriate money, food, or disused real estate from group entities like banks, corporations, and the state, where any individual losses are indirect and widely dispersed.

After that, I distinguished between sufficiently just property laws, severely unjust property laws, and partially unjust property laws. Unsurprisingly, I argued that egalitarian vigilantism is very difficult to justify when property laws are sufficiently just but is fairly easy to justify when property laws are severely unjust.

When societies are partially unjust, there are additional presumptions against vigilantism which imposes property losses on innocent bystanders to injustice and when there are legal forms of resistance to injustice which have a non-negligible probability of proving effective. But there is no overriding presumption against the expropriation of property from groups which culpably perpetrate injustice when the prospects for reducing inequality through democratic or otherwise legal means are remote. So, if you're an egalitarian, why not rob a bank?

3 OPTIMAL EQUALITY AND FORCED LABOUR

3.1 Introduction

G. A. Cohen argued that we could bring about a justly egalitarian society without sacrificing economic efficiency or legal freedom of occupational choice.¹ He said that such a society could be realised if people were motivated by an egalitarian ethos to work in productive occupations for an egalitarian wage. But two counterarguments purport to show that Cohen's vision is incoherent, or, along with auxiliary premises, morally tyrannical.

The first counterargument, which I call *the efficiency dilemma*, says that a *merely egalitarian* ethos would be insufficient to induce people to freely choose productive occupations in order to realise a Pareto optimal level of equality. Realising a Pareto optimal level of equality would require both an egalitarian ethos *and* a *Paretian ethos* (or, as I shall call it, a *productive ethos*).² Accordingly, Cohen faces a dilemma. First, a society with a merely egalitarian ethos would not bring about a Pareto optimal level of equality, leaving everyone worse off than they could otherwise be. Second, if justice requires an egalitarian ethos *and* a productive ethos, then justice is not merely equality, contrary to what Cohen claims.

The second counterargument, which I call *the forced labour objection*, says that Cohen's claims about justice and labour, along with auxiliary premises, commit him to the unpalatable view that a just society could contain legally coercive job allocation (*forced labour*, for short).³ This is because, so the objection goes, Cohen lacks the normative resources to distinguish between a society in which distributive justice is legally enforced with respect to money and a society in which distributive justice is legally enforced with respect to labour. But forced labour is intuitively inimical to justice.

¹ Cohen, *Rescuing Justice and Equality*, Chapter Five.

² Patrick Tomlin, "Internal Doubts About Cohen's Rescue of Justice", *The Journal of Political Philosophy* (2010), 18.2, 228-47, and Jonathan Quong "Justice Beyond Equality", *Social Theory and Practice*, 36.2 (2010), 315-40.

³ Cécile Fabre, "Distributive Justice and Freedom: Cohen on Money and Labour", *Utilitas*, 22.4 (2010), 393-412, Gerald Lang, "Rawlsian Incentives and the Freedom Objection", *The Journal of Social Philosophy*, 47.2 (2016), 231-49 (pp. 241-247), and Michael Otsuka, "Freedom of Occupational Choice", *Ratio*, 21.4 (2008), 440-53, (p.446).

In this chapter, I argue that, although each objection draws attention to ambiguities in Cohen's arguments, both can be overcome in ways that strengthen Cohen's claims without abandoning the egalitarian spirit of his arguments.⁴ The chapter is divided into three main parts. In the first part, I explain Cohen's claims in a bit more detail. In the second part, I set out the efficiency dilemma and provide two arguments to show that, with some revisions, Cohen's claims can be defended from both horns of the dilemma. In the third part, I set out two formulations of the forced labour objection and provide several lines of argument in reply which strip the objection of rhetorical force.

3.2 Equality, Pareto Optimality, & Freedom of Occupational Choice

In *Rescuing Justice and Equality*, Cohen claimed that distributive justice is realised by equality of access to advantage, where advantage includes both resources and welfare. But under realistic conditions, the full realisation of equal access to advantage would typically require sacrificing either Pareto optimality or legal freedom of occupational choice.⁵

This is because in an equal society there would be no inequality creating incentives to induce people to efficiently allocate their labour by making good use of their productive talents. And if people do not freely make good use of their productive talents, then the only way to realise a Pareto optimal level of equality would be to legally conscript people into doing so, in violation of legal freedom of occupational choice. As a result, under realistic conditions, we must typically choose to sacrifice one of equality, Pareto optimality, or legal freedom of occupational choice.

But Cohen pointed out that these trade-offs are typically only required because people are unwilling to act in ways in which they are capable of acting. Cohen argued that we

⁴ I do not provide a complete defence of Cohen's claims. There may be other objections which prove successful, but I set them aside. For example, some have objected that an egalitarian moral obligation would constrain autonomy. See Paula Casal, "Occupational Choice and the Egalitarian Ethos", *Economics and Philosophy*, 29.1 (2013), 3-20. Others have objected that Cohen's ethos is implausibly morally demanding. See Titelbaum, "What Would a Rawlsian Ethos of Justice Look Like?" For a wider discussion of the productive requirements of justice, see Lucas Stanczyk, "Productive Justice", *Philosophy and Public Affairs*, 40.2 (2012), 144-64.

⁵ A distribution is Pareto optimal if nobody can be made better off without making at least one person worse off (if there is no Pareto superior distribution available). One distribution, A, is weakly Pareto superior to another, B, if at least one person is better off and nobody is worse off in A than B, and strongly Pareto superior if everybody is better off in A than B.

could bring about a Pareto optimal level of equality without sacrificing legal freedom of occupational choice if people were morally motivated by an egalitarian ethos to freely make good use of their productive talents.

According to Cohen, acting in accordance with an egalitarian ethos would not impose intolerable demands on people because the benefits and burdens of work are, along with income, goods to be distributed in accordance with distributive justice. Accordingly, in a just society, the benefits and burdens of work would be fairly distributed as part of the overall distribution of access to advantage. In such a society, people are asked to make “good use” of their productive talents through their labour, including their choice of occupation, until doing so asks them to bear an unfair burden in comparison to others; that is, a burden which would worsen inequality beyond the permissible range of inequality which is granted by the existence of a personal prerogative to refrain from doing what equality optimally requires.⁶

3.3 The Efficiency Dilemma

Jonathan Quong and Patrick Tomlin have independently pointed out that that a merely *egalitarian* ethos would not induce people to bring about a Pareto optimal level of equality.⁷ *Equality*, by itself, would typically be consistent with people choosing any

⁶ Cohen draws upon Joseph Carens’s ideas when proposing the egalitarian ethos. However, the two accounts differ in significant respects. For one, Carens is concerned to establish how an equal distribution of *income* could be maximised under feasible conditions whereas Cohen is concerned to establish how an equal and Pareto optimal distribution of *advantage* could be brought about. Second, Carens does not include labour burdens in his account. Third, Carens is concerned to show that production and distribution are independent and the amount which is produced need not be influenced by the principle of distribution which is in place. For Cohen, production and distribution are not independent.

Carens’s view is set out in Joseph H. Carens, “The Egalitarian Ethos as a Social Mechanism” in *Distributive Justice and Access to Advantage* ed. by Alexander Kaufman (Cambridge: Cambridge University Press, 2014), “An Interpretation and Defence of the Socialist Principle of Distribution”, *Social Philosophy and Policy*, 20.1 (2003), 145-77, and “Rights and Duties in an Egalitarian Society” in *Political Theory*, 14.1 (1986), 31-49.

⁷ Tomlin, “Internal Doubts About Cohen’s Rescue of Justice”, and Quong, “Justice Beyond Equality”.

Andrew Williams has also made a similar point, using the terms *wide ethos* and *narrow ethos* in Andrew Williams, “Incentives, Inequality, and Publicity” in *Philosophy and Public Affairs*, 27.3 (1998), 225-47 (p.235). In *Rescuing Justice and Equality*, Cohen said that he thinks justice requires the wide ethos (or productive ethos), p.370.

occupation. In fact, equality alone does not even require each person to make a productive contribution to society; it merely requires that, taking into account each person's overall bundle of resources and welfare, nobody is unfairly worse off than anybody else.

An *egalitarian* ethos would induce the widespread belief across the population that each person must make choices which ensure that access to advantage is equally distributed. It would not, however, induce the widespread belief that each person has an obligation to make a productive contribution to society. Accordingly, a society with freedom of occupational choice and a Pareto optimal level of equality would require both an egalitarian ethos *and* a *productive ethos*. A *productive* ethos would induce the widespread belief across the population that each person has an obligation to make use of their productive talents in ways bring about a Pareto optimal level of equal access to advantage.

Now, it does not follow from the fact that an egalitarian ethos only induces people to realise equality that a merely egalitarian ethos *never* governs occupational choice. Equality would still govern occupation choice when some occupational choices would reduce inequality. Cohen's example of the person who can choose to become a doctor or a gardener illustrates this point.⁸ In *that* example, the agent in question is better off than the rest of the population *whichever* occupation is chosen, but *everyone else's* access to advantage is improved if the agent chooses to be a doctor, so that choosing to be a doctor rather than a gardener both reduces inequality and brings about a Pareto superior distribution.

This ambiguity over whether Cohen thought that *justice* requires a productive ethos gives rise to a dilemma. The first horn of the dilemma is as follows: if justice requires a merely egalitarian ethos, then, under realistic conditions, labour would be inefficiently distributed and Cohen would seemingly have failed to show how equality, Pareto optimality, and freedom of occupational choice could be realised together. As a result, a just society could be one in which everyone is worse off than they could otherwise be.

The second horn of the dilemma is the following: if justice requires both an egalitarian ethos *and* a productive ethos, then this is inconsistent with Cohen's claim that equality of access to advantage is both necessary and sufficient for the realisation of distributive justice. According to Cohen, Pareto optimality is not a requirement of justice. A society

⁸ Cohen, *Rescuing Justice and Equality*, pp.184 to 185.

could be justly equal, but Pareto suboptimal (even if, all things considered, we sometimes ought to bring about an unjust but Pareto optimal distribution).⁹

Quong says that his observation is significant because “once we realize that the ethos directing people to make certain occupational choices is not required as a matter of equality, then Cohen cannot claim that the ethos is a matter of justice”.¹⁰ Moreover, according to Quong, a productive ethos is beyond what justice, or morality more generally, could require and so realising Cohen’s vision would require morally supererogatory choices on the part of each person in society.

Now, as a matter of exegesis, it is not clear that Cohen did defend the claim that justice *requires* a productive ethos. Although there is no doubt that Cohen thought that justice required an *egalitarian* ethos, one might instead interpret Cohen as trying to show that, under realistic conditions, a justly egalitarian distribution could be brought about *without sacrificing* legal freedom of occupational choice or Pareto optimality. Cohen may be interpreted as defending the weaker conditional claim that *if* individuals freely chose occupations which helped to bring a Pareto optimal level of access to advantage (motivated by, for example, “principled commitment and fellow feeling”), then the trilemma, at the bar of state policy, of choosing between equality, freedom and Pareto optimality would disappear.¹¹

Ambiguity concerning whether Cohen thought the realisation of justice requires a productive ethos may be explained by the distinction between Cohen’s arguments as an internal criticism of the difference principle and Cohen’s arguments as an independent account of distributive justice. The accounts can differ in the following important respects. Formulated as an internal critique of the difference principle, Cohen holds that if people in a Rawlsian society *really* believed in the difference principle, then this would motivate them to choose occupations which maximised the economic prospects of the least well-off group in society without requiring inequality-creating incentives to do so.

⁹ It should be noted that Cohen does not rule out the possibility that an equality conserving Pareto improvements are requirement of distributive justice (in *Rescuing Justice and Equality*, pp.322 to 333). However, this possibility is clearly in tension with his insistence that equal access to advantage is sufficient for the realisation of distributive justice.

¹⁰ Quong, “Justice Beyond Equality”, p.328.

¹¹ In support of this interpretation, I would cite the following passage: “all that the [...] ethos requires is that one does not violate *equality* in career choices” *Rescuing Justice and Equality*, p.371.

More precisely, the claim is that for any equal distribution D1, if there is an unequal, strictly Pareto superior distribution, D2, made possible by at least one person working more productively for an inequality-creating incentive, then there is a further possible distribution, D3, which is both equal and strictly Pareto superior to D1 and Pareto incomparable to D2, made possible by at least one person working more productively without an inequality creating incentive. As a result, according to Cohen, a *consistently Rawlsian* society would have a productive ethos.

On this view, a society organised around the difference principle would not contain inequality-creating income incentives, other than to influence those choices which fall within the bounds of people's personal prerogative to refrain from doing what justice requires, *and* those which are *strictly* necessary (that is, necessary apart from the voluntary choices of individuals) to maximise the prospects of the least well-off. More precisely, an unequal distribution would be just if, and only if, (1) the worse off could not be made better off and (2) the inequality was not produced by people intentionally refusing to act in ways in which they are capable of acting, aside from actions which are morally protected by their personal prerogative.

Now, Cohen insists that the difference principle *is not* a principle of distributive justice. Formulated as an independent account of what a just society would look like, Cohen's claims can be interpreted differently. On such a formulation, an unequal distribution would be unjust even if it was strictly necessary to maximise the prospects of the least well-off. On this view, justice requires an egalitarian distribution but is indifferent between equality at a high level and equality at a low level. For example, a distribution could be Pareto optimal, but unjust. By contrast, a distribution could be justly equal, but Pareto suboptimal. As a result, a just society requires (only) an egalitarian ethos.

Exegesis aside, I will now assume, for the sake of argument, that Cohen endorses the following two claims.

(Egalitarian Claim) Distributive justice is equality of access to advantage.

(Productive Ethos Claim) Distributive justice requires a productive ethos.

The efficiency dilemma shows that these two claims are inconsistent. I will argue that Cohen's account can overcome both horns of the dilemma but only by rejecting one or the other of the claims. However, I argue that neither alternative is especially costly to Cohen's overall position because both routes retain the egalitarian spirit of his arguments.

3.3.1 Reply to the First Horn

A defender of Cohen could retain the egalitarian claim, relinquish the productive ethos claim, but nevertheless maintain that that equality, Pareto optimality, and legal freedom of occupational choice would typically (though not always) be co-realised under realistic conditions. Once different metrics of justice are disentangled, it becomes apparent that it would typically be in everybody's rational self-interest to make good use of their productive talents through their choice of occupation, even if justice did not require them to do so. This offers a way to overcome the first horn of the dilemma.

Consider the following illustration in which the occupational choices dictated by equality and Pareto optimality diverge. Suppose we can move from a suboptimal state of equality (D1) to either a Pareto superior state of inequality (D2) or a Pareto optimal level of equality (D3) which is Pareto incomparable to D2.

	D1	D2	D3
Karl	1	4	3
Everyone Else	1	2	3

To fill in the details, suppose Karl is employed at the Ministry of Equality and he can choose to instead work at the workers' co-operative farm. Other things being equal, Karl would prefer to work in the Ministry of Equality, but, as he is a talented farmer but an untalented clerk, his talents would be put to better productive use in the workers' co-operative.

If Karl continues to work at his desk job at the Ministry of Equality, the distribution will remain D1. If Karl works at the workers' co-operative for an inequality creating income (because he is unmoved by an egalitarian ethos, say), he would bring about D2. If he chooses to work at the workers' co-operative for an egalitarian wage, this will bring about D3.

The fact that D3 is Pareto superior to D1 is a reason to prefer it from an impartial perspective, all things considered, but there is no *egalitarian* reason to prefer it. So a *merely* egalitarian ethos would not move Karl to bring about D3 by changing occupation. Here, a merely egalitarian ethos would permit Karl to continue to work in the Ministry of Equality. By contrast, a *productive ethos* would move Karl to change his occupation in order to bring about D3. It appears, therefore, that a merely *egalitarian* ethos would fail to realise both equality and Pareto optimality.

