The London School of Economics and Political Science

# Competing Claims, Risk and Ambiguity

Thomas Rowe

A thesis submitted to the Department of Philosophy, Logic and Scientific Method of the London School of Economics for the degree of Doctor of Philosophy, London, March 2017

## Declaration

I certify that the thesis I have presented for examination for the MPhil/PhD degree of the London School of Economics and Political Science is solely my own work other than where I have clearly indicated that it is the work of others (in which case the extent of any work carried out jointly by me and any other person is clearly identified in it).

The copyright of this thesis rests with the author. Quotation from it is permitted, provided that full acknowledgement is made. This 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 55,692 words.

## Statement of conjoint work

I confirm that Chapter 3 was jointly co-authored with Alex Voorhoeve and I contributed 60% of this work.

Fore

Thomas Rowe

#### Abstract

This thesis engages with the following three questions. First, how should the presence of risk and ambiguity affect how we distribute a benefit to which individuals have competing claims? (In line with common use in decision theory, a case involves *risk* when we can assign at least subjective probabilities to outcomes and it involves *ambiguity* when we cannot assign such probabilities.) Second, what is it about the imposition of a risk of harm *itself* (that is, independently of the resulting harm), such as the playing of Russian roulette on strangers, which calls for justification? Third, in the pursuit of the greater (expected) good, when is it permissible to foreseeably generate harms for others through enabling the agency of evildoers?

Chapters 1 through 3 of the thesis provide an answer the first question. Chapter 1 defends the importance of a unique complaint of unfairness that arises in risky distributive cases: that sometimes individuals are better off at the expense of others. Chapter 2 defends a view called *Fairness as Proper Recognition of Claims* which guides how a decision-maker ought to act in cases where individuals have unequal claims to a good. Chapter 3 considers how the presence of ambiguity affects distributive fairness, and defends an egalitarian account of the evaluation of ambiguous prospects.

Chapter 4 provides an answer to the second question through a defence of the *Insecurity Account*, which is a unique way in which impositions of risks of harm can be said to harm individuals, namely by rendering the victim's interests less secure.

Chapter 5 provides an answer to the third question by defending what I call the *Moral Purity Account*, to explain when it is permissible to provide aid in cases where individuals are harmed as a foreseeable consequence of the provision of such aid.

### Acknowledgements

I would like to thank Alex Voorhoeve for being an excellent supervisor. Alex provided extremely helpful feedback on very many pieces of my work. Alex always had time for me when I had questions about my ideas, and would return extremely detailed feedback often within days and sometimes within hours. I found our discussions very rewarding, and from them I learnt a lot. I am extremely grateful for his support throughout the PhD.

My second supervisor, Mike Otsuka, provided very helpful feedback on multiple drafts of three of my chapters. I am very grateful for the care and attention he paid to my work. I would also like to thank Luc Bovens for taking a keen interest in my work on Chapter 5, and for providing very helpful feedback on a draft.

My thanks also go to Richard Bradley and those who were part of the AHRC Managing Severe Uncertainty research project. I learnt a lot from the project, and found being part of it a very rewarding experience. I also received my PhD funding from this project, for which I am extremely grateful.

I had the good fortune to be in the company of some great people whilst at the LSE. Thanks to my fellow PhD students for providing a wonderful atmosphere. Special thanks to Chris Marshall, Adam White, Asbjørn Aagaard Schmidt, Silvia Milano, Phillipe van Basshuysen, Nicolas Wuethrich, Todd Karhu, Goreti Faria and Bastian Steuwer. I am particularly grateful to Chris, Adam and Phillipe for their sanity-preserving company – more often than not in the White Horse. I thank Goreti and Adam for kindly taking some of my teaching in order to give me more time to spend on the thesis.

Unending gratitude goes to my parents, Sue and Martin, and to my sisters Sophie and Lucy. I thank them for letting me pursue whatever it was that I was interested in. Their support has been of immense value to me.

Finally, and most of all, many thanks to Eva-Maria Lohwasser for her continual love and support over the duration of the thesis, especially in the last few months. I am forever grateful.