But I believe that this conclusion is too hasty. A society with a merely egalitarian ethos could still realise a Pareto optimal level of equality under realistic conditions without sacrificing legal freedom of occupational choice. That is because it would typically be in the rational self-interest of each person to make optimal use of their productive talents when doing so would bring about an equality conserving strict Pareto improvement in the currency of justice.

We must bear in mind what Jonathan Wolff has called the *real Pareto maxim*.¹² Distributions which are Pareto suboptimal in one currency may be optimal under another currency. While a distribution may be Pareto optimal in terms of *access to advantage*, where “advantage” includes both resources, including income, and welfare, including labour burdens, it may nevertheless be Pareto suboptimal in terms of *income*.

Assume for the sake of argument that the currency of justice is income, labour burdens are excluded from the currency of justice, and Karl is unmoved by a productive ethos. Now it looks morally demanding for Karl to work at the workers’ co-operative for an egalitarian wage. For although Karl has a higher income in D3, we know that, unmoved by a productive ethos, he would prefer to work at the Ministry of Equality than to receive this additional income while working at the workers’ co-operative. Karl could only bring about a Pareto optimal level of equality of income by acting contrary to his preferences.¹³

Now assume for the sake of argument that the currency of justice is access to advantage, labour burdens are included in the currency of justice, and Karl is motivated by an egalitarian ethos. Now the picture is different. It now appears to be in Karl’s rational self-interest to bring about D3 by working at the workers’ co-operative. For in D3, Karl receives more welfare and resources than in D1, and any additional labour burdens imposed in D3 are, by stipulation, included in the calculation of Karl’s share of the currency of justice.

This point is made more starkly when we generalise from the case of a single individual choosing a job to occupational choices across an entire population. We all benefit from the efficient distribution of labour, even if that distribution allocates a job which would

¹² Jonathan Wolff, “Levelling Down” in *Challenges to Democracy: The PSA Yearbook 2000*, eds., Keith Dowding, Jim Hughes, and Helen Margetts (Macmillan, 2001), p.18.

¹³ A requirement to bring about a Pareto optimal distribution of income would clearly run contrary to equality of access to advantage. Indeed, Cohen says that this “conjures up a nightmare scenario in which the duty I advocate is interpreted as making the productive work as much as they can to make the wages of the less well paid as high as possible”. (*Rescuing Justice and Equality*, p.402)

not be our first choice, other things being equal. If people were motivated *both* by a moral commitment to equality *and* their rational self-interest, an equal and Pareto optimal distribution could contingently be brought about in a society with an egalitarian ethos.

To illustrate, suppose Vladimir, Leon, and Mikhail are on a camping trip and must decide how to divide their labour. Vladimir prefers to fish, Leon prefers to build campfires and Mikhail prefers to hunt for rabbits. Unfortunately they are all unskilled at their respective preferred tasks and the overall camping experience is generally underwhelming. However, Vladimir is talented at hunting for rabbits, Leon is an excellent fisherman, and Mikhail is skilled at building campfires. By reallocating their labour, they all enjoy a much better camping experience, all things taken into account, despite the fact that they forego their preferred tasks.

Each person can offer the following justification for choosing an occupation which makes an efficient contribution to the social product. “It is in our collective and individual interest(s) to distribute our labour in accordance with our talents and skills. Although we may have to forego what would be first occupational preference, other things being equal, we are all much better off in terms of access to advantage when labour is efficiently distributed”.

Two objections may be raised to this line of argument. First, one might object that if it is in the rational self-interest of Karl to bring about D3 rather than D1, this renders the egalitarian ethos redundant: we simply don’t need an ethos in order to bring about the optimal outcome. But this is mistaken. An egalitarian ethos would still be required in order to bring about D3 *rather than* D2. A commitment to both equality and self-interest would make possible the optimal equal outcome of D3.

Second, one might object that it is often *not* in the rational self-interest of a person to make good use of their productive talents when they would prefer to do a less productive job for the same income. For example, Jonathan Quong writes: “the relatively small increases in income that the talented would receive [by making Pareto improving occupational choices in an equal society] might not compensate for having to make radically different career choices to the ones the talented would prefer to make”.¹⁴

But if an increase in a talented person’s income *did not* compensate for the decrease in welfare incurred by foregoing their first occupational preference, and labour burdens are

¹⁴ Jonathan Quong, “Contractualism, Reciprocity and Egalitarian Justice”, *Politics, Philosophy, and Economics*, 6.1 (2007), 75-105, (p.86).

included in the currency of justice, then choosing the more productive occupation would *not* lead to an equality conserving strict Pareto improvement in the currency of justice; D3 would not be Pareto superior to D1. The absence of sufficient compensation would render the two situations Pareto incomparable. By contrast, if choosing the more productive occupation *did* constitute an equality conserving strict Pareto improvement, then it would generally be in the rational self-interest of the talented person to choose the more productive occupation.

That being said, while the occupational choices dictated by optimal equality and rational self-interest would often correlate, people's preferences could sometimes diverge from what a Pareto optimal level of equal access to advantage requires. So while D3 may be Pareto superior to D1 in terms of access to advantage, it may be Pareto incomparable to D1 in terms of actual preference satisfaction.

My argument would have a more comprehensive reach if we assume that the currency of justice is *opportunity for welfare*, rather than access to advantage, and we assume that it is always in the rational self-interest of people to increase their own opportunity for welfare. But I will not explore this thought any further here. The important point is that a merely egalitarian ethos would not be strictly *inimical* to Pareto optimality.

3.3.2 Reply to the Second Horn

A second way for Cohen to overcome the efficiency dilemma is to avoid the second horn of the dilemma by retaining the productive ethos claim, but relinquishing the egalitarian claim. Instead, Cohen could, and, in my view, should claim that distributive justice is a second order property which combines the values of equality and welfare (as well as community and legitimacy), while attaching lexical priority to equality. This offers a way to overcome the second horn of the dilemma while retaining the egalitarian spirit of Cohen's vision.

The claim that justice is merely distributive equality has several shortcomings, because not only does it imply that an equal society in which everybody is very badly-off is more just than an unequal society in which everybody is very well-off, it also implies that a society in which everybody is equally badly-off is no less just than a society in which everybody is equally well-off. Moreover, it implies that, when an equal distribution is unavailable, we should be indifferent, as a matter of justice, between a Pareto optimal level of inequality and a Pareto suboptimal state which is unequal to the same degree. This is not to say that Cohen's egalitarianism is flawed as a view about the comparative

fairness of *outcomes*, but as a complete account of *distributive justice*, Cohen's egalitarianism has these counterintuitive implications.

Abandoning the egalitarian claim might tempt us towards the view that equality is only one aspect of distributive justice, to be balanced alongside distinct, potentially conflicting, values. Accordingly, we might say that justice is a complex, second order value which includes the optimal balance of relevant first order values, like welfare, efficiency, and freedom. For example, we could say that (3,2) is *more just* than (1,1), because everyone is better off, even though it is worse, in one respect, because the former includes some unfairness whereas the latter does not.¹⁵

However, if we were to take this route, we would be abandoning the egalitarian spirit of Cohen's arguments. Cohen insists that moving from a state of equality to a Pareto superior state of inequality creates distributive injustice, even if we ought to bring about such a state, all things considered. But there is a further intermediate position available which allows us to insist on the primary importance of equality to justice, while avoiding some of the unpalatable implications outlined above.

On Cohen's view, we can seemingly determine whether a distribution of benefits and burdens is just on its *internal aspects*, where those internal aspects are the distributive pattern that obtains. In other words, we do not need to know what other distributions are feasible to determine whether or not a distribution is just.¹⁶

We should, as egalitarians, instead adopt an alternative view, according to which we can only determine whether a distribution is just by comparing it to possible alternatives. One might maintain that equality is *necessary* for justice, so that an equal distribution will always be more just than an unequal one (even if the unequal distribution is strongly Pareto superior to the equal one). But we might also say that equality is sometimes *insufficient* for justice, when a Pareto superior state of equality is available.

We should hold that unfair inequality is *always* unjust, but *equality* is also unjust when a Pareto superior state of equality is possible through people acting in ways in which

¹⁵ This response is suggested by Tomlin (in "Internal Doubts About Cohen's Rescue of Justice", p.242) and resembles claims made by Michael Otsuka in "Equality, Ambition, and Insurance", *Proceedings of the Aristotelian Society Supplementary Volumes*, 78 (2004), 151-66, p.164.

¹⁶ The distinction between internal aspects and essential comparability is identified in Temkin's *Rethinking the Good*, Chapter 7.

they are capable of acting.¹⁷ After all, the claim that distributive justice is a complex value with different aspects does not tell us anything about the respective strength of the different values which comprise it.

This revision is supported by two powerful intuitions. First, given two outcomes containing the same population and roughly the same amount of goods, the more equal outcome is more just. Second, given two outcomes containing the same population which are equivalent from an egalitarian perspective (because both are perfectly equal or contain the same degree of inequality), in which one is Pareto superior to the other, the Pareto superior outcome is more just.

Amending Cohen's claims about equality and justice is consonant with remarks Cohen makes in some later works, in which he seems to soften his stance on the claim that equality of access to advantage is the only aspect of distributive justice. In *Why Not Socialism?*, he suggested that the egalitarian principle may need to be supplemented with a principle of *community*, because an egalitarian principle would allow large inequalities to be produced by responsible choices and option luck, and such distributive inequalities may serve to undermine fraternal social relations.¹⁸ Moreover, in "Fairness and Legitimacy in Justice", he suggested that fairness and *legitimacy* may be two distinct aspects of distributive justice, because inequality brought about by option luck and bilateral gifts may create a distribution which is unfair but legitimate, where legitimacy is the property possessed by a distribution about which nobody has the right to complain.¹⁹

The proposed revision to Cohen's argument allows us to claim that a productive ethos *is* required as a matter of justice, without abandoning the egalitarian spirit of Cohen's vision, while overcoming the shortcomings of equality of access to advantage as a complete theory of distributive justice.

At this point, Quong might object that if the productive ethos *is* supposed to be required by justice, then this simply creates a new problem for Cohen. For a productive ethos is beyond what any theory of justice, or morality in general, could plausibly require. But as

¹⁷ Indeed, Cohen contemplates this possibility, in *Rescuing Justice and Equality* p.322, and Quong offers it on Cohen's behalf, in "Justice Beyond Equality", p.331.

¹⁸ G. A. Cohen, *Why Not Socialism?* (Princeton: Princeton University Press, 2009), Chapter 2.

¹⁹ G. A. Cohen, "Fairness and Legitimacy in Justice, And: Does Option Luck Ever Preserve Justice?" in *On The Currency of Egalitarian Justice, and Other Essays in Political Philosophy* ed. By Michael Otsuka (Princeton: Princeton University Press, 2011).

I pointed out in my reply to the first horn of the dilemma, this is not so. For, if the currency of justice is access to advantage, and labour burdens are included in the currency of justice, then it is generally in the rational self-interest of each person with egalitarian beliefs to work productively for an egalitarian income.

3.3.3 Summary

We are now in a position to take stock of the potential replies to *the efficiency objection*. First, we can claim that a society with an egalitarian ethos *could* realise a Pareto optimal level of equality under realistic conditions, even if justice does not require it. If the currency of justice is access to advantage, labour burdens are a part of the currency of justice, and people are motivated by a belief in equality in addition to their rational self-interest (which coincides with a concern for the welfare of others) then a Pareto optimal state of equality could typically be brought about.

Second, we could dispense with the claim that justice is merely equality and instead adopt the view that justice is a second order property encompassing a variety of values, including equality, welfare, community, and legitimacy, with equality enjoying lexical priority. Then we are in a position to claim that justice requires a productive ethos, but retain the egalitarian spirit of Cohen's vision.

3.4 The Forced Labour Objection

While most egalitarians would welcome the removal of inequality through coercive taxation, most would condemn the removal of inequality through forced labour. But some have argued that, by the lights of Cohen's claims, there is no moral difference between the legal enforcement of equality with respect to money and the legal enforcement of equality with respect to labour; both involve coercing the unwilling into bringing about equality.²⁰ If this is so, one may regard this claim as a reason to revise

²⁰ Note that this differs from Robert Nozick's claim that redistributive taxation *is*, or is "morally on a par with", forced labour, in *Anarchy, State, and Utopia* (Oxford: Blackwell, 1974), p.169. The objection to Cohen is that, given his endorsement of coercive egalitarian taxation, he cannot *object* to the use of forced labour in the name of equality.

one's views on the permissibility of coerced labour, or regard the argument as a *reductio ad absurdum* of Cohen's vision of a just society.²¹

For the sake of argument, I assume that this objection, if successful, would constitute a *reductio ad absurdum*. In what follows, I will distinguish between two different ways of posing this objection more precisely, and argue against each in turn. The first thing we can ask is whether Cohen has the normative resources to oppose the realisation of equality through forced labour, *all things considered*. The second thing we can ask is the more difficult question of whether Cohen has the normative resources to oppose the realisation of equality through forced labour *as a matter of justice*.

3.4.1 The First Formulation

Cohen tried to morally distinguish between coercive taxation and forced labour by pointing to the fact that the latter requires the state to intrusively manipulate and command people through the use of intimate truths about them, while the former does not.

But Michael Otsuka has objected that this way of distinguishing between coercive taxation and forced labour does not adequately explain the moral repugnance of the latter, on the grounds that if, contrary to fact, people's talents and preferences were transparent to all, it would still be intuitively objectionable for the state to force people into occupations which bring about optimal equality, even though doing so would not require intrusion into people's private thoughts.

Moreover, Otsuka says that, on Cohen's view, it is difficult to see why it would be wrong to coercively override a persons' unwillingness to do what egalitarianism requires because "based as it is on mistaken ethical convictions and damaging as it is to the least well off, this recalcitrance does not appear to ground a decisive objection to the state's so commanding him."²²

Otsuka says that self-ownership largely provides the grounds for drawing a clear moral disparity between coercive taxation and forced labour, because, in contrast to property rights in oneself, people do not enjoy stringent rights of ownership over an inegalitarian

²¹ Cécile Fabre has stated that she believes that justice does require the legal conscription of people into certain occupations, in *Whose Body is it Anyway? Justice and the Integrity of the Person* (Oxford: Oxford University Press, 2006), Chapter 3: "A Civilian Service".

²² Otsuka, "Freedom of Occupational Choice", p.446.

share of natural resources or income that can be derived from those resources (including in combination with one's labour).²³ Cohen, however, denies that we have rights of self-ownership over our bodies and labour and therefore cannot appeal to this idea to distinguish between taxation and forced labour.²⁴

So the challenge for Cohen is the following: according to Cohen, both income and labour burdens are part of the currency of justice, and justice requires us to bring about equality. Cohen says that it is permissible for the state to coercively enforce a distribution of income which realises equality of access to advantage. If, as Cohen believed, we do not have rights of self-ownership, and the fact that forced labour requires the acquisition and use of private information about people is not a decisive reason to reject it, on what normative grounds can we advocate forcing someone to relinquish some of their income for the sake of equality but oppose forcing someone to sacrifice some of their labour for the sake of equality?

I believe that we can answer this question in three ways that make Cohen's position more palatable, without appealing either to rights of ownership over one's labour or the fact that a state policy of forced labour would require intrusion into people's private thoughts.

First, one could simply concede that forced labour could promote justice but maintain that it would be morally repugnant to employ it, due to intolerable costs to freedom, autonomy and rights against interference which are not equivalent to rights of self-ownership, costs which are not imposed by redistributive taxation.²⁵ In this respect, justice is no different to other values; since maximising utility or freedom at the expense

²³ Moreover, Otsuka points out that forced labour can be morally repugnant even if it does not violate a person's right of self-ownership. For example, the state could force someone into a particular occupation by threatening to withhold the resources necessary for their survival. He suggests that there is an important sense in which your life is not your own to lead when you are forced through the threat of starvation to spend all of your labour hours in a job which has been legally assigned to you.

²⁴ In G. A. Cohen, *Self-Ownership, Freedom, and Equality* (Cambridge: Cambridge University Press, 1995).

²⁵ Cohen has suggested that he would, under certain circumstances, consider the implementation of an egalitarian principle of distribution morally repugnant. See G. A. Cohen, "How to Do Political Philosophy" in *On the Currency of Egalitarian Justice and Other Essays in Political Philosophy*, p.234.

Similarly, in Nagel's, "Equality", p.108, he notes that "the promotion of equality may require objectionable means [...] greater equality may be attainable only by more general coercive techniques, including ultimately the assignment of work by public administration instead of private contracts".

of other values would also sometimes be morally repugnant. We should not “let justice be done though the heavens fall”.²⁶

Second, one could posit that only moral duties (though not all duties) are permissibly enforceable by the state and deny that justice requires people to undertake particular occupations entails that they are under a moral *duty* to undertake it. People are not generally morally required to impose large burdens on themselves for the sake of bestowing an accumulation of small benefits on many other worse off individuals even if doing so would make things go best from an impartial perspective, and it is generally not thought to be permissible to force people to do so. One might maintain, however, that there is a moral duty to redistribute income, when the costs of doing so to each person are not large and the benefits to others are great.

Third, one could say that the decisive reason to favour coercive redistributive taxation is not to remove intrinsic disvalue, but rather to remove the *instrumentally* bad effects of inequality, like inequality of power, social status or its detrimental effects on community. Such a position would be perfectly plausible, since few egalitarians would advocate, all things considered, a state policy of levelling down income or destroying wealth when doing so achieves no gains in terms of social equality. By contrast, one might maintain, forced labour in the name of equality would do little to remove the instrumentally bad effects of inequality.

3.4.2 The Second Formulation

A second way of formulating the forced labour objection is to say that Cohen lacks the normative resources to oppose the realisation of equality through forced labour *as a matter of justice*. This objection has been posed by Cécile Fabre, who argued that, by the lights of Cohen’s claims, a society in which distributive equality is legally enforced

²⁶ One might find this response unsatisfactory when it comes to egalitarianism as a theory of *justice*, rather than a view about the goodness of outcomes. We want to object to forced labour on the grounds of justice, rather than at the bar of an all things considered judgement about what is morally permissible. Some would regard it as a *reductio ad absurdum* of a theory of justice which allowed for its implementation to be morally repugnant.

through forced labour could be just.²⁷ On this formulation, the forced labour objection is a claim about what a just society *is*. But intuitively, for me at least, a society in which equality is enforced by forced labour would be an unhappy parody of a *just* society.

Fabre suggests that there are typically three kinds of reasons to provide the legal freedom to refrain from doing what justice requires: enforcing the requirements of justice might be morally impermissible, all things considered, coercively enforcing the requirements of justice might be practically infeasible, or trying to coercively reduce injustice might be counterproductive by unintentionally worsening it.

In terms of exegesis, it is unclear whether Cohen objected to the use of forced labour because he thought it was *unjust* or merely because he thought that forced labour was impractical or impermissible, all things considered. But I assume for the sake of argument that a plausible defence of Cohen will establish that a just society would not contain forced labour.

The remaining challenge is to provide an account of how and why coercively enforcing the requirements of justice through forced labour would be counterproductive, in contrast to redistributive taxation. Fabre considers and rejects four attempts to do so. I will briefly summarise the arguments that Fabre considers and rejects, before explaining what I take to be better replies which Fabre overlooks.

First, Fabre considers the idea that legally forcing people into productive jobs might deter them from acquiring the skills that they would need to perform those jobs in the first place. Coercive job allocation might therefore unintentionally but foreseeably worsen injustice rather than reduce it.

Fabre points out however, that in the absence of coercion, people might be deterred from acquiring productive skills in non-ideal societies in which equality is coercively enforced, due to the absence of incentives to do so. She also observes that even in a society in which people held the egalitarian conviction that they ought to acquire the skills that justice requires, they might not be sufficiently motivated to act upon those beliefs. She points out that a society of weak-willed egalitarians in which the

²⁷ Fabre's argument applies whether justice requires optimal equality or merely equality. Sometimes forced labour would serve to reduce inequality, rather than bring about an efficient distribution of labour. Fabre cites the example of forcing a well-off person to be a doctor in order to provide aid to those who are very badly off. Such conscription may serve to reduce inequality even if it does not also contribute to bringing about optimal equality.

requirements of justice are coercively enforced might be more just than a society of weak-willed egalitarians without any legal coercion.

Second, Fabre considers the claim that the state could not acquire the information needed to legally conscript people into their most productive occupation. This epistemic constraint means that forcing people into particular jobs runs a high risk of worsening inequality rather than reducing it by burdening people with work that leaves them worse off than others.

Fabre replies that we can make rough estimations of the burdensomeness of imposing a particular occupation on a person, much in the same way that we can roughly judge that, on average, higher income corresponds to higher welfare and justify egalitarian taxation on those grounds. Moreover, she points out that an epistemic aversion to imposing occupations is to the disadvantage of the worse off who stand to gain from the additional benefits provided by an increase in the social product and, on balance, a legally coercive policy is more likely to reduce inequality than legal freedom of occupational choice.

Third, Fabre considers the claim that it is better if people do the egalitarian thing due to egalitarian motivations, rather than due to state coercion. Fabre points out, however, that although it is usually better that people do the right things for the right reasons we do not normally regard that as a moral barrier to legally enforcing certain moral requirements. She also adds that being legally compelled to do the right thing can gradually induce people to do the right thing without coercion (pointing to examples including legally coercive health and safety requirements in the workplace and anti-discrimination legislation, which play some part in bringing about a change in attitudes and voluntary behaviour).

Fourth, Fabre looks at the claim that coerced labour impermissibly uses people as a mere means to the end of justice. In reply, she says that one cannot hold that legally compelling people into conforming to the requirements of justice uses them as a mere means: she says to do so would condemn the legal prohibition of murder and assault, let alone coercive taxation. She suggests that the Kantian objection to forced labour must appeal to the fact that it involves the state intrusively interfering in people's private lives, and rejects this on the similar grounds to Otsuka.

It is, of course, objectionable, all things considered, to force someone to live in a way which runs contrary to how *they* choose to live their life. But, I repeat, Fabre's formulation of the forced labour objection does not deny *that* claim. The objection says that, *by the lights of Cohen's account of distributive justice*, we cannot oppose coercive

job allocation *as a matter of justice*. In my view, Fabre's formulation of the forced labour objection can be overcome in three ways.

I think that one can coherently maintain that forced labour is unjust, even if it produces justice in distribution. There is a difference between egalitarian justice as an evaluative view which tells us which distributions are just and which are unjust and principles which tells us what we ought to do as a matter of justice; a difference between acting for reasons of justice and acting justly and unjustly. Justice can be both a deontological property of actions and an axiological property of distributions and it may sometimes be unjust to reduce injustice (for example, levelling down might be unjust even if it reduces distributive injustice).

Relatedly, Fabre's version of the forced labour objection can be resisted by distinguishing between a just *distribution* and a just *society*. One could agree that forced labour could reduce distributive injustice, but deny that it follows that a just *society* could contain forced labour. Cohen wrote:

A just society [...] is one whose citizens affirm and act upon the correct principles of justice but justice in distribution [...] consists in a certain egalitarian profile of rewards. It follows that, as a matter of logical possibility, a just distribution might obtain in a society that is not itself just.²⁸

In a society in which distributive justice is realised by coercive job allocation, it need not be true that citizens affirm and act upon the correct principles of justice. This does, of course, commit one to the view that a society in which equality is realised through coercive income redistribution is less than fully just. But this is not a big bullet to bite; indeed, Cohen says that in a perfectly just society, coercive taxation would not be required.²⁹

Finally, Fabre's formulation of the forced labour objection can be repelled by drawing on my proposal, in reply to the efficiency dilemma, that Cohen's claims be amended in order to judge a Pareto suboptimal state of equality unjust when a Pareto optimal state of equality is possible through people acting in ways in which they are capable of acting. I would reject the claim that there is no difference, as a matter of justice, between a society in which equality is coercively enforced by redistributive taxation and a society in which equality is coercively enforced by forced labour on the following grounds.

²⁸ Cohen, *If You're An Egalitarian, How Come You're So Rich?*, pp. 131-132.

²⁹ Cohen, *Rescuing Justice and Equality*, p.221.

In non-ideal situations in which people will not voluntarily do what justice requires, forced labour achieves equality by suppressing welfare. For any society in which equality is enforced by forced labour, there will typically be a feasible, equal, and Pareto superior state available which does not include forced labour. By contrast, this is not typically true of redistributive taxation.

I believe that this response captures the core of people's intuitive distaste for forced labour: we think that *part* of the injustice of coercive job allocation resides in the contingent fact that it would typically make everyone's life equally bad. So the response has intuitive support and a plausible rationale. But to develop this line of response in more detail, we must first distinguish between *varieties of non-ideality*. Of course, there are many degrees of non-ideality along the spectrum from a perfectly just world to the opposite extreme but we can delineate some general types.³⁰

First, there are worlds in which every person does what they are required to do, as a matter of justice, and morality more generally, without compulsion.³¹ For instance, imagine a perfectly egalitarian society in which everybody holds egalitarian beliefs and voluntarily does what justice requires of them. In such a world, a coercive state would not be required, but there may nevertheless be a state for practical reasons (for example, to solve co-ordination problems). Cohen says that in a truly just society, people's occupational preferences would *coincide* with what a Pareto optimal level of equality requires, rather than justice requiring people to act contrary to their prudential preferences due to their moral convictions.³²

Second, there are morally blemished worlds in which every person does what they are required to do as a matter of *justice*, but do not fully comply with the requirements of morality more generally. Again, in these kinds of worlds, there would be no need for coercive taxation or forced labour, even though such worlds would be morally imperfect. In these kinds of worlds, we could try to remove the moral imperfection of the world through state coercion, this but in many instances this would only compound

³⁰ This discussion draws on similar distinctions made in Hillel Steiner's "Levels of Non-Ideality", *The Journal of Political Philosophy* (2017).

³¹ We might even imagine a more perfect world than this: morally spotless worlds in which every individual does what is morally best, including carrying out morally supererogatory choices.

³² He says that "[O]ne focus of the egalitarian ethos is to make conscious focus on the worst-off unnecessary. What rather happens is that people [...] unreflectively live by, principles that restrain the pursuit of self-interest" in *Rescuing Justice and Equality*, p.73.

the world's non-ideality by doing something which does not fall within the moral jurisdiction of the state.

Third, there are morally tainted worlds in which some people *do not* voluntarily do what is required of them by justice (or morality in general) but all requirements of justice are coercively enforced by the state. In such worlds, justice may still not be fully realised because it would be impermissible for the state to coercively enforce justice across the board. For example, a society might have perfectly just anti-discrimination laws, but unjust discrimination nevertheless takes place in those spheres of everyday life in which it would be impermissible for the state to coercively interfere.

We can divide these third kinds of worlds into two types. In one kind of world, people *believe* that they ought to do what egalitarianism requires, but, due to the psychological burden of carrying out acts of self-sacrifice, are not moved to so act. For example, imagine a society of weak-willed egalitarians in which everybody wants to do the job that justice requires of them but cannot muster the will to do so.

In such a society, it may be objectionable to coercively give people the extra push required do the egalitarian thing that they justifiably believe that they ought to do. But it is *less* objectionable than coercing people who *lack* egalitarian beliefs. If a policy of state job allocation was agreed with democratic unanimity, for example, nobody would have grounds to complain once that had been legally assigned a job. So it cannot be *that* kind of non-ideality we have in mind when we recoil at the idea of forced labour.

The type of non-ideality in which we especially object to the use of forced labour is one in which people generally *don't* believe that they have an egalitarian obligation to do the job that optimal equality requires, and the state forces them do it anyway. Consider a world, not unlike our own, in which many people do not have egalitarian beliefs and would recoil at the idea of the state conscripting them into certain lines of work in the name of equality. Considering forced labour, Fabre says

[T]hat which A is [forced] to do neither makes her worse off than the less fortunate nor violates her personal prerogative – any more than giving a share of her material resources to the less fortunate would do.³³

But this is misleading. There is typically a difference in burdensomeness between someone choosing an occupation because of their egalitarian convictions and being forced to forego their preferred occupation in spite of their inegalitarian convictions.

³³ Fabre, "Distributive Justice and Freedom", p.399.

Moreover, there is a difference in burdensomeness between being coerced into expending one's labour in a way one disprefers and having legal conditions attached to the accumulation of income.

Many people in the actual world do not enjoy their jobs, but they are willing to do it for the sake of a higher income, even if, other things being equal, they would prefer to do a different job. It would be odd to say that, were we to force such individuals to do the same job for a lower income, this would not impose a large burden, merely because they are prepared to do the job for a price. One's life would go much worse on plausible accounts of well-being if a person were forced to perform a particular occupation that, in the absence of egalitarian convictions, they do not want to do.³⁴

Consider an example. Suppose Bartleby is a talented scrivener. Compare two situations. In the first situation, Bartleby is morally motivated to copy legal documents due to his justified conviction that copying legal documents fulfils his obligation to make a productive contribution to society. In the second situation, Bartleby lacks the belief that he has an egalitarian obligation to make a productive contribution to society. He would prefer not to copy legal documents but the state coerces him into doing so. In this situation, Bartleby is forced to spend his working days doing something he would prefer not to do. It seems that Bartleby's life is much worse in the second situation than the first.

While forced labour has deleterious effects on one's well-being that, under realistic conditions, reduce distributive injustice, the same cannot be said for the confiscation of one's income. It is implausible to say that taxing a person with inegalitarian beliefs would have an equally deleterious effect on their well-being as forcing them to take an occupation that they object to. Forcing someone to spend their life toiling in an occupation that they don't want to do diminishes their well-being. No comparable case can be made against redistributive taxation which confiscates income but leaves an individual free to otherwise lead their lives as they choose.

We are now in a position to draw on these claims and combine them with the claims I advanced in response to the efficiency dilemma in order to complete this line of response

³⁴ My reply to the forced labour objection resembles a remark made by Cohen in the course of discussing whether *the ethical solution* is oppressive when he says: "it blights a person's life to require her to do a job that she disprefers [...] where the appeal is to the severity of the deprivation that the dispreferred job imposes, then no problem for egalitarianism is raised, because egalitarians take into account, when expecting or not expecting a person to do a job, the full costs of doing so for that person", *Rescuing Justice and Equality*, p.213.

to the forced labour objection. To repeat, in response to the efficiency dilemma, I claimed that we could revise Cohen's view such that justice always prefers a high level of equality to a low level of equality. Above, I argued that forced labour under certain types of non-ideal situations achieves equality at the price of suppressing well-being. As a result, as a matter of contingent fact, forced labour would typically bring about an unjustly low level of equality in comparison to some available alternatives, including those in which freedom of occupational choice is legally guaranteed and equality is brought about by the coercive redistribution of income.

Accordingly, a society containing forced labour would typically not be just, when compared to the possible alternatives. By contrast, redistributive taxation typically would *not* bring about a low level of equality compared to the available alternatives. On the contrary, redistributive taxation would typically increase the sum total of welfare in society in addition to reducing inequality, due to the diminishing marginal utility of money. We can say that a society containing legally coercive job allocation will typically be unjust because Pareto superior states of equality will be contingently available. This feature distinguishes forced labour from the egalitarian redistribution of income.

One might object to this line of reasoning in the following way: although it may contingently be true that the realisation of equality through forced labour would lead to a Pareto suboptimal level of equality, it is also contingently the case that the realisation of equality through forced labour is better, from the perspective of distributive justice, than any possible alternative.