# **Table of Contents**

| List of Tables  | 7      |
|---|--------|
| Introduction  | 8      |
| Chapter 1 – Risk and the Unfairness of Some Being Better (            | Off at |
| the Expense of Others   | 16     |
| 1.1 Introduction  | 16     |
| 1.2 Single-Person Case and Two-Person Intrapersonal Case              | 17     |
| 1.3 Independent Risks and Inequality                                  | 25     |
| 1.4 Inversely Correlated Risks and the Complaint that Some are Better | Off at |
| the Expense of Others   | 35     |
| 1.5 Competing Claims and the Separateness of Persons                  | 47     |
| 1.6 Conclusion  | 51     |
| Chapter 2 – Fairness and the Satisfaction of Unequal Claims           | 52     |
| 2.1 Introduction  | 52     |
| 2.2 Fair Distribution and Unequal Claims                              | 53     |
| 2.2.1 Claims  | 53     |
| 2.2.2 Unequal Claims and the Direct Receipt of an Indivisible Good    | 54     |
| 2.2.3 Lazenby's Argument  | 55     |
| 2.2.4 Proportionate Satisfaction                                      | 57     |
| 2.3 Three Objections to Lazenby's Fairness Ranking                    | 61     |
| 2.3.1 Marginally Stronger Claims Objection                            | 61     |
| 2.3.2 Multi-Person Objection  | 64     |
| 2.3.3 Relational Unfairness Objection                                 | 66     |
| 2.4 An Account of Weighted Lotteries                                  | 67     |
| 2.4.1 The Argument  | 67     |
| 2.4.2 Premise One   | 68     |
| 2.4.3 Premise Two   | 71     |
| 2.4.3.1 Broome and Surrogate Satisfaction                             | 74     |
| 2.4.4 Fairness as Proper Recognition versus FLOS                      | 76     |
| 2.4.5 Summary   | 78     |
| 2.5 Objections and Replies  | 78     |
| 2.5.1 Does Fairness Require a Weighted Lottery?                       | 78     |
| 2.5.2 "But the Claim is Stronger!"                                    | 79     |
| 2.5.3 Very Unequal Claims   | 80     |

| 2.6 Conclusion  |
|---|
| Chapter 3 – Ambiguity and Fairness90                                  |
| 3.1 Introduction  |
| 3.2 Ellsberg Cases and Distributive Fairness                          |
| 3.2.1 Pareto Under Ambiguity and Fairness92                           |
| 3.3 Inequality in Final Well-being and Individual Versus Social-Level |
| Ambiguity   |
| 3.4 Conclusion  |
| <b>Chapter 4 – When is it Permissible to Impose a Risk of harm?</b>   |
| 4.1 Introduction  |
| 4.2 Accounts of the Moral Significance of Risking109                  |
| 4.2.1 No Harm Account   |
| 4.2.2 The Wronging Account118   |
| 4.2.3 The Autonomy Account122   |
| 4.2.4 The Preference Account127                                       |
| 4.2.5 Summary   |
| 4.3 Risk Impositions and the Insecurity of Interests                  |
| 4.3.1 The Insecurity Account132                                       |
| 4.3.2 Premise One134  |
| 4.3.3 Premise Two135  |
| 4.3.4 Objections and Replies136                                       |
| 4.4 Conclusion  |
| Chapter 5 – Providing Aid and Foreseeing Harm                         |
| 5.1 Introduction  |
| 5.2 Motivation  |
| 5.3 Theorising the Aid Case146  |
| 5.3.1 Complicity  |
| 5.3.2 The Doctrine of Double Effect148                                |
| 5.3.3 Rubenstein's Account150   |
| 5.4 The Moral Purity Account155                                       |
| 5.5 Rubenstein and Responsibility172                                  |
| 5.6 Conclusion  |
|   |
| Bibliography176   |

# List of Tables

| Table 1.1: Final utilities for Single-Person Case 18                 |
|--|
| Table 1.2: Final utilities for Two-Person Intrapersonal Case20       |
| Table 1.3: Final utilities for Fully Independent Risk Case           |
| Table 1.4: Final utilities for Modified Independent Risk Case29      |
| Table 1.5: Final utilities for Inversely Correlated Case             |
| Table 1.6: Final utilities for Certainty Case 48                     |
| Table 1.7: Final utilities for Intra- Versus Interpersonal Trade-off |
|  |
| Case   |

### Introduction

#### **Core Questions**

This thesis engages with the following three questions. First, how should the presence of risk and ambiguity affect how we distribute a benefit to which individuals have competing claims? (In line with common use in decision theory, I shall say that a case involves *risk* when we can assign at least subjective probabilities to outcomes and that it involves *ambiguity* when we cannot assign such probabilities.) Second, what is it about the imposition of a risk of harm *itself* (that is, independently of the resulting harm), such as the playing of Russian roulette on strangers, which calls for justification? Third, in the pursuit of the greater (expected) good, when is it permissible to foreseeably generate harms for others through enabling the agency of evildoers?