For in a society with legal freedom of occupational choice, many people are forced, not by legal conscription, but by economic circumstance, to do jobs that they would rather not do. If we are to condemn the realisation of equality through legally coercive job allocation on the grounds that it suppresses people's welfare by forcing them to do jobs that they don't want to do, must we also condemn the realisation of equality through the labour market in which some people's welfare is similarly suppressed?

I grant that there may be circumstances in which we must choose between legally coercive job allocation and equivalent unfreedom in the labour market, and, in such circumstances we could not oppose the use of forced labour as a matter of distributive justice. It is also true, for example, that in some circumstances of extreme inequality, forcing someone to provide assistance to the worse off would reduce inequality by improving the lot of the worse off and decreasing the gap between them.

Often we must choose between many non-ideal possible states of the world and reach a judgement about which is best, all things considered. Nevertheless, I think that my reply

to Fabre demonstrates that, in many contingent circumstances, forced labour would be inimical to the realisation of a Pareto optimal level of equality of access to advantage. Accordingly, as a matter of contingent fact, a society containing forced labour would typically not be just. As a result, my reply strips Fabre's objection of rhetorical force.

3.4.3 Summary

To summarise, one formulation of the forced labour objection says that, because Cohen does not believe in self-ownership, and the fact that forced labour involves violation of people's private thoughts is not a decisive reason to oppose it, Cohen lacks the normative resources to oppose the realisation of equality through forced labour, *all things considered*.

I argued that we can distinguish between the coercive enforcement of a just distribution of income and a just distribution of labour in four ways: (1) concede that forced labour could promote justice but oppose it, all things considered, due to the costs to freedom and autonomy, (2) deny that people are morally *required* to realise justice through their occupational choices, in contrast to choices concerning their income, and (3) defend redistributive taxation on grounds other than the fact that it promotes distributive equality.

The second formulation of the objection says that Cohen lacks the normative resources to oppose the realisation of equality through forced labour *as a matter of justice*. I argued that we can resist this by (4) maintaining that forced labour is unjust even though it may promote justice in distribution, (5) distinguishing between a just *distribution* and a just *society*, and (6) pointing to the differences between redistributing income and welfare across different varieties of non-ideal societies.

3.5 Conclusion

G. A. Cohen claimed that we could bring about an egalitarian society without sacrificing economic efficiency or legal freedom of occupational choice if people were moved by an egalitarian ethos to make good use of their productive talents. *The efficiency objection* says such a society would require a *productive ethos*, in addition to an egalitarian ethos, and such an ethos is both beyond what justice could require, and is inconsistent with Cohen's claim that justice is equality. The forced labour objection says that Cohen's claims commit him to the view that a just society could contain forced labour in the name of equality.

We are now in a position to combine my arguments to form two potential lines of response to these objections, both of which allow us to retain the egalitarian spirit of Cohen's claims. First, we could retain the claim that justice is equality, relinquish the claim that justice requires a productive ethos, and draw a distinction between a just distribution and a just society (as well as opposing the use of forced labour, all things considered). By doing this, we can hold on to the claims that equality, Pareto optimality, and freedom of occupational choice could be co-realised under realistic conditions and reject the claim that a just *society* could contain forced labour.

Alternatively, we can relinquish the claim that justice is (only) equality and instead claim that justice is a second order property encompassing other values, including Pareto optimality, while affording lexical priority to equality. That allows us to maintain that justice requires a productive ethos and reject the claim that a just distribution could be brought about by forced labour, due to the detrimental effects of forced labour on welfare in non-ideal societies.

3.6 Appendix The Equality Objection

The central purpose of this chapter was to provide a defence of Cohen's claim that we could bring about a Pareto optimal level of equality without using or endorsing legally coercive job allocation. I defended this claim from the efficiency dilemma and the forced labour objection. However, there is a further significant counterargument to Cohen's claim that warrants a response, which I will call *the equality objection*.

Egalitarians want to reduce distributive unfairness. But they also want to abolish stigmatising differences in social status and inegalitarian social hierarchies.³⁵ Some have argued that these two aspects of egalitarian concern can come into conflict. They point out that reducing *distributive* inequality can sometimes worsen *social* inequality, because doing so requires the humiliating and undignified unveiling of people's flaws and shortcomings. For example, Jonathan Wolff has pointed out that in order to realise distributive fairness in the actual world, we may need to intrusively gather information about the worse off in order to scrutinise whether their disadvantage is due to their

³⁵ See, for example, O'Neill, "What Should Egalitarians Believe?", and T. M. Scanlon, "The Diversity of Objections to Inequality" in *The Difficulty of Tolerance* (Cambridge: Cambridge University Press, 2003).

responsible choices or due to brute bad luck, dividing people into the “deserving poor” and “undeserving poor”.³⁶

Of course, an egalitarian is not committed to saying that we should do whatever it takes to bring about a fair distribution. So, *if* realising a fair distribution requires treating people in an undignified way, or carries other morally intolerable costs, then we should not realise a fair distribution either as a matter of equality, or all things considered. There are *egalitarian* constraints against the full realisation of distributive fairness.

Emily McTernan has argued that this observation raises a problem for Cohen’s solution to the trilemma of co-realising equality, Pareto optimality, and legal freedom of occupational choice. She argues that a society with a productive ethos would be likely to contain inegalitarian hierarchies which rank people in terms of their productive contribution to society, where one’s social status is tied to the extent of one’s contribution. Call this *the equality objection*.

Now, it should first be clarified that this objection does not apply to a society with a merely *egalitarian* ethos, since an egalitarian ethos does not induce the belief that each person ought to make a productive contribution to society, it merely induces the belief that everybody should be equal in terms of access to advantage (or so I have argued). The objection *does*, however, apply to a society with a *productive ethos*.

McTernan thinks that two kinds of inegalitarian hierarchies would be likely to emerge in a society with a productive ethos.³⁷ First, there would be an inegalitarian hierarchy between those who fully carry out their productive obligation and those who, for good reasons, do not. She cites the example of people who choose to refrain from taking paid employment in favour of providing full-time unpaid care for the young, the old, and the infirm. McTernan says that, despite being admirable, this is “often unlikely to be an individual’s most socially productive role”.³⁸ I think the same might be said of untalented artists, musicians, poets, gardeners, philosophers, and many others.

In a society with a productive ethos, one can imagine that people might raise the following complaint against those who choose not to take paid employment, or choose

³⁶ Jonathan Wolff, “Fairness, Respect, and the Egalitarian Ethos”, *Philosophy and Public Affairs*, 27.2 (1998), 97-122.

³⁷ Emily McTernan, “The Inegalitarian Ethos: Incentives, Respect and Self-Respect”, *Politics, Philosophy and Economics*, 12.1 (2013), 93-111.

³⁸ McTernan, “The Inegalitarian Ethos”, p.99.

an occupation which does not make good use of their talents: "I am contributing my fair share by making good use of my productive talents, even though, egalitarian obligation aside, I would prefer to do a different job. Why should others be allowed to shirk their obligation by refusing to make a fair productive contribution?"

To illustrate, it is a social norm in the United Kingdom at the time of writing that people who receive an income from society ought to take a job if one is available and they are capable of doing it. Those who receive an income without working, even when they are capable of doing so, are, rightly or wrongly, subject to social disapproval. We can imagine that something similar might happen if it was a social norm that everybody ought to make good use of their productive talents for the common good. McTernan points out that in an ideally egalitarian society, people should be free to refrain from carrying out their productive obligation in full (including the choice not to take paid employment) without suffering diminished social status.

Second, McTernan says that there would be an inegalitarian hierarchy between "the untalented", who make a comparatively small productive contribution to society, and "the talented" who contribute a comparatively large amount. This thought is motivated by the observation that, in a socialist market economy of the kind which may be required to realise Cohen's vision in practice, pre-tax incomes would reveal the size of each person's productive contribution to society. Knowledge that one's productive contribution to society is smaller than others might diminish people's sense of self-respect, and lead "the talented" to regard "the untalented" as their social inferiors. She concludes that

Cohen appeared to combine capitalism's efficiency with socialism's equality and motivation for the common good. But in incorporating capitalism's market structure, Cohen preserves its [unjust] hierarchical ordering [and] its inegalitarian consequences.³⁹

McTernan contemplates two ways in which someone might try to resist her arguments. First, she considers whether choosing not to make good use of one's productive talents falls within the sphere of moral permissions granted by each person's personal prerogative to refrain from doing what justice requires. She dismisses this on the grounds that the personal prerogative is not so extensive as to permit people to opt out of paid employment altogether.

³⁹ McTernan, "The Inegalitarian Ethos", p.107.

Second, she considers whether providing unpaid care for the young, old, and infirm should simply be counted as fulfilment of one's productive obligation. But she rejects this idea on the grounds that typically people who voluntarily opt out of paid employment would be capable of making a more productive contribution to society. Or, if they really are incapable contributing more, would feel ashamed of this fact, because it reveals a lack of "talent".

I think that this is an insightful objection. It may well be the case that, the facts being as they are, it would be impossible to realise an egalitarian society of the kind described by Cohen in the actual world without giving rise to inegalitarian social hierarchies. Cohen argued that what we ought to *believe* about the requirements of justice is independent of whether justice could, or should, be realised. Perhaps the ethos serves an illustrative function, by helping us to clarify what we ought to believe about what justice requires, but we should not try to bring about such a society, once all aspects of equality are taken into account, or all things considered.

However, I will now argue that Cohen's claims can be defended against McTernan's challenge. Since McTernan specifically mentions social inequality between "the talented" and "the untalented", and between the optimally productive and justifiably unproductive, I will restrict my reply to those two kinds of hierarchies. I will not consider other kinds of inegalitarian hierarchies that might arise. Before proceeding, I will say a little bit more about the constitutive parts of a productive ethos in order to illuminate the following discussion. The ethos can be broken down into four distinct parts.

(1) First, an ethos consists of widely shared *beliefs* and *attitudes*. A productive ethos would induce all or most people believe to that justice requires a Pareto optimal level of equality, and that each has an obligation to make choices, including occupational choices, which contribute to bringing that about. It would induce people to hold attitudes which contribute to the realisation of justice.

(2) Second, an ethos also requires *motivations*. An ethos would ensure that people are generally motivated to act upon their egalitarian beliefs. It is possible that people might have egalitarian beliefs but are not ultimately motivated to act upon them, (for example, due to weakness of the will).

(3) Third, an ethos consists of a set of *choices*. That is, the choices required of individuals to ensure that an optimal level of equality is realised. These could come apart from egalitarian beliefs and motivations, since one might make the

choices that egalitarianism requires, but be motivated by something other than egalitarian beliefs. Moreover, one might hold egalitarian beliefs and be motivated to act upon them, but fail to successfully make egalitarian choices, due to imperfect information or poor decision-making.

(4) Finally, an ethos requires that all or most people hold egalitarian beliefs, and are motivated to act on those beliefs, and act upon them successfully. If only a small number of people in a society hold egalitarian beliefs which they successfully act upon, equality would not be realised.

What McTernan seems to have in mind is situations in which most people hold egalitarian *beliefs*, or, in her words, have “internalised” the productive ethos, but not everybody makes the *choices* that are required to realise a Pareto optimal level of equality.

First, I will argue that we don’t have good reasons to think that social inequality would arise between the “the talented” and “the untalented” in a society with a productive ethos. To start, we should remember that “talent”, in this context, refers to the skills and abilities that are in lesser supply within a particular distribution of skills and consumer preferences; it does not imply that there is something intrinsically valuable about a particular kind of work. In a socialist market economy, high pre-tax wages would indicate which talents are less commonplace than others; it would not denote which talents or occupations are valuable. So why should we believe that people with less “talent” would feel ashamed of this fact?

The objection works if we assume that people hold particular beliefs about the relationship between “talent” and worth. But we should not assume that the population of a society with a productive ethos would hold beliefs which resemble the ones that many people have in the unjustly unequal societies that exist in the actual world. In anticipation of this point, Jonathan Wolff wrote

what is considered shameful is socially relative and contingent [...] What counts as a productive talent may vary from society to society, but what seems less variable is that those who are unable to make a significant contribution may feel at least somewhat ashamed of this fact [...] Now in an enlightened society of equals such attitudes might be considered an unfortunate fact about our barbaric prehistory. This, though, is psychological speculation and we have little, if any, good reason to believe that it is true.⁴⁰

⁴⁰ Wolff, “Fairness, Respect, and the Egalitarian Ethos”, pp.114-115.

Without wishing to deny that we can only speculate about the attitudes that the residents of an imaginary society would hold, it seems less likely that people in an egalitarian society *of the kind described by Cohen* would hold these attitudes. By stipulation, in a society with a productive ethos, most people would hold egalitarian beliefs and attitudes.⁴¹ A part of holding egalitarian beliefs and attitudes is to believe that natural talent endowments are *arbitrary* from a moral perspective. People guided by a productive ethos would believe that it is a matter of *brute luck* what one's natural talent endowments are, and how abundant or scarce those talents happen to be in the circumstances in which they find themselves.

It is difficult to see how, on the one hand, people could believe that we ought to eliminate morally arbitrary disadvantage and, on the other hand, believe that an unchosen lack of marketable skills as something to be ashamed of. If people *didn't* believe that natural endowments are arbitrary from a moral perspective, it is difficult to see why they would be motivated to make a productive contribution to society for an egalitarian wage without coercion. I would posit that believing that a lack of "talent" is something to be ashamed of seems inconsistent with the kinds of beliefs people would have *in a society with a productive ethos*.

To be sure, it's possible, indeed common, for people to regard some aspect of themselves as arbitrary from a moral point of view but also a shortcoming to be regarded with shame or embarrassment. For example, people often regard physical imperfections as something to be ashamed of while at the same time believing that those things are morally arbitrary. And McTernan rightly points out that people regarding their talents as inappropriate objects of both shame and pride would alienate people from important aspects of their personality.⁴²

But it is not *talents* per se that people with egalitarian beliefs would regard as inappropriate objects of shame or pride, but the *marketability* of those talents. One can be both proud of their talents and indifferent to their marketability. For example, a person might be proud of their philosophical talents and at the same time feel no shame about the fact that those talents are not marketable.

Moreover, in *any* society with a market economy which provides information about how much everybody contributes, there could be social division between those who

⁴¹ For a more comprehensive defence of this point, see Joanna Firth, "What's So Shameful About Shameful Revelations?", *Law, Ethics, and Philosophy*, 1 (2013), 31-51.

⁴² McTernan, "The Inegalitarian Ethos", p.206.

contribute more and those who contribute less. The best way to break the links between “talent”, productive contribution, and perceived worth, is to try to change people’s beliefs, in order to recognise that the *market value* of unchosen talents are arbitrary from a moral perspective. A society in which egalitarian beliefs are widespread would surely score better on *that* front than many others we can imagine.

I now turn to whether we should expect an inegalitarian hierarchy to emerge between those who carry out their productive obligation in full, and those who do not: either by choosing not to take paid employment or by choosing not to make the most of their talents. I will argue that we should not expect such a hierarchy to emerge.

Take the example of providing unpaid care that McTernan uses to illustrate her central claim. We need not assume that choosing to provide unpaid care for the young, old, and infirm, counts as a failure to fulfil one’s productive obligation. The fact that a person chooses not to contribute *to producing wealth* to the full extent that they are able to, does not, by itself, establish that the choice fails to fulfil their productive obligation. This would conflate productive efficiency with Pareto optimality.

The point of the productive ethos is *not* to ensure that society produces as much wealth as possible, but to ensure that we cannot make anyone better off, *in terms of access to advantage*, without making someone worse off. A society could have a Pareto optimal distribution of access to advantage but a very low productive output in terms of wealth. For example, if everyone in society preferred a lot of leisure time, then a Pareto optimal distribution of advantage could exist without much wealth being produced. In fact, there may be cases where we should level down the amount of wealth in society for the sake of realising a Pareto optimal distribution of advantage.

To illustrate, consider Cohen’s example of a person who can choose to between working as a doctor and working as a gardener.⁴³ Suppose we consider an analogous case of someone who can choose either to be a doctor or to refrain from paid employment in order to raise their children. One might want to conclude that since Cohen says that egalitarian justice requires the agent to choose doctoring over gardening he must also say that justice requires people to choose to be a doctor instead of opting out of paid employment to raise their children.

But in Cohen’s example, the agent is better off than others *regardless* of whether they choose to be a doctor or a gardener, but by choosing to be a doctor, they reduce the gap between themselves and others by benefitting people. Importantly, Cohen stipulates that

⁴³ Cohen, *Rescuing Justice and equality*, pp.184-185.

the agent would *not* be required to be a doctor if doing so would make them worse off than others. The case also excludes, by stipulation, other choices which may be available, including the choice to refrain from paid employment altogether.

For an analogous case in which the agent can choose whether to either doctor or provide unpaid care, we would need more information about the situation. Is the person who opts out of paid employment going to be better off than others regardless of what they do? Are the recipients of care going to be better or worse off if the person chooses to undertake paid employment instead? How much would it blight their lives to take paid employment contrary to their preferences? Without knowing the answers to these questions, we cannot say whether or not opting out of paid employment would count as a failure to fulfil one's productive obligation.⁴⁴

McTernan says that if people were motivated to provide unpaid care from a sense of social obligation, it would introduce "one thought too many", and would simply show that people are acting for the wrong kind of reasons. The same could be said of artists, musicians, doctors, nurses, teachers, and many others occupations.

But that's a question of motivation. True, it would be odd if people were motivated to give up a paid job in order to care for their children because they thought it was their social duty to do so. But, as I pointed out, people can make *choices* in accordance with what egalitarianism requires even if they are not motivated to do so by egalitarian beliefs. Moreover, people could provide care for loved ones for the right kinds of reasons, but still think that, by happy coincidence, that they are making a contribution to society.

Second, even if we assume for the sake of argument that occupational choices which appear to be suboptimal *are* in fact suboptimal, those who make such choices would be no better or worse off than anybody else, when all aspects of advantage are taken into account. Their overall share of the currency of justice would be the same as everybody

⁴⁴ Although I argued that justice requires Pareto optimal equality, a weaker reading of Cohen's claims says that justice permits, but does not require, a move from Pareto suboptimal equality to a state of Pareto superior equality, but forbids moves from equality to Pareto superior inequality. In Cohen's vision of a just society, the incentive to do what one believes to be one's egalitarian obligation replaces inequality-creating income incentives. So if an inequality-creating Pareto improvement is possible through at least one person working more productively for an inequality-creating incentive, then an equality preserving Pareto improvement is possible through at least one person working more productively without an inequality-creating incentive. This would help to bring about roughly the same degree of economic efficiency that would be brought about by the incentive of a higher income.

else's. It is therefore difficult to see why people who voluntarily choose to refrain from productively contributing to the full extent to which they are able, would regard this choice as something to be ashamed of. And it is difficult to see why others would disapprove of this choice, given that *they* have freely chosen to contribute to the full extent to which they are able, and are compensated accordingly with other aspects of advantage.

Third, it would not be transparent to each whether or not other people are making good use of their productive talents because each person's talent endowment is not transparent. People would simply lack knowledge of whether or not each person is carrying out their productive obligation, even if they *would* regard such choices as something to be ashamed of (which, I have claimed, they would not).

In summary, I do not think we have good reasons to worry that a society with a productive ethos would contain social division between the more productive and the less productive. To be sure, in possible non-ideal societies between the actual world and Cohen's imagined one, some of these problems may arise, and I cannot *show* that *the equality objection* is mistaken. But insofar as we can speculate about what would happen in societies unlike any that currently exist, I do not think that we have good reasons to worry that the realisation of optimal equality through a productive ethos would be inimical to social equality.

Finally, we should bear in mind the following. If it *were* true that a society with a productive ethos would be unjustly hierarchical, in which worth was tied to productive contribution, this would show that, under realistic conditions, we could not co-realise social equality, distributive fairness, Pareto optimality, and legal freedom of occupational choice. It would not, however, show that we should not aim to bring about such a society. For other kinds of societies may be even worse, both regarding inequality and taking everything into account.

4 KILLING AND DUCKING: SELF-DEFENCE AND MORAL EQUIVALENCE

4.1 Introduction

I now turn away from the redistribution of benefits and burdens in the domain of distributive justice towards a related problem concerning the imposition and redistribution of *harm* in the domain of normative ethics. More exactly, I turn to a problem in the morality of imposing and redistributing harm in self-defence.

Many people think that two acts of self-defence can differ in permissibility even if they bring about the same distribution of harm, because the mode of agency used to bring about an outcome can make a moral difference. Consider the following cases.

Killing

You are trapped at the bottom of a deep well. An innocent person is, through no fault or choice of their own, falling down the well. If they land on you, they will survive, but you will be killed. You can survive only by vaporising the falling person with a ray gun.¹

Ducking

You are trapped at the bottom of a deep well. An innocent person is, through no fault or choice of their own, falling down the well. If they land on you, they will survive, but you will be killed. You can survive only by stepping aside to allow the person to fall to their death.²

Some people judge it morally permissible to defend your life in both cases. But according to *the moral immunity thesis*, it would be *wrong* to kill the falling person in the first case, on the grounds that doing so is morally indistinguishable from killing an innocent

¹ Nozick, *Anarchy, State, and Utopia*, p. 34.

² The morality of “ducking” harm is explored in Christopher Boorse and Roy A. Sorensen, “Ducking Harm”, *The Journal of Philosophy*, 85.3 (1988), 115-34.

bystander in the course of defending one's life.³ Instead, you must allow yourself to be killed by the falling threat.

By contrast, some proponents of the moral immunity thesis plausibly maintain that it is permissible to evade the threat to your life in the second case, even though you foresee that an innocent person will be killed as a result. On the face of it, these two positions appear easy to reconcile; after all, there are clear moral differences between the two cases. In the first case, you kill an innocent person by creating a new and distinct harmful sequence which violently imposes upon the bodily sovereignty of the victim. In the second case, you allow an innocent person to die as a side-effect of redistributing harm, without even touching the victim.

However, I believe that, contrary to appearances, the two judgements cannot be easily reconciled by the lights of the moral immunity thesis. I defend this scepticism by arguing for the following central claim.

The moral immunity thesis maintains that killing an innocent *threat* in self-defence is morally indistinguishable from killing an innocent *bystander* in self-defence. If this is so, then both killing an innocent threat and killing an innocent bystander in self-defence are, in terms of permissibility, indistinguishable from *ducking* a threat.

I shall defend this central claim on the grounds that, by the lights of the moral immunity thesis, there is no moral feature which is both (a) essentially present in ducking but typically absent from defensive killing, and (b) makes a difference to the permissibility of acting in self-defence.⁴

If this central claim is true, then it shows that the moral immunity thesis reaches a highly counterintuitive conclusion. This would give us a good reason to suspect either that the moral immunity thesis is mistaken in some way, or, less plausibly, that the judgement that ducking is permissible should be revised.⁵

³ Michael Otsuka, "Killing the Innocent in Self-Defence", *Philosophy and Public Affairs*, 23.1 (2004), 74-94, and Jeff McMahan, "Self-Defence and the Problem of the Innocent Attacker", *Ethics*, 104.2 (1994), 252-90.

⁴ The significance of the distinction between moral properties which are essentially present, and only contingently present, was identified in Frances Kamm, "Harming, Not Aiding, and Positive Rights", *Philosophy and Public Affairs*, 15 (1986), 3-32.

⁵ John Martin Fischer and Mark Ravizza defend the view that ducking harm is morally equivalent to killing in "Ducking Harm and Sacrificing Others", *Journal of Social Philosophy*, 25.3 (1994), 135-45.

The chapter is divided into three main sections. In what remains of the first section, I explain the distinctions between innocent threats, innocent bystanders, killing and ducking. In the second section, I set out the moral immunity thesis in more detail and provide preliminary support for my claim that, by the lights of that account, ducking is morally equivalent to killing threats and bystanders, in terms of permissibility. In the third section, I consider five ways in which the moral immunity thesis might be defended from my claim, and argue that they do not succeed.

All references to victims, threats and bystanders will refer to innocent people and all references to killing and ducking will refer to acts of self-defence. References to “moral equivalence” refer to moral equivalence in terms of permissibility. One act may be easier or harder to justify than another even though both acts have the same moral status – are morally equivalent - in terms of permissibility. I focus on cases in which the indivisible harm at stake is the loss of life to only one person, but it may be possible to generalise the arguments to cases of non-lethal harm in which the expected harm or burden suffered by any one individual is of an equivalent magnitude.

4.2 Innocent Threats, Innocent Bystanders, Killing & Ducking

An innocent threat is someone who threatens lethal harm without justification but is not morally responsible for doing so, because they did not intend to pose a threat and exercise no agency; they pose a threat simply through the movement of their own body, which they cannot control.⁶ An innocent bystander is someone who plays no causal role in the creation or sustainment of a threatening sequence of events and has not done anything else to make them liable to defensive harm.

One kills an innocent threat or bystander when one creates, sustains, or inserts somebody into, a sequence of events which results in the death of a person.⁷ One ducks a lethal threat when, without killing as so-defined, one evades a threat to one’s life with the foreseeable result that an innocent person is killed instead. This includes moving

This was the subject of a reply by Christopher Boorse in “Ducking Trolleys”, *The Journal of Social Philosophy*, 25.3 (1994), 146-52.

⁶ We can distinguish innocent threats from innocent *aggressors* or *attackers*, who pose a threat through their intentional agency, but who are not morally responsible for their agency at the time. I concentrate on innocent threats and do not make any attempts to morally differentiate threats from aggressors.

⁷ I follow Otsuka in employing this definition of killing. “Killing the Innocent in Self-Defence”, p.76, footnote 7.

oneself out of the path of a moving threat (such as a projectile) but also removing oneself from a threatening situation with the foreseeable effect of exposing another person to the impending hazard.

The killing and ducking distinction is a factual one which correlates with several morally significant distinctions, some of which are essential to the factual classification of each kind of self-defence and others which are contingently, but unusually, present. For example, it is true by definition that ducking a threat does not involve creating a new threat.

But one might feel differently about the permissibility of ducking a threat when certain contingent features are present. Consider evasive actions which violate a widely accepted norm or make use of an object or resource other than one's own body. Imagine a sinking ship without enough lifeboats to save all of the passengers. One might find it objectionable if someone were to jump an orderly queue and grab a lifeboat for themselves, thereby condemning somebody else to drown. This would violate a norm of holding a fair procedure to distribute the scarce resources required for survival.

Ducking may also be objectionable when the resultant harm is intended or when the ducker was morally responsible for another action which created the threat. For example, suppose you intend the death of a rival and you place yourself into the path of an oncoming projectile, obscuring your rival's view of the danger to come. At the moment before the projectile collides with you, you step aside and ensure that it fatally collides with your rival. These might be unusual examples of impermissible ducking. Nevertheless, "pure" ducking cases, which lack these atypical features, are intuitively permissible.

4.3 The Moral Immunity Thesis

Three desiderata of an account of self-defence are that it (1) tells us when it is permissible and when it is wrong to kill in self-defence, (2) provides a rationale which explains why a particular act of self-defence is permissible or wrong, and (3) provides results which coincide with our intuitive judgements in a range of cases.

According to the moral immunity thesis⁸:

⁸ I use the phrase "moral immunity thesis" to refer specifically to the claim that killing an innocent threat is wrong because of its similarity to killing an innocent bystander. This claim is typically associated with proponents of "the moral responsibility account of liability to defensive killing", which

- (1) It is wrong to kill innocent bystanders in the course of defending one's own life (*the inviolability of a bystander thesis*).
- (2) Other things being equal, killing a person who innocently threatens one's life is morally equivalent to killing an innocent bystander in the course of defending one's own life, in terms of permissibility (*the moral equivalence thesis*).
- (3) It is wrong to kill innocent threats in the course of defending one's own life.

The rationale for these claims is that it is permissible to kill someone in self-defence without their consent only if they are morally responsible for posing a threat without justification; otherwise, it is wrong. This is because people who are not morally responsible for posing a threat retain a right not to be killed and acting in defence of one's own life is not a sufficient justification for overriding that right.

The striking feature of the account is that it says that one is morally required to allow oneself to be killed when the only alternative is to kill an innocent person in the course of preserving one's life. Yet some proponents of the moral immunity thesis maintain that ducking an imminent threat is morally distinguishable from killing. For example, Michael Otsuka says: "It is, I think, permissible to duck out of the way of an oncoming javelin even if you foresee that, by ducking, you will allow the javelin to impale a Bystander standing behind you [and] to move out of the way of a falling Threat even if you foresee that doing so will result in that Threats death".⁹ And Jeff McMahan says that ducking a threat "is self-preservation by a form of redirection. It is generally permissible – primarily, I believe, because it normally involves allowing unintended harm to occur."¹⁰

Notice that the moral equivalence thesis need not deny that there are *some* moral differences between killing threats and bystanders. It simply states that if there are moral differences, they are not sufficient to render self-defence permissible in one case but not the other. Similarly, I do not deny that killing in self-defence is harder to justify than

says that moral responsibility for posing a threat to someone is a necessary condition for liability to defensive harm. But that account does not, by itself, tell us when there is sufficient justification to kill or bring about harm to someone in any particular case. That is why I refer to *the moral immunity thesis* instead. Thanks to Susanne Burri for emphasising this point.

⁹ Otsuka, "Killing the Innocent in Self-Defence", p.76, footnote 7.

¹⁰ McMahan, "Self-Defence and the Problem of the Innocent Attacker", p.253.

ducking. What I deny is that those differences translate to a difference in permissibility in self-defence cases, by the lights of the moral immunity thesis in particular.¹¹

My defence of this claim is largely a negative one: I will argue against five ways in which one might try to morally differentiate killing from ducking, by the lights of the moral immunity thesis. However, before proceeding with my negative argument, I will briefly offer a positive one which provides some preliminary support for my claim. Consider the following case which acts as an intermediary between the two cases which I presented at the beginning of the chapter.

Barrier

You are trapped at the bottom of a deep well. An innocent person is, through no fault or choice of their own, falling down the well and will crush and kill you if they land on you. They will survive the fall if they land on you. You can survive only by interposing a barrier between yourself and the falling threat. The barrier will shield you from harm, but the falling threat will be killed upon colliding with it.¹²

According to the moral immunity thesis, it would be wrong to interpose the defensive shield, since doing so would kill the person who innocently threatens one's life. This case is analogous to one presented by Otsuka in which he says that it would be wrong to continue to hold a flagpole which will impale a falling person: "I do not believe that continuing to hold the flagpole is, except perhaps trivially, morally less bad than shooting one's ray gun".¹³

I do not think it is plausible to maintain that interposing the barrier is wrong but stepping aside to allow the person to fall to their death is permissible. There may be superficial differences between the two actions, but this does not seem to me to amount to a difference in moral status. One can therefore proceed, by a transitive relation of

¹¹ I do not, for instance, argue that the factors discussed lack moral significance or that because killing and ducking are equivalent in permissibility when it comes to self-defence, they are equivalent in other situations. Nor do I argue that my claim extends to other accounts of the morality of self-defence.

¹² This is inspired by an example in Jonathan Quong, "Agent Relative Prerogatives to Do Harm", *Criminal Law and Philosophy*, 10 (2016), 815-829, (p.819).

¹³ Otsuka, "Killing the Innocent in Self-Defence", p.89

moral equivalence, from the claim that killing the threat with a ray gun is wrong to the claim that ducking the falling threat is wrong.¹⁴

Now, to resist this claim, a proponent of the moral immunity thesis must establish that there is some moral factor typically present in ducking which differentiates it from killing, or vice versa, without also undermining the moral equivalence thesis or the inviolability of a bystander thesis. In the rest of the chapter, I consider five moral features which a proponent of the moral immunity thesis can use to try to draw a moral disparity between killing and ducking and argue that none of the attempts succeed in a way which coheres with the account.