Chapters 1 through 3 of the thesis provide an answer the first question. Chapter 1 defends the importance of a unique complaint of unfairness that arises in risky distributive cases: that sometimes individuals are better off at the expense of others. Chapter 2 defends a view called *Fairness as Proper Recognition of Claims* which guides how a decision-maker ought to act in cases where individuals have unequal claims to a good. Chapter 3, which is joint work with Alex Voorhoeve, considers how the presence of ambiguity affects distributive fairness, and defends an egalitarian account of the evaluation of ambiguous prospects.

Chapter 4 provides an answer to the second question through a defence of the *Insecurity Account*, which is a unique way in which impositions of risks of harm can be said to harm individuals, namely by rendering the victim's interests less secure.

Chapter 5 provides an answer to the third question by defending what I call the *Moral Purity Account*, to explain when it is permissible to provide aid in cases where individuals are harmed as a foreseeable consequence of the provision of such aid.

#### **Competing Claims**

The theme of "competing claims" is the thread that runs through the thesis. The first three chapters go some way towards providing an account of how one ought to act in distributive cases with competing claims, taking into account variables such as the type of risks – whether they are, for instance, correlated or independent – (Chapter 1), the number of individuals with claims of varying strength (Chapter 2) and various levels of probabilistic information (Chapter 3).

Before proceeding, I shall briefly comment on the nature of competing claims.<sup>1</sup> Suppose I have a medicine, which I myself have no use for, and an individual needs this medicine in order to survive. This individual has a *claim* on me that they receive the medicine. The strength of this claim can be measured by how much this individual's well-being would be increased by receiving the good and by how badly off they are compared to others (Voorhoeve, 2014: 66). Individuals do not have claims to a good if their interests are not at stake. Now suppose that there are two individuals who both require the medicine in order to survive. If and only if the claims of these individuals cannot be jointly satisfied, then these claims are *in competition* with one another. The chapters in this thesis revolve around such cases. It is worth noting that claims may be influenced by other factors as

<sup>&</sup>lt;sup>1</sup> The competing claims approach has its roots in the work of Thomas Nagel (1979: 106-27; 1991). The most popular exposition of a competing claims approach can be found in the development of contractualist moral theory (Scanlon, 1998).

well, such as the nature of the agency involved in producing a benefit or harm. Indeed, Chapter 5 considers how one ought to act when one set of individuals has a claim to medical assistance, and another group has a claim against being killed by a villainous aggressor.

One particular virtue of the competing claims approach is that it respects both the unity of individuals and what is often called the "separateness of persons".<sup>2</sup> The former establishes that one has reason to balance benefits and burdens for an individual with an eye towards their prospective good. (This is true both across times in an individual's life – e.g. a burden today may be outweighed by a somewhat larger benefit later in their life – and across an individual's potential futures – e.g. a chance of a burden for them may be outweighed by an equal chance of a greater benefit to them). The latter establishes that, because individuals lead separate lives, it is inappropriate to balance benefits and burdens *across* individuals as if they were a super-individual. A present or possible burden for an individual can be compensated by a future or possible benefit, whereas a burden to one individual cannot straightforwardly be compensated by giving a benefit to another person. The views defended in the first three chapters are appropriately sensitive to the unity of the individual and the separateness of persons.

#### **Risk and Ambiguity**

Sven Ove Hansson has claimed that throughout history "moral theorizing has referred predominantly to a deterministic world in which the morally

<sup>&</sup>lt;sup>2</sup> Utilitarianism famously fails to respect the separateness of persons because it treats a collection of individuals as if they were a super-individual (Rawls, 1999: 26-7). Otsuka (2012) presses separateness of persons objections to prioritarianism. Voorhoeve and Fleurbaey (2012) offer a competing claims account which respects the separateness of persons.

relevant properties of human actions are both well-determined and knowable. ... Ethics still lives in a Newtonian world." (2003: 291). The project of determining how one ought to act under conditions of risk and ambiguity is of great importance, since we do not live in such a world. There has been a recent blossoming in the ethics of risk literature, particularly in distributive ethics (see, e.g. Broome, 1984; Otsuka and Voorhoeve, 2009; Parfit, 2012; Hyams, 2015), and the ethics of imposing risks of harm (see, e.g. McCarthy, 1997; Finkelstein, 2003; Lenman, 2008; Oberdiek, 2009 & 2012). This thesis makes contributions to these two important areas.<sup>3</sup>

Furthermore, the thesis enters a new area in the ethics of risk, namely distributive ethics under conditions of *ambiguity*. An individual is in an ambiguous situation when they are not able to assign precise subjective probabilities to states of the world and the associated possible outcomes of their actions. The vast majority of the ethics of risk literature considers cases where it is at least possible in principle to assign precise probabilities. This altogether bypasses the common scenario in which it is often not possible to establish precise probabilities of states of the world. Distributive ethics under conditions of ambiguity is therefore of great practical importance.

#### **Outline of the Thesis**

The thesis can be categorised roughly into two parts. The first part (Chapters 1-3) examines questions of distributive fairness under conditions of risk and

<sup>&</sup>lt;sup>3</sup> Another recent off-shoot of the ethics of risk literature concerns the ability of contractualism to accommodate risk. Since contractualism, in the form outlined by Scanlon, focuses on undiscounted complaints of individuals to a particular policy, it seems to be "too confining", since activities that impose a small risk of death on everyone whilst also giving each a small benefit (e.g. the benefits everyone receives from permitting air travel, vs. the small risk that there will be one death from falling plane debris) will be outlawed because the complaint of the one who will die will outweigh the small inconvenience of those who forgo a small benefit (Scanlon, 1998; Reibetanz, 1998; Ashford, 2003; James, 2012; Frick, 2015; Kumar, 2015). I will bracket an explicit discussion of this off-shoot of the literature.

ambiguity. In such cases, individuals have competing claims to a good and a morally-motivated decision-maker has to decide how to distribute the good. The second part (Chapters 4-5) considers the relationship between imposing risk, harming, and benefitting. The outlines of the five chapters are as follows.

Chapter 1 asks how a morally-motivated decision-maker ought to act under conditions of risk when (i) only the interests of one person are at stake; and (ii) the interests of multiple other people are at stake. The chapter defends a competing claims account to guide how a decision-maker ought to act in such cases. Against recent writings of Michael Otsuka, Alex Voorhoeve, and Marc Fleurbaey, it argues that one ought to maximise each person's expected utility not merely in single-person cases without inequality, but also in many-person cases with guaranteed (or likely) outcome inequality, when such inequality is the result of independent risks. At the centre of this account is a distinct complaint of unfairness: that some will (or might be) made better off at the expense of others. This occurs when the fates of individuals are tied in such a way that inequality that is to the detriment of the person who is worse off is guaranteed (or likely) to obtain. The chapter therefore considers how different types of risk (e.g. independent and inversely correlated risks) can make a moral difference to how one should act.

Chapter 2 continues on the theme of distributive ethics under conditions of risk, and asks what fairness requires when individuals have unequal claims to a good. Some have argued that what fairness requires depends on the strength of claims of those who have their interests at stake (Broome, 1990; Piller, forthcoming). On this view, when claims are very unequal, fairness

requires giving the good directly to the person with the strongest claim, and when claims are slightly unequal, fairness requires the use of a weighted lottery. Others have questioned the contribution to fairness made by weighted lotteries and have argued that fairness requires that the person with the strongest claim ought to receive the good directly (Hooker, 2005: 348-9; Lazenby 2014). The chapter argues that fairness requires a weighted lottery in such cases, and in demonstrating this, outlines and defends an account of distributive fairness called *Fairness as Proper Recognition of Claims*. This account provides a justification for lotteries in cases with unequal claims and is, it will be argued, appropriately sensitive to changes to the number and strength of claims. The chapter makes a contribution to the recent fair distribution literature by providing a novel account for the justification of weighted lotteries in cases with indivisible goods.<sup>4</sup>