This is an argument by elimination and there could, of course, be additional features which I do not consider; my argument does not show that the moral immunity thesis is mistaken. However, it does (if persuasive) undermine its appeal. The moral immunity thesis draws appeal from its coherence with other moral convictions, like the belief that it is morally impermissible to kill an innocent bystander when acting in self-defence. If, as I contend, the account counterintuitively implies that ducking a lethal threat is equivalent to killing the innocent in self-defence, then we have a good reason to be sceptical of the account, particularly if other plausible accounts of the permissibility of self-defence do not have similarly counterintuitive implications.

4.3.1 The Right of Self-Ownership

An immediately appealing way of distinguishing between killing and ducking is by pointing to a right of self-ownership over one's own body, and the correlative moral permissions and immunities which flow from this right. This right is significant in two respects when it comes to self-defence.

First, a right of self-ownership grants moral permission to do what one wishes with one's own body, providing that one does not violate the rights of others. So, one might say: "By ducking a threat, I act permissibly, because my body belongs to me and I am entitled to do as I please with it, as long as I do not use it to violate the rights of others. I do not enjoy the same rights over other people's bodies or objects to which others have a rightful claim, or which have no prior claims attached to them".

¹⁴ A very similar argument is briefly made by Helen Frowe, in *Defensive Killing* (Oxford: Oxford University Press, 2014), pp. 64-66.

Second, a right of self-ownership provides moral immunity from having one's body imposed upon without one's consent. One might maintain that the right not to be imposed upon is what could morally prohibit killing the innocent in self-defence, even when such harm is a side-effect of one's action. One could say: "It is unjustifiable to impose upon another person in self-defence, but it is justifiable to shift a threat to another innocent person in self-defence if one can do so without imposing upon them. Killing involves wrongful imposition, whereas the act of ducking a threat does not involve such imposition."¹⁵

To test the first claim, let's consider an act of self-defence in which one ducks a threat by using something other than (merely) one's own body.

Car

You are trapped inside an abandoned car on trolley tracks through no fault or choice of your own. The car does not belong to anybody. A trolley hurtles towards the car and will kill you if it collides with the car. An innocent person is tied further down the trolley tracks, and is afforded protection by the car. You drive the car out of the path of the trolley, which then foreseeably kills the innocent person.¹⁶

This case seems morally indistinguishable from ducking a projectile using only one's own body. Yet in this case, rather than merely their own body, the agent makes use of an object which belongs to nobody. This implies that it is not the right of self-ownership which exclusively grants moral permission to evade a threat to one's life. Rather, one is sometimes permitted to evade a threat even when doing so requires the use of an object to which nobody has an exclusive claim.

Now let's consider the second aspect of the self-ownership thesis. If a right against bodily imposition helped to differentiate between permissible and wrongful ways of killing others in self-defence, then it should be harder to justify defensively killing someone through directly imposing on their bodily sovereignty than it is to kill them

¹⁵ Indeed, Otsuka says that "it is a right of self-ownership not to be imposed upon [which makes killing threats and bystanders wrong] rather than a right not to have made use of something to which one has a rightful claim – whether it be one's body or the space one occupies" in "The Moral Responsibility Account of Liability to Defensive Killing", in *The Ethics of Self-Defence* ed. By Christopher Coons and Michael Weber (Oxford: Oxford University Press, 2016), p.14.

¹⁶ This case is borrowed from Kai Draper, "Rights and the Doctrine of Doing and Allowing", *Philosophy and Public Affairs*, 33.3 (2005), 253-80 (p.262).

without such imposition and, moreover, this difference in justifiability must amount to a difference in permissibility, other things being equal. We can test this claim by considering a case of killing a bystander as a result of acting in self-preservation, when there is no clear bodily imposition.

Nitrous Oxide

You are trapped in an airtight room. Nitrous oxide is being released into the room which will soon suffocate you because of the resultant displacement of oxygen. An innocent bystander, whose life is not presently endangered, is trapped in a nearby room. You can turn a dial which will disable the release of nitrous oxide into the room. However, the dial will also foreseeably release nitrous oxide into the nearby room which, through the displacement of oxygen, will kill the innocent person trapped inside.

Consider some of the morally significant aspects of this case. First, the agent kills the victim without imposing on their bodily sovereignty. Second, the harm inflicted upon the victim is a side-effect of action aimed at self-preservation, and the side-effect itself does not contribute to the survival of the agent. Third, the harm is inflicted by unblocking a threat (aspects which I examine in more detail below). My own intuitive judgement is that this is no different in permissibility from killing a bystander as a side-effect through direct imposition upon their body: if one is permissible then so is the other and if one is impermissible then so is the other.

One might respond to this case in one of three ways. First, one might say that the *Nitrous Oxide* case is an example of the violation of a person's right of self-ownership. This suggests that it is not imposing upon someone's bodily sovereignty without their consent which is morally amiss but also causing physical harm to take place, without laying a hand on the victim. Yet this would render the harm brought about by ducking a violation of a right of self-ownership since on a plausible difference-making counterfactual account of causation, ducking is a cause of the harm to the victim: if you duck, the victim will be killed. If you do not, the victim will survive.

Second, one might say that a right of self-ownership is *not* violated in this case on the grounds that the agent does not touch the victim. This would concede that the presence or absence of bodily imposition is not crucial to determining the permissibility of acting in self-defence. Note that this claim would be curious since initiating threats which go on to inflict harm on a person - for example, pulling the trigger of a gun which is aimed at someone - would typically count as violations of their right of self-ownership, even if the agent never comes into contact with the victim. What matters is that actions which

flow from one's intentional agency play a relevant causal role in the harm that ultimately befalls the victim.

Third, one might grant that the act of self-defence does not violate rights of self-ownership on the grounds that the harm imposed upon the victim is a side-effect of one's act of self-defence and the side-effect itself does not contribute to the survival of the agent. One might maintain that one violates another person's right of self-ownership only when using them or harming them as a mere means.¹⁷ This claim requires a more detailed examination, which I will now provide.

4.3.2 The Side-Effect Principle

A second way of drawing a clear and plausible moral distinction between killing and ducking is by pointing to the fact that harm brought about by ducking a threat is a side-effect of action aimed at self-preservation. Ducking a threat involves, in Warren Quinn's terminology, indirect harm.¹⁸ By contrast, when killing a threat or bystander, the agent directly harms a person in order to secure survival: whether by opportunistically *using* the victim as a mere means (what Quinn calls *direct manipulative agency*) or by *eliminating* a threat or obstacle that the victim presents (what Quinn calls *direct eliminative agency*).¹⁹

To elaborate, some instances of killing in self-defence opportunistically exploit the presence of an innocent bystander as a means of securing survival: consider grabbing a bystander to use them as a human shield against a threat.²⁰ Other instances of killing do not opportunistically use a bystander as a means, but do eliminate a bystander whose presence is an obstacle to survival; consider a case in which a person is driving at speed

¹⁷ Otsuka advocates this formulation of the right of self-ownership in *Libertarianism Without Inequality*, p.15.

¹⁸ Warren Quinn, "Actions, Intentions, and Consequences: The Doctrine of Double Effect", *Philosophy and Public Affairs*, 18 (1989), 334-51. In light of Quinn's convincing revisions to the doctrine of double effect, I set aside the distinction between intending harm and merely foreseeing harm.

¹⁹ Judith Jarvis Thomson, "Self-Defense" in *Philosophy and Public Affairs*, 20.4 (1991), 283-310 (pp.289-290). It is true that in such cases one inflicts harm on an innocent threat as a means of securing survival but you do not use *them*, or exploit their *presence*, for your own advantage. Helen Frowe draws the distinction between harming as a means and using as a means, in "Equating Innocent Threats and Bystanders", *Journal of Applied Philosophy*, 25.4 (2008), 277-90.

²⁰ Thomson, "Self-Defense", pp.289-290.

to elude a threat and runs over an innocent bystander on a narrow bridge who obstructs their path to safety.

A further class of cases, which are important for present purposes, are those in which harm is inflicted on a bystander as a side-effect, in which the death of the victim does not causally contribute to the survival of the agent. One might plausibly contend that using a person as a mere means is harder to justify than eliminating them, and, moreover, that eliminating someone is harder to justify than killing them as a side-effect. One might further hold that the difference in justifiability between direct harm (in either sense) and harming as a side-effect translates to a difference in permissibility when it comes to self-defence.

Accordingly, one might say the following: “it is unjustifiable to directly harm a person by opportunistically exploiting their presence, or by eliminating the threat or obstacle that they innocently present. It is, however, justifiable to kill someone as a mere side-effect of action aimed at self-preservation.”²¹

This picture is complicated by the existence of ducking cases in which the harm that befalls the victim does not fall easily into any of the categories described above. There are cases of apparent ducking which may also be classified as opportunistically exploiting a bystander. For example, suppose an agent is threatened by a non-human projectile and the only way to survive is to step behind a bystander, without touching them, so as to benefit from their presence as a human shield.²² The moral status of this action is intuitively unclear.

Even if we set aside this complication, I do not think that the moral immunity thesis can appeal to these claims in order to differentiate ducking from other kinds of self-defence. The reason for this is that the most powerful argument in favour of the moral

²¹ Some reject the moral immunity thesis by claiming that, contrary to what the account says, it is permissible to employ eliminative agency in defence of one’s life, but impermissible to employ opportunistic agency in defence of one’s life. This difference in permissibility can be used to draw a moral disparity between opportunistically using innocent bystanders as human shields and eliminating the danger posed by an innocent threat. See Jonathan Quong, “Killing in Self-Defence”, *Ethics*, 119.3 (2009), 507-37, for this argument.

²² Similarly, Boorse and Sorensen open their article on ducking with the example of a grizzly bear in pursuit of two people. The first person outruns the second, who is subsequently caught and eaten by the bear. This provides a welcome distraction which allows the first person to escape unharmed. In this case, I would be inclined to describe the harm suffered by the victim as a side-effect, even though the victim’s death contributes to the survival of the agent.

equivalence thesis rests on the claim that the indirect killing of a bystander as a side-effect of one's act of self-defence is morally equivalent to eliminating an innocent threat. Consider the following example of such an indirect killing.

Dynamite

You are trapped on trolley tracks and a runaway trolley is hurtling towards you. It will kill you if it collides with you. The only way to save yourself from the threat of the trolley is to hurl a stick of dynamite in order to destroy it. Regrettably, the explosion will foreseeably kill a person who stands nearby.

The moral immunity thesis seems committed to condemning this indirect killing of a bystander. Indeed, Otsuka says that "even in cases in which the Bystander's body is of no use to you, but you know that you will survive only if you initiate a sequence of events that you know will kill her, it is impermissible to do so". So it seems that the moral immunity thesis does not discriminate between indirectly harming as a side-effect and direct eliminative harm.

It is true that one could infer the moral impermissibility of killing an innocent threat exclusively from its similarity to eliminating a bystander who obstructs one's path to safety. However, if it were permissible to indirectly kill an innocent person as a side-effect of defending one's own life, then the moral immunity thesis would have to be qualified to permit those instances of killing a bystander.²³ But this revision would clearly be unsupported by the underlying rationale for the account. As innocent bystanders who pose no threat or disadvantage to anybody, their moral status is not reduced and, accordingly, they retain a right not to be killed.

One might, however, employ more finely tuned distinctions between varieties of harming as a side-effect. For instance, one might distinguish hurling the stick of dynamite from ducking on the grounds that the harm from hurling the bomb is "causally sidestream from" the outcome of self-preservation while the harm from ducking is easier to justify because it is causally downstream from the outcome of self-preservation.²⁴

Similarly, we could say that, in the dynamite case, the death of the bystander is an event caused by a new sequence of events that one has initiated, whereas in a ducking case, the

²³ Otsuka suggests that he accepts this revision in "The Moral Responsibility Account of Liability to Defensive Killing", pp. 55 to 57.

²⁴ McMahan, *The Ethics of Killing*, p.408.

death of the bystander is the flip-side of shifting a pre-existing threat.²⁵ I now turn to consider the moral significance of these distinctions between causal sequences in more detail.

4.3.3 The Redistribution Principle

One might try to vindicate a difference in permissibility between killing and ducking by pointing to moral differences in how harmful sequences causally unfold. For example, one might say the following. “The presumption against *creating, sustaining, or inserting somebody into* a threatening sequence is much stronger than the presumption against *merely redistributing* a pre-existing threat from one person to another. Other things being equal, this difference in justifiability amounts to a difference in permissibility when it comes to self-defence. Killing involves creating, sustaining, or inserting somebody into a lethal sequence but ducking merely redistributes a pre-existing threat from one person to another”.

To be clearer about what these terms mean, you *create* a lethal sequence of events if your act sets in motion a sequence of events which culminates in lethal harm. You *sustain* a lethal sequence of events if your actions or inactions causally contribute to the continuation of a lethal sequence that would otherwise come to a halt. You *insert somebody into* a threatening sequence if you move somebody into the path of an existing threat. You *redistribute* a threat when you switch the victim of a pre-existing threat without interfering with the threat itself. I assume for the sake of argument that these distinctions are metaphysically robust.

This explanation coheres well with some other widely held moral beliefs. For example, many people believe that it is permissible to *redirect* a threat away from five people towards one person. By contrast, many believe that it is wrong to *insert* one person into the path of a threat (for example, by moving them into the path of a runaway trolley) in order to save five others from being killed.²⁶ But note that redistributing a threat differs from *redirecting* a threat. In both classes of cases, the victim of a pre-existing threat is

²⁵ Frances Kamm, *Morality, Mortality Volume II: Rights, Duties and Status* (Oxford: Oxford University Press, 1996), p.174.

²⁶ There is disagreement about whether redirecting a trolley away from oneself towards one other innocent person is permissible. For instance, Thomson thinks that substituting a bystander by redirecting a threat is wrong (“Self-Defence”, pp.289-290), but Quong thinks it is permissible (“Killing in Self-Defence”, p.512).

changed, but redirection moves a threat into the path of the victim and redistribution substitutes one victim for another, without interfering with the threat.

The first problem with trying to differentiate killing from ducking using these distinctions is that the factual categories are not exclusive. *Inserting somebody into* a causal sequence is, like ducking, a variety of redistributing a pre-existing threat. So one cannot differentiate killing from ducking merely on the grounds that the latter redistributes harm.

We can test the claim that redistributing lethal harm in self-defence is generally permissible by considering a case of killing a bystander through the redistribution of a threat. Consider the following case.

Rope

You are standing on trolley tracks and unfortunately a runaway trolley is speeding towards you. Fortunately, you are able to step out of the path of the trolley. However, you are tied to a nearby innocent bystander by a piece of rope. By stepping out of the path of the trolley, you will drag the bystander into its path (the innocent bystander does not shield you: their death does not contribute to your survival).

My own judgement is that, although easier to justify, this act of redistribution does not differ in permissibility from grabbing somebody to use them as a human shield against a projectile.

In any case, by itself, the distinction between creating and redistributing a threat is factual, rather than normative; an underlying rationale is required in order to morally distinguish between them. I submit that there are two plausible moral rationales for distinguishing between creating and redistributing a threat in self-defence, both of which create problems for the moral immunity thesis.

First, the moral difference might be explained by pointing to the fact that redistributing a threat brings about harm as a side-effect. But, as I argued above, the moral immunity thesis seems committed to the claim that the difference in justifiability between direct killing and side-effect killings does not translate to a difference in permissibility. Here, the search for a rationale to explain why the presumption against redistributing a pre-existing threat and harming as a side-effect seems weaker than other ways of causing harm becomes circular.