Chapter 3 considers the ethics of distribution under conditions of ambiguity. It argues against the use of a version of the *ex ante* Pareto principle applied to ambiguous prospects, or "Pareto under Ambiguity", for short. It rejects this principle on two grounds. One is familiar from discussions of the *ex ante* Pareto principle under risk, namely that it conflicts with egalitarian concerns (see, e.g. Fleurbaey and Voorhoeve 2013). The other is novel: we show that prospects that are ambiguous at the individual level may be far less so at the social level. An ambiguity-averse distributor, who rightly considers the

<sup>&</sup>lt;sup>4</sup> This is an underrepresented position within the contemporary debate on the ethics of fair distribution. Broome (1990) and Piller (forthcoming) consider the fairness of weighted lotteries for *some* cases of unequal claims, but my account argues that fairness *requires* the use of weighted lotteries in such cases. To my knowledge, this is not a position that has been extensively defended. However, weighted lotteries have often been defended in the literature of "saving the greater number", where it is either possible to save one individual, or a group of individuals (Kamm, 1993; Otsuka, 2006; Saunders, 2009). The cases I discuss do not involve options where it is possible to save a whole group, but rather involve separate individuals.

latter level, may therefore correctly favour a set of ambiguous individual prospects that eliminates social ambiguity, even when one would not prefer this set of individual prospects on any individual's behalf.

Chapters 1 and 3 outline new reasons to depart from what would maximise each individual's expected utility in distributive cases under conditions of risk and ambiguity. Chapter 1 argues that one should sometimes depart from what is best for each individual taken separately when there is a complaint that one is better off at the expense of others. Chapter 3 argues that one should sometimes depart from what is best for each individual because of social ambiguity-aversion.

Chapter 4 moves onto the question of when it is permissible to impose a risk of harm on others. The chapter asks what it is that gives impositions of risk themselves (that is, even when they do not result in material harm) moral significance, and thereby be the sort of thing that can be impermissible. The chapter examines, and rejects, four candidate views of the moral significance of risking. The first states that risks of harm are not *themselves* harms, and therefore, that if imposing a risk of harm were to possess moral significance it is not because of its harmfulness (Perry, 1995, 2003). A second view states that the potential impermissibility of imposing a risk arises from the fact that a risky act may itself wrong the victim. A third view is that an imposition of a risk of harm's impermissibility stems from the fact that it curtails the victim's autonomy (Oberdiek, 2009, 2012). A fourth view uses a preferencebased account to argue that because these risks are dispreferred by the person exposed to them, an imposition of risk of harm is *itself* a harm, and thereby possesses moral significance (Finkelstein, 2003). In contrast to all these accounts, the chapter outlines and defends a novel account of the

moral significance of imposing a risk of harm, called the *Insecurity of Interests Account*. It is argued that what gives an imposition of a risk its moral significance is that it renders the victim's interests less secure. As well as providing a unique account of what it is that is morally significant about impositions of risk, the account complements the value of chances view that is defended in Chapter 1 by outlining why it is that risks of harm can have disvalue for an individual.

Chapter 5 considers a case of competing claims between those who wish to be saved from harm by rescuers and those who wish to avoid being harmed by evildoers. The chapter asks the following question: In the pursuit of the greater (expected) good, when is it permissible to foreseeably contribute to harm on innocent others through enabling the agency of evildoers? The chapter outlines and defends what is called the *Moral Purity Account* to both explain when it is permissible to provide aid in such cases and determine when aiders bear moral responsibility for the harms imposed by another. This is contrasted with a recent account by Jennifer Rubenstein (2015) which she calls the *Spattered Hands* view. According to this view, aiders sometimes have a responsibility to grudgingly accept contributing to injustices perpetrated by others. When they do so, they bear a degree of moral responsibility for those harmful actions. It is argued that the *Moral Purity* Account is superior to Rubenstein's account, and it is claimed that when suitably motivated, aiders are not blameworthy for the actions of evildoers. This chapter stakes out terrain in an area that is relatively under-theorised, and has great practical importance.

## Chapter 1

# Risk and the Unfairness of Some Being Better Off at the Expense of Others<sup>5</sup>

#### 1.1 Introduction

How should a morally-motivated decision-maker act under conditions of risk when (i) only the interests of one person are at stake; and (ii) the interests of multiple other people are at stake?<sup>6</sup> This chapter defends a competing claims account to guide how a decision-maker ought to act in such cases. At the centre of this account is a distinct complaint of unfairness: that sometimes some are better off at the expense of others.