Second, and more plausibly, the difference might be explained by maintaining that we are especially responsible for what we *do*, rather than what we allow to happen. Creating, sustaining and insertion all typically involve *doing* harm whereas redistribution of a pre-existing threat by ducking merely *allows* harm (this does not explain a moral difference between the *redirection* and creation of a threat, because both of those might plausibly be regarded as varieties of doing harm). Since this explanation requires a more thorough examination, I move on to consider it separately.

4.3.4 Killing & Letting Die

One might suggest that the creation and redistribution distinction only approximates the moral boundary between permissibility and impermissibility in self-defence cases. Perhaps the creation and redistribution distinction is significant only insofar as it sometimes correlates with the distinction between killing and letting die. Accordingly, one might say the following. “Killing an innocent person in self-defence is harder to justify than letting someone die in self-defence. Ducking a lethal threat does not *kill* anybody; it merely allows harm to befall an innocent person. At worst, ducking is the justified removal of protection.”²⁷

Now, it would be a mistake to assume that ducking is a case of letting die *because* it seems easier to justify than killing; our factual classification of cases into either killing or letting die is easily influenced by our moral appraisal of a situation. First, we need a factual account of which acts count as killing, which acts count as letting die, and which acts fall into neither category. Second, we need a moral explanation of when and why killing is harder to justify than letting die when it comes to self-defence.

Let’s start with the factual account. There are at least two broad types of letting die. First, one can forbear to prevent a lethal sequence of events from occurring. Second, one can remove a protective barrier, thereby unblocking or releasing a lethal threat. The latter is typically called *enabling* harm.²⁸ Enabling harm is conceptually intermediate

²⁷ I believe it is easier to justify *withdrawing aid* that one has deliberately intervened to provide, than to remove protection that one is unintentionally providing. The former involves restoring a situation to the status quo that obtained before one’s intervention (in terms of the distribution of prospective harm). It undoes one’s own well-intentioned actions. The latter involves disrupting the status quo in one’s own favour, shifting a prospective harm from oneself to another person.

²⁸ These distinctions were drawn by Philippa Foot in “The Problem of Abortion and the Doctrine of Double Effect” reprinted in *Virtues and Vices and Other Essays in Moral Philosophy* (Oxford: Clarendon Press, 2002), pp. 26-27.

between killing and letting die through forbearance to prevent, but there is disagreement over whether enabling death is morally equivalent to killing, morally equivalent to letting die through forbearance, or has a moral status which lies somewhere between the two.²⁹

Ducking is an action, rather than an inaction, which exposes someone to a threat who was otherwise protected from it. Rather than being a forbearance to prevent a harmful sequence of events, ducking is better described as the enabling of harm through the removal of protection from a pre-existing threat.

Ducking does, however, differ from other paradigmatic cases of enabling harm because it involves situations in which there is an active threat and an agent can shift it from themselves onto another person. This differs from unblocking a threat which is being “held back”. It is not true of enabling harm in general that harm will inevitably occur.

The fact that harm to someone is inevitable, and that the act of enabling is aimed at self-preservation, makes enabling by ducking more easily justifiable than enabling harm in a non-threatening situation, or for some end other than self-defence, other things being equal. Nevertheless, ducking is a form of enabling harm, rather than a forbearance to prevent harm. So, a defence of the moral immunity thesis would need to say the following. “Killing an innocent person in self-defence is impermissible, but letting someone die by removing a protective obstacle in the course of preserving one’s life is permissible.”

Now, in order to establish that the killing and letting die distinction justifies a permission to duck a threat, but not to kill, we need a case of killing a threat or bystander in which harm is enabled in self-defence, to see whether that is permissible. Consider the following case adapted from an example provided by Kasper Lippert-Rasmussen.³⁰

²⁹ Arguments for the claim that enabling harm is often morally equivalent to letting die can be found in Samuel C. Rickless, “The Moral Status of Enabling Harm”, *Pacific Philosophical Quarterly*, 92 (2011), 66-86.

Arguments for the contrary claim that enabling harm is sometimes morally equivalent to killing can be found in Jonathan Bennett, “Morality and Consequences”, in *The Tanner Lectures on Human Values Volume Two* (Salt Lake City: University of Utah Press, 1981), p.89, Shelly Kagan, *The Limits of Morality* (Oxford: Oxford University Press, 1991), p.101, Kasper Lippert-Rasmussen, “Are Enabling and Allowing Harm Morally Equivalent?”, *Utilitas*, 27.3 (2015), 365-83, and Jason Hanna, “Enabling Harm, Doing Harm, and Undoing One’s Own Behaviour”, in *Ethics*, 126 (2015), 68-90.

³⁰ Lippert-Rasmussen, “Are Enabling and Allowing Harm Morally Equivalent?”, p. 372.

Hungry Bear

You and an innocent bystander are trapped in cages which do not belong to anybody. A hungry bear lurks outside the cages. Unfortunately for you, the door to the cage which contains you is unlocked, and the bear is able to get inside. The other cage, which contains the bystander, is secured by a heavy padlock which prevents the bear from reaching him. You remove the padlock to the bystander's cage in order to secure your own cage. The bear eats the bystander.

Taking the lock seems wrong, despite the fact that it involves the redistribution of a threat through the removal of protection. Moreover, the death of the bystander is a side-effect of the agent taking the lock. This suggests that the fact that ducking a threat involves enabling harm by removing a protective object does not, by itself, make ducking a threat justifiable.

One might resist this conclusion by denying that the *Hungry Bear* case should be classed as a case of enabling harm. One might maintain that the killing and enabling distinction is more complicated than the foregoing discussion suggests because it correlates with a further underlying feature: whether a person deprives somebody of an object that they are rightfully entitled to use.³¹

In McMahan's detailed examination of removing or withdrawing aid or protection, he says that the killing and letting die distinction depends upon a variety of subfactors, suggesting that some cases of enabling harm by removing a barrier count as killing, whereas others count as letting die.³² He proposes that an agent kills somebody if they remove an obstacle that another agent or event provided, *or* the protective obstacle is operative and self-sustaining. By contrast, an agent lets someone die if they have provided the protection, and the protection requires further contributions from the agent to remain effective.³³

³¹ The suggestion that the distinction between killing and letting die is underpinned by the deeper moral significance of rights is made in Fiona Woollard, "If This is My Body: A Defence of the Doctrine of Doing and Allowing", *Pacific Philosophical Quarterly*, 94 (2013), 315-41, Timothy Hall, "Doing Harm, Allowing Harm, and Denying Resources", *Journal of Moral Philosophy*, 5 (2008), 50-76, Kai Draper, "Rights and the Doctrine of Doing and Allowing" and Kagan, *The Limits of Morality*, pp.101-106.

³² Jeff McMahan, "Killing, Letting Die, and Withdrawing Aid", *Ethics*, 103.2 (1993), 250-79.

³³ McMahan, "Killing, Letting Die, and Withdrawing Aid", p.257.

These distinctions help explain why taking the lock to the cage seems wrong. Since the lock to the cage was not provided by the agent, does not require further contributions from the agent to remain effective, and is already actively providing protection from a potential harm, the agent's interference with the lock is a variety of killing.

However, while McMahan's account successfully accommodates the *Hungry Bear* case, a problem for this view is presented by the *Car* case described earlier. In that case, the agent is not responsible (either morally or causally) for providing the protection, and the protection provided by the car is both operative and self-sustaining; it requires no further contribution from the agent to remain effective. So, on McMahan's account, the *Car* case is, factually, an instance of killing.

Yet the *Car* case seems intuitively morally equivalent to "pure" cases of ducking harm – that is, ducks which involve only the use of the one's body. Pure cases of ducking are examples of enabling on McMahan's account. Since the *Car* case does not seem to differ in permissibility from a "pure" ducking case, one can infer that the killing and enabling distinction, as formulated by McMahan, does not make a difference to the permissibility of acting in self-defence. Whichever formulation of the killing and enabling distinction is employed, the moral immunity thesis runs into difficulties.

4.3.5 Property Rights and the Fair Distribution of Chances

One might resist the conclusion reached above by objecting that, contrary to appearances, the *Car* case is morally inequivalent to a "pure" ducking case. One could do this by combining the self-ownership thesis explored above with an additional account of rights over the external world. One could say: "When an indivisible object or resource is at stake, to which nobody has an exclusive claim, the right thing to do is to hold a fair lottery to decide who gets to use it. By contrast, it is justifiable to duck a threat without a lottery when evasion requires *only* the use of one's body. That is because each person has a property right in their own body; it is not a common resource to be used by others."

One might accordingly maintain that in the *Car* case, the right thing to do would be to hold a fair lottery to decide whether to leave the car in place or allow the person trapped inside to move it away from the threat.³⁴ This runs counter to the intuitive judgement

³⁴ The view that defensive harm against innocent threats may be justified after a fair lottery is defended by Susanne Burri in "The Toss-Up Between a Profiting, Innocent Threat and His Victim", *The Journal of Political Philosophy*, 23.2 (2015), 146-65. She does, however, say that one is permitted to duck a

that it is permissible to move the car without holding a fair lottery. However, one might be willing to bite this bullet in the interest of reconciling the various judgements presented so far. Moreover, the analysis is supported by two good arguments.

First, people's judgements may shift between cases in which an agent removes a protective object which belongs to them, cases in which it belongs to the victim, and cases in which neither the agent, nor the victim, has a prior claim to the object.³⁵ For instance, suppose we modify the *Car* case so that the car is the rightful property of the prospective victim further down the tracks. One might think it is wrong to move the victim's own property in self-preservation when it is protecting them from a threat. One might also think that if neither the agent nor the prospective victim has an exclusive claim to the car, the right thing to do is to distribute the opportunity to use the object equally.³⁶ Ducking a threat usually differs from other cases of enabling harm in an important respect: one's own body is the object to be removed and so the other person had no claim to use it.

I would resist this first move for two reasons. First, although something does seem morally amiss about contributing to the death of a bystander by using an object or resource to their detriment, this does not seem sufficiently important to override the partiality that each person has over their own lives.³⁷ Even when fairness requires that the prospect of harm be equally distributed, prospective victims may usually give priority to their own lives by securing survival without a fair lottery, as long as this does not involve actions which are independently wrong. If that were correct, then it would be permissible to use the car to evade the threat without a lottery. From this, one could conclude that the *Car* case is, after all, morally equivalent to a "pure" duck.

threat without a fair lottery, on the grounds that one may only impose a chance of harm if an agent would be morally required to suffer harm in order to avoid imposing it on others.

³⁵ For examples of this shift in judgement, see Woollard, "A Defence of the Doctrine of Doing and Allowing", p.318, and Draper, "Rights and the Doctrine of Doing and Allowing", pp.259-263.

³⁶ For his part, Otsuka denies that fairness can justify overriding moral constraints against bodily incursion (in "Killing the Innocent in Self-Defence", p.76, footnote 9) but he says that fairness can justify wresting a life-saving benefit from somebody after a favourable and fair coin toss, if one can do so without incursion on the person's body (in "The Moral Responsibility Account of Liability to Defensive Killing", p.13).

³⁷ This proposed solution to the ducking puzzle is raised in D. W. Haslett, "Boulders and Trolleys", *Utilitas*, 23.3 (2011), 268-87 (p.280). The solution is also contemplated by Boorse and Sorenson in "Ducking Harm", pp.122-124.

Moreover, once it is conceded that fairness requires us to equally distribute the chance to use an indivisible object, it becomes difficult to resist the conclusion that fairness requires the equal distribution of the indivisible good of continued survival itself. After all, it is not the use of the lock or car *per se* that agents are interested in, it is the prospects for survival offered by the use of those objects. Each person's claim to life-saving resources is a surrogate for a claim to the outcome in which they live, and fairness requires that the prospects of each outcome should be fairly distributed.

A second argument for the claim that the *Car* case is inequivalent to a "pure" ducking case is to suggest that the *Car* case is structurally equivalent to the *Hungry Bear* case. In each case, the agent moves an object which is benefitting the potential victim, and which they do not have an exclusive right to use. Yet while the *Car* case seems permissible, it seems wrong to take the padlock in the *Hungry Bear* case. So, if the cases are indeed factually equivalent, one of the intuitive judgements must yield to the other.

I would, however, resist this move. It is true that the cases are alike in several respects. However, they differ in the following way. In the *Car* case, the fate of each person is tied together by the object; the car which affords protection for the person on the tracks is *the very same object* which traps the other person in the path of the threat. Indeed, the presence of the car is part of the threatening sequence from the perspective of the agent trapped inside.³⁸ It would be better from their perspective if the car was absent from the scene.

By contrast, although the lock in the *Hungry Bear* case could afford protection to either person, it is not a part of the causal sequence which exposes the agent to the threat. The presence of the lock is potentially beneficial to the agents (or at least, to whichever agent uses it). I conclude from this that the agent's claim to take the lock in the *Hungry Bear* case is much weaker than the agent's claim to move the car, and that this explains the difference in judgement between the cases.

4.4 Two Objections

I have concluded my examination of the five moral features to which I think one might appeal in order to morally distinguish killing from ducking. However, there are two final lines of defence to which a proponent of the moral immunity thesis may appeal to resist my arguments. First, one might protest that while individual features might establish moral parity between in individual cases of killing and ducking, it is the

³⁸ Thanks to Fiona Woollard for emphasising this point.

accumulation or cluster of these morally significant features which amounts to a difference in moral status. Since each feature identifies one respect in which one act may be harder to justify than another, these features, taken together, amount to a difference in permissibility, even if none of the features alone can mark a boundary between permissible and impermissible self-defence.

Typically, ducking enables harm as a side-effect of redistributing of a pre-existing threat. Yet even when we take all of these features together and export them into a case of killing a bystander, we still find that they do not appear to make a difference to permissibility. This is illustrated by the *Nitrous Oxide* and *Hungry Bear* cases, both of which contain each of those factual features, but nevertheless appear to be impermissible acts of self-defence.

A second line of defence for the moral immunity thesis is to deny that ducking is a cause of harm, or, alternatively, to maintain that its causal contribution to harm is less morally significant than killing. Directly “acting on” someone seems to be a kind of causation with moral significance, whereas bringing something about by “acting at a distance” intuitively lacks the same moral significance. For example, one might suggest there is no “spatiotemporally continuous causal chain” from ducking to the harm that befalls the victim.³⁹

Yet this line of resistance would have potentially wide-reaching implications for the moral status of enabling harm through the removal of protection: it would be highly counterintuitive to deny that unblocking a threat by removing a protective obstacle is a cause of any subsequent harm. Alternatively, if ducking *is* a cause of harm, but a morally insignificant one, this would be best explained by the further moral factors I have already explored (for example, the differences between direct harms and side-effect harms, or the differences between creating a new threat and redistributing a pre-existing threat).

Before concluding, I must point out that the mere fact that the moral immunity thesis has a counterintuitive implication is insufficient reason to reject it, since other accounts of the morality of self-defence may have similarly counterintuitive implications. However, as I see it, the moral status of ducking presents a difficulty for the moral immunity thesis in particular; other accounts of the morality of self-defence are well placed to accommodate the judgement that ducking is usually permissible, without equivalently counterintuitive implications of their own.

³⁹ Draper, “Rights and the Doctrine of Doing and Allowing”, pp. 266-267.

For example, Jonathan Quong has argued that it is wrong to kill someone in the course of defending one's own life if one uses a person or their rightful property (including the space that they occupy) as a means of securing survival. Accordingly, it is wrong to use people as human shields and ride roughshod over innocent bystanders who obstruct one's path to safety.

By contrast, he says it is permissible to kill the innocent in self-defence when one does so without using them or their property. So it is permissible to kill people who innocently threaten one's life, to duck threats and even to kill bystanders as a side-effect of action aimed at self-preservation. The rationale for these claims is simply that each person is morally permitted to exercise partiality when acting in defence of their life, provided that certain conditions are met. This account has a plausible rationale and provides results which coincide with considered judgements in a range of cases.