The chapter has three subsidiary aims. The first is to demonstrate how the type of risk that one is exposed to can make a moral difference to how one should act. There has been little systematic discussion of this question. The second is to outline and defend a unique complaint of unfairness: that sometimes some are better off at the expense of others. I argue that competing claims obtain if and only if some might be better off at the expense of another. The third is to contrast my account of competing claims with that of the recent view of Michael Otsuka, Alex Voorhoeve and Marc Fleurbaey. My view differs from theirs in the respect that it limits the variety of cases in which competing claims complaints obtain.

<sup>&</sup>lt;sup>5</sup> Earlier versions of this chapter were presented at the Brave New World conference in Manchester, and at the Warwick Graduate Conference in Legal and Political Theory. I am grateful to the audiences for helpful comments.

<sup>&</sup>lt;sup>6</sup> By morally-motivated stranger I am referring to a private individual who can at very little cost to herself come to the assistance of other strangers.

The structure of the chapter is as follows. In Sections 1.2 and 1.3 I argue that when risks are independent, i.e. when the prospects for one person in a gamble do not depend on the prospects of another, it is permissible for a morally-motivated decision-maker to provide an expected well-being maximising alternative, regardless of the number of people involved, and that the potential for outcome inequality does not itself give rise to complaints of unfairness. In Section 1.4 I argue that, in cases where risks are inversely correlated, i.e. when the prospects for one person have an inverse relation to the prospects of another, a source of individuals complaints against a situation in which they are disadvantaged is that some are better off at their expense. This complaint, I argue, should lead us to be averse to inequality in inversely correlated cases. I contrast my account to Otsuka, Voorhoeve and Fleurbaey's recent view. In Section 1.5 I demonstrate how the account guides action in cases involving certainty, as well as how it respects the distinction between the unity of the individual and the separateness of persons. Section 1.6 concludes.

#### **1.2 Single-Person Case and Two-Person Intrapersonal Case**

Consider the following case:

Single-Person Case: Ann, a young child, who is currently in full health, will soon go completely blind through natural causes (utility = 0.65) unless Tessa, a morally-motivated stranger, provides one of two available treatments. Treatment *A* will either, with 50 percent probability, leave Ann blind (0.65) or instead, with 50 percent probability, fully cure her (utility =1); treatment *B* will restore Ann to partial sight for sure (0.8). (See Table 1.1)

|             | <i>S1</i> (0.5) | S2 (0.5) |
|-------------|-----------------|----------|
|             | Ann             | Ann      |
| Treatment A | 0.65            | 1        |
| Treatment B | 0.8             | 0.8      |

Table 1.1: Final utilities for Single-Person Case.

Before judging this case, I need to clarify the measure of well-being employed. I shall assume a measure of utility derived from idealized preferences satisfying the Von Neumann-Morgenstern axioms. According to this measure, a prospect has higher expected utility for a person "just in case it would be preferred for that person's sake after rational and calm deliberation with all pertinent information while attending to her selfinterest only. (A person's expected utility is just the probability-weighted sum of her utility in each state of the world.) One prospect has the same expected utility as another for a person just in case such deliberation would yield indifference between the two prospects" (Otsuka and Voorhoeve, forthcoming: 8-9).

Now, supposing only Ann's interests are at stake, which treatment should Tessa select? Treatment *B* gives Ann a set outcome for sure, but it offers a lower expected well-being than treatment *A*. In this case, it is possible to offer the following prudential justification for providing *A*: "I did the best I could for you given the information I had at the time" (Voorhoeve and Fleurbaey, 2016: 935). This provides, I believe, a strong reason to provide *A*. Moreover, insofar as the individual's interests are considered in isolation

from others, there is no countervailing reason to favour *B*. As Michael Otsuka puts it, there are "no interpersonally comparative or otherwise distributive considerations … that tell in favour of paying heed to anything other than what is in this [person's] rational self-interest" (Otsuka, 2015: 5). Why, when one is considering this individual's interests alone, ought one to depart from what, given the information available, anyone concerned exclusively with the individual's interests (including the individual herself) would rationally regard as the best one can do for her?