The crucial difference between Quong's account and the moral immunity thesis as it pertains to ducking is that Quong's account permits the killing of bystanders as a side-effect of action aimed at self-preservation. Consequently, the arguments that I provide in Section 4.3.2 do not apply to Quong's view. This might move us to conclude that the moral immunity thesis has shortcomings which should push us to favour Quong's account instead.

4.5 Conclusion

At the outset, I contrasted two cases in which a person who is endangered by an innocent threat can defend themselves using different modes of agency in each case: by killing the innocent threat in the first case and by ducking the threat in the second case. I pointed out that some maintain that it would be *wrong* to *kill* the innocent threat, on the grounds that doing so is morally indistinguishable from killing an innocent *bystander* in the course of defending one's own life, but that most believe that it is permissible to *duck* the threat.

I then set out my central claim: *if* killing an innocent threat is morally equivalent, in terms of permissibility, to killing an innocent *bystander*, then both are equivalent to *ducking* a threat. As a result, if the moral immunity thesis is correct, then ducking the threat is *impermissible*. Since this is highly counterintuitive, I suggested that we should instead reject (or revise) the moral immunity thesis.

I offered a brief positive defence of this claim and then considered five ways in which a proponent of the moral immunity thesis may seek to resist it: by appealing to the right

of self-ownership, the side-effect principle, the redistribution principle, the distinction between killing and letting die, and the moral significance of fair chances and property rights. Using counterexamples, I argued that none of these factors make a difference to the permissibility of self-defence without undermining the moral immunity thesis.

Although ducking a threat intuitively differs in moral status from killing innocent threats and bystanders, this intuition is difficult to accommodate with the moral immunity thesis without relinquishing either the moral equivalence thesis or the inviolability of a bystander thesis. Since this is a serious shortcoming, I submit that the account should be seriously revised or rejected in favour of accounts which do not have this shortcoming.

BIBLIOGRAPHY

Adler, Matthew, “The Pigou-Dalton Principle and the Structure of Distributive Justice” (Working Paper, 2013)

Amiel, Yoram and Frank Cowell, *Thinking About Inequality* (Cambridge: Cambridge University Press, 2009)

Aquinas, Thomas, *Summa Theologia*, 2.2, 66.7

Atkinson, Anthony B., and Andrea Bandolini, “Unveiling the Ethics Behind Inequality Measurement”, *The Economic Journal*, 125 (2015), 209-34

Bennett, Jonathan, “Morality and Consequences”, *The Tanner Lectures on Human Values Volume Two* (Salt Lake City: University of Utah Press, 1981)

Boorse, Christopher and Roy A. Sorensen, “Ducking Harm”, *The Journal of Philosophy*, 85.3 (1988), 115-34

Boorse, Christopher, “Ducking Trolleys”, *The Journal of Social Philosophy*, 25.3 (1994), 146-52

Burri, Susanne, “The Toss-Up Between a Profiting, Innocent Threat and His Victim”, *The Journal of Political Philosophy*, 23.2 (2015), 146-65

Butt, Daniel, “On Benefitting From Injustice”, *Canadian Journal of Philosophy*, 37.1 (2007), 129-52

Carens, Joseph H., “An Interpretation and Defence of the Socialist Principle of Distribution”, *Social Philosophy and Policy*, 20.1 (2003), 145-77

Carens, Joseph H., “Rights and Duties in an Egalitarian Society”, *Political Theory*, 14.1 (1986), 31-49

Carens, Joseph H., “The Egalitarian Ethos as a Social Mechanism”, ed. by Alexander Kaufman, *Distributive Justice and Access to Advantage* (Cambridge: Cambridge University Press, 2014)

Casal, Paula, "Occupational Choice and the Egalitarian Ethos", *Economics and Philosophy*, 29.1 (2013), 3-20

Christiano, Thomas, *The Constitution of Equality* (Oxford: Oxford University Press, 2008)

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

Cohen, G. A., "Fairness and Legitimacy in Justice, And: Does Option Luck Ever Preserve Justice?" in *On the Currency of Egalitarian Justice and Other Essays in Political Philosophy* ed. by Michael Otsuka (Princeton: Princeton University Press, 2011)

Cohen, G. A., "How To Do Political Philosophy" in *On the Currency of Egalitarian Justice and Other Essays in Political Philosophy*, ed. by Michael Otsuka (Princeton: Princeton University Press, 2011)

Cohen, G. A., *If You're an Egalitarian, How Come You're So Rich?* (Cambridge, MA: Harvard University Press, 2002)

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

Cohen, G. A., *Self-Ownership, Freedom, and Equality* (Cambridge: Cambridge University Press, 1995)

Cullity, Garrett, *The Moral Demands of Affluence* (Oxford: Clarendon Press, 2004)

Dalton, Hugh, "The Measurement of the Inequality of Incomes", *The Economic Journal*, 30.199 (1920), 348-61

Draper, Kai, "Rights and the Doctrine of Doing and Allowing", *Philosophy and Public Affairs*, 33.3 (2005), 253-80

Dworkin, Ronald, *Sovereign Virtue* (Cambridge, MA: Harvard University Press, 2000)

Edmundson, William, "Distributive Justice and Distributed Obligations", *The Journal of Moral Philosophy* (2017)

Edmundson, William, "Ought We Do What We Ought to Be Made to Do?" in *Reasons and Intentions in Law and Practical Agency*, ed. by George Pavlakos and Veronica Rodriguez-Blanco (Cambridge: Cambridge University Press, 2015)

- Estlund, David, "Liberalism, Equality, and Fraternity in Cohen's Critique of Rawls", *The Journal of Political Philosophy*, 6.1 (2008), 99-112
- Fabre, Cécile, "Distributive Justice and Freedom: Cohen on Money and Labour", *Utilitas*, 22.4 (2010), 393-412
- Fabre, Cécile, *Whose Body is it Anyway? Justice and the Integrity of the Person* (Oxford: Oxford University Press, 2006)
- Firth, Joanna, "What's So Shameful About Shameful Revelations?", *Law, Ethics, and Philosophy*, 1 (2013), 31-51
- Fischer, John Martin and Mark Ravizza, "Ducking Harm and Sacrificing Others", *Journal of Social Philosophy*, 25.3 (1994), 135-45
- Foot, Philippa, *Virtues and Vices and Other Essays in Moral Philosophy* (Oxford: Clarendon Press, 2002)
- Frowe, Helen, "Equating Innocent Threats and Bystanders", *Journal of Applied Philosophy*, 25.4 (2008), 277-90
- Frowe, Helen, *Defensive Killing* (Oxford: Oxford University Press, 2014)
- Hall, Timothy, "Doing Harm, Allowing Harm, and Denying Resources", *Journal of Moral Philosophy*, 5 (2008), 50-76
- Hanna, Jason, "Enabling Harm, Doing Harm, and Undoing One's Own Behaviour", *Ethics*, 126 (2015), 68-90
- Haslett, D. W., "Boulders and Trolleys", *Utilitas*, 23.3 (2011), 268-87
- Hobsbawm, Eric J., *Bandits* (London: Abacus, 2001)
- Hooker, Brad, *Ideal Code, Real World: A Rule-Consequentialist Theory of Morality* (Oxford: Clarendon Press, 2000)
- Horton, Keith, "Fairness and Fair Shares", *Utilitas*, 23.1 (2011), 88-93
- Julius, A.J. "Basic Structure and the Value of Equality", *Philosophy and Public Affairs*, 31.4 (2003), 321-55
- Kagan, Shelly, *The Limits of Morality* (Oxford: Oxford University Press, 1991)

Kamm, Frances, *Morality, Mortality Volume II: Rights, Duties and Status* (Oxford: Oxford University Press, 1996)

Kamm, Frances, "Harming, Not Aiding, and Positive Rights", *Philosophy and Public Affairs*, 15 (1986), 3-32

Lang, Gerald, "Rawlsian Incentives and the Freedom Objection", *The Journal of Social Philosophy*, 47.2 (2016), 231-49

Lippert-Rasmussen, Kasper, "Are Enabling and Allowing Harm Morally Equivalent?", *Utilitas*, 27.3 (2015), 365-83

Maysh, Jeff, 'The Uncatchable', *BBC News Magazine*, 25 September 2014

McMahan, Jeff, "Killing, Letting Die and Withdrawing Aid", *Ethics*, 103.2 (1993), 250-79

McMahan, Jeff, "Self-Defence and the Problem of the Innocent Attacker", *Ethics*, 104.2 (1994), 252-90

McTernan, Emily, "The Inegalitarian Ethos: Incentives, Respect and Self-Respect", *Politics, Philosophy and Economics*, 12.1 (2013), 93-111

Miller, David, "Equality and Justice", *Ratio*, 10.3 (1997), 222-37

Miller, David, "Taking Up The Slack? Responsibility and Justice in Situations of Partial Compliance", *Responsibility and Distributive Justice* ed. by Carl Knight and Zofia Stemplowska (Oxford: Oxford University Press, 2011)

Murphy, Liam B., "Institutions and the Demands of Justice", *Philosophy and Public Affairs*, 27.4 (1999), 251-91

Murphy, Liam, *Moral Demands in Non-Ideal Theory* (Oxford: Oxford University Press, 2000)

Nagel, Thomas, "Equality", *Mortal Questions* (Cambridge: Cambridge University Press, 1979)

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

O'Neill, Martin, "Constructing a Contractualist Egalitarianism: Equality After Scanlon", *The Journal of Moral Philosophy*, 10 (2013), 429-61

O'Neill, Martin, "What Should Egalitarians Believe?" *Philosophy and Public Affairs*, 36.2 (2008), 119-56

Otsuka, Michael, "Freedom of Occupational Choice", *Ratio*, 21.4 (2008), 440-53

Otsuka, Michael, "Prerogatives to Depart From Equality", *Royal Institute of Philosophy Supplement*, 58 (2006), 95-111

Otsuka, Michael, "Equality, Ambition, and Insurance", *Proceedings of the Aristotelian Society Supplementary Volumes*, 78 (2004), 151-66

Otsuka, Michael, "Killing the Innocent in Self-Defence", *Philosophy and Public Affairs*, 23.1 (2004), 74-94

Otsuka, Michael, "The Moral Responsibility Account of Liability to Defensive Killing" in *The Ethics of Self-Defence* ed. by Christopher Coons and Michael Weber, (Oxford: Oxford University Press, 2016)

Otsuka, Michael, *Libertarianism Without Inequality* (Oxford: Clarendon Press, 2003)

Øverland Gerhard, and Christian Barry, "Do Democratic Societies Have a Right to Do Wrong?", *The Journal of Social Philosophy*, 42.2 (2011), 111-31

Parfit, Derek, "Equality or Priority?" in *The Ideal of Equality*, ed. by Matthew Clayton and Andrew Williams (Basingstoke: Palgrave, 2002)

Parfit, Derek, *On What Matters* (Oxford: Oxford University Press, 2011)

Penalver, Eduardo Moises and Sonia K. Katyal, *Property Outlaws: How Squatters, Pirates and Protestors Improve the Law of Ownership* (New Haven: Yale University Press, 2010)

Porter, Thomas, "The Division of Moral Labour and the Basic Structure Restriction", *Politics, Philosophy and Economics*, 8 (2009), 173-99

Quinn, Warren, "Actions, Intentions, and Consequences: The Doctrine of Double Effect", *Philosophy and Public Affairs*, 18 (1989), 334-51

Quong, Jonathan "Justice Beyond Equality", *Social Theory and Practice*, 36.2 (2010), 315-40

Quong, Jonathan, "Agent Relative Prerogatives to Do Harm", *Criminal Law and Philosophy*, 10 (2016), 815-829

Quong, Jonathan, "Contractualism, Reciprocity and Egalitarian Justice", *Politics, Philosophy, and Economics*, 6.1 (2007), 75-105

Quong, Jonathan, "Killing in Self-Defence", *Ethics* 119.3 (2009), 507-37

Radin, Margaret Jane, "Property and Personhood", *Stanford Law Review*, 34.5 (1982), 957-1015

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

Rickless, Samuel C. "The Moral Status of Enabling Harm", *Pacific Philosophical Quarterly*, 92 (2011), 66-86

Rothman, Benny, *The 1932 Kinder Trespass: A Personal View of the Kinder Scout Mass Trespass* (Altrincham: Willow Publishing, 1982)

Scanlon, T. M., "The Diversity of Objections to Inequality" in *The Difficulty of Tolerance* (Cambridge: Cambridge University Press, 2003)

Scott, James C. *Two Cheers for Anarchism* (Princeton: Princeton University Press, 2012)

Segall, Shlomi, *Why Inequality Matters: Luck Egalitarianism, Its Meaning and Value* (Cambridge: Cambridge University Press, 2016)

Shelby, Tommie, "Justice, Deviance, and the Dark Ghetto", *Philosophy and Public Affairs*, 35.2 (2007), 126-60

Singer, Peter, "Famine, Affluence and Morality", *Philosophy and Public Affairs*, 1.1 (1972), 229-43

Smilansky, Saul, "On Practicing What We Preach", *American Philosophical Quarterly*, 31.1 (1994), 73-79

Stanczyk, Lucas, "Productive Justice", *Philosophy and Public Affairs*, 40.2 (2012), 144-64

Steiner, Hillel, "Levels of Non-Ideality", *The Journal of Political Philosophy* (2017)

- Stemplowska, Zofia, "Non-ideal Theory" in *A Companion to Applied Philosophy* ed. by Kasper Lippert-Rasmussen, Kimberley Brownlee, and David Coady (Chichester: Wiley Blackwell, 2017)
- Tan, Kok Chor, "Justice and Personal Pursuits", *Journal of Philosophy*, 101.7 (2004), 331-62
- Temkin, Larry, "Exploring the Roots of Egalitarian Concerns", *Theoria*, 69 (2003), 125-51
- Temkin, Larry, *Inequality* (Oxford: Oxford University Press, 2003)
- Temkin, Larry, *Rethinking the Good: Moral Ideals and the Nature of Practical Reasoning* (Oxford: Oxford University Press, 2012)
- Thomson, Judith Jarvis, "Self-Defense", *Philosophy and Public Affairs*, 20.4 (1991), 283-310
- Titelbaum, Michael G., "What Would a Rawlsian Ethos of Justice Look Like?", *Philosophy and Public Affairs*, 36.3 (2008), 289-322
- Tomlin, Patrick, "Internal Doubts About Cohen's Rescue of Justice", *The Journal of Political Philosophy*, 18.2 (2010), 228-47
- Unger, Peter, *Living High and Letting Die: Our Illusion of Innocence* (Oxford: Oxford University Press, 1996)
- Valentini, Laura, "Ideal and Non-Ideal Theory: A Conceptual Map" in *Philosophy Compass*, 7.9 (2012), 654-64
- Vallentyne, Peter, "Justice, Interpersonal Morality, and Luck Egalitarianism", *Distributive Justice and Access to Advantage* ed. by Alexander Kaufman (Cambridge: Cambridge University Press, 2014)
- Waldron, Jeremy, *Liberal Rights: Collected Papers, 1981 – 1991* (Cambridge: Cambridge University Press, 1993)
- Wellman, Christopher Heath, and A. John Simmons, *Is There a Duty to Obey the Law? For and Against* (New York: Cambridge University Press, 2005)

Williams, Andrew, "Incentives, Inequality, and Publicity", *Philosophy and Public Affairs*, 27.3 (1998), 225-47

Wolff, Jonathan, "Fairness, Respect, and the Egalitarian Ethos", *Philosophy and Public Affairs*, 27.2 (1998), 97-122

Wolff, Jonathan, "Levelling Down" in *Challenges to Democracy: The PSA Yearbook 2000*, ed. by Keith Dowding, Jim Hughes, and Helen Margetts (Macmillian, 2001)

Wollner, Gabriel, *On Permissibly Resisting Moderate Injustice* (Working Paper)

Woollard, Fiona, "If This is My Body: A Defence of the Doctrine of Doing and Allowing", *Pacific Philosophical Quarterly*, 94 (2013), 315-41