For these reasons, in this chapter, I shall assume the following answer is correct: in Single-Person Case, Tessa should maximize Ann's expected wellbeing and hence choose *A*.<sup>7</sup> The correctness of this answer has recently been much debated.<sup>8</sup> I shall not revisit this debate here. Instead, I shall pursue the following, less thoroughly discussed question: How should a view that accepts the premise that one should maximize an individual's expected wellbeing when one considers her fate in isolation deal with multi-person risky cases, in which the risky alternative that is in the expected interests of each person, taken separately, will generate outcomes in which some end up better off and others worse off than they would under a less risky alternative?<sup>9</sup> This question is important because some who accept my

<sup>&</sup>lt;sup>7</sup> The measure does not presuppose anything about what the nature of well-being is in itself. For instance, one may think that well-being consists in the satisfaction of preferences or the presence of happiness and absence of suffering. The measure is consistent with a decisionmaker maximizing whatever it is they take well-being to be. I will use the terms "utility" and "well-being" as synonymous in this article.

<sup>&</sup>lt;sup>8</sup> Otsuka and Voorhoeve (2009) and Otsuka (2015) defend the permissibility of maximizing well-being in the Single-Person Case. McCarthy (2008; forthcoming), Greaves (2015), and Voorhoeve and Fleurbaey (forthcoming) defend the view that one ought to maximize this person's expected well-being. For the contrary view that one is permitted to be risk averse in a person's well-being, see, e.g. Parfit (2012: 432) and Bovens (2015: 404).

<sup>&</sup>lt;sup>9</sup> Many views accept this premise. Most noteworthy are the views defended by Otsuka and Voorhoeve (2009; 2011; forthcoming), Voorhoeve and Fleurbaey (2012; 2013; 2016) and Johann Frick (2013: 130-133; 2015: 186-191).

premise hold that in such multi-person risk cases, each individual has a claim only to what would maximize their expected well-being, and those who end up worse off as a result of the distributor's choice of such an alternative have no complaint (Frick, 2013: 144-5; 2015: 181-8). Against such a view, I shall argue that within a claims-based framework there can sometimes be reason to select an alternative that does not maximize each person's expected well-being, namely when choosing what would maximize each person's expected well-being would ensure that some end up better off at the expense of others.

Now consider the following case, inspired by Voorhoeve and Fleurbaey (2012: 386):

*Two-Person Intrapersonal Case*: This case is identical to the preceding case with the addition of an extra person, Bill, who has partial sight for sure (0.8), no matter what treatment Tessa selects for Ann. Bill's well-being is completely unaffected by Tessa's action. If Tessa selects treatment *A*, then Ann will with 50 percent probability either remain blind (0.65) or instead, with 50 percent probability, be fully cured (1); treatment *B* will restore Ann to partial sight for sure (0.8). (See Table 1.2.)

Table 1.2: Final utilities for Two-Person Intrapersonal Case:

|             | S1 ( | 0.5) | S2 (0.5) |      |  |
|-------------|------|------|----------|------|--|
|             | Ann  | Bill | Ann      | Bill |  |
| Treatment A | 0.65 | 0.8  | 1        | 0.8  |  |
| Treatment B | 0.8  | 0.8  | 0.8      | 0.8  |  |

In this case, Tessa ought to, in my view, select A. Only Ann's interests are at stake in this case, just like in Single-Person Case. Just as it was reasonable to select a treatment in Single-Person Case to maximize Ann's expected wellbeing, it should also be the case here. The only difference to the structure of the example is that now there is an extra individual, Bill, who is completely unaffected by Tessa's decision. If we believe that only people who have their well-being at stake in a gamble have a complaint, or potential complaint, against the actions that a decision-maker will take, then Tessa ought to select A, as in Single-Person Case. To lend support to this idea, an interpretation of contractualism, the Complaint Model, states that: "a person's complaint against a principle must have to do with its effects on him or her, and someone can reasonably reject a principle if there is some alternative to which no other person has a complaint that is as strong" (Scanlon, 1998: 229).<sup>10</sup> The absence of any effects on Bill from Tessa's action means that Bill does not have a complaint against the provision of A. This captures an important feature of the separateness of persons.

One way to argue against the selection of A in this case is to claim that the morally-motivated decision-maker ought to care about the fact that there will be inequality if A is selected. For instance, there may be a brute luck egalitarian reason to favour the selection of B for this ensures that both Ann and Bill will have a well-being level of 0.8. Brute luck egalitarians believe that "it is bad, or objectionable, to some extent—because unfair—for some to be worse off than others through no fault or choice of their own" (Temkin, 1993: 13). Even if Ann and Bill were on separate continents and have

<sup>&</sup>lt;sup>10</sup> However, Scanlon also argues that there may be other considerations other than wellbeing which may lead to grounds for complaints, such as complaints of unfairness (1998: 219).

absolutely no relationship with one another, inequality between them would be bad to the extent that it did not follow from a choice of theirs.

One could object to this reason for favouring *B* on the grounds that it is an alternative that is not in the expected best interests of Ann, and *only* her interests are at stake. The fact that some are worse off than others through no fault or choice of their own can sometimes generate a reason for selecting a more egalitarian alternative over an alternative that does not necessarily guarantee outcome equality. But in Two-Person Intrapersonal Case only the interests of Ann are at stake. It would be difficult to justify her choice if Tessa selected an option that guaranteed equality, but failed to give Ann what was in her best expected interests. Following this, Two-Person Intrapersonal Case lends strong support to the claim that the mere pattern of inequality resulting from a gamble is no basis for individual complaints. This is because bill can't reasonably complain about the ensuing outcome inequality since nothing that is selected for Ann would affect his well-being. Although inequality obtains if *A* is selected, the only effective claim that can be lodged to Tessa is that she do what is in the best expected interests of Ann.

Some brute luck egalitarians may argue that we should not always be guided by individual complaints, but should sometimes be guided by *impersonal value*. Perhaps brute luck equality is impersonally valuable. My response is that although it may be granted that brute luck inequality can sometimes provide reasons for action, there is no good reason for such impersonal considerations to outweigh the individual reasons that Ann possesses in favour of the selection of *A*. Given that only Ann's interests are at stake in this case, there is no good reason to allow anything other than the

rational self-interest of Ann to determine what treatment ought to be selected.<sup>11</sup> If *A* is selected then there will be outcome inequality for sure, but there will be no reasonable complaints against this. Bill could complain and say that the fact that there is this impersonal complaint against outcome inequality should mean that *B* is selected. It may be the case that the fact of inequality means that in *some* respect the state of affairs is bad for luck egalitarian reasons, but the presence of a prudential justification lends support to selecting *A* in this case, as in Single-Person Case.

A relevant distinction can be drawn here between *consequentialist* complaints and *personal* complaints. Consequentialist complaints track the alleged badness of patterns of consequences. It is implausible that such complaints should determine what Tessa ought to do in Two-Person Intrapersonal Case because only the interests of Ann are at stake, so only considerations relevant to Ann's fate should determine how Tessa ought to act. Further, considerations of the pattern of inequality do not bear on the well-being level of individuals. Bill will remain at the same level of well-being if treatment *A* or *B* is selected, since his interests are not at stake. Personal complaints track individuals' reasons for complaint. The only claim in Two-Person Intrapersonal Case is on behalf of Ann for Tessa to select the alternative that is in her expected best interests. If Ann ends up badly off, a prudential justification can be given to her. If Ann ends up better off than Bill, Bill has no reasonable personal complaint that he can raise, because he could not have been made better off.

<sup>&</sup>lt;sup>11</sup> As such, it may be the case that brute luck inequality can provide *reason* for action, but such a consideration fails to outweigh the considerations in favour of selecting what is best for Ann, since only Ann's well-being is at stake in Tessa's choice.

I do not think that the inequality present in Two-Person Intrapersonal case is a sufficient reason to favour treatment *B*. This is because only Ann's interests are at stake, and as such only her interests should guide how Tessa ought to act. The claims of individuals are what should guide action in cases such as Two-Person Interpersonal Case because claims are responsive to the level of well-being that an individual has at stake (i.e. the prospective increase in or decrease in that individual's well-being), unlike impersonal considerations such as the pattern of inequality, which is not sensitive to such considerations.

Ittay Nissan-Rozen (2017) has argued that a distributor has a pro tanto reason to discard impersonal reasons when deciding what reasons to take into account when distrusting a resource to individuals that individuals have a claim to. In support of this claim, Nissan-Rozen appeals to the Kantian demand to treat individuals as ends in themselves and not as mere means. The impersonal reason of equality has nothing to do with Ann's well-being. By selecting an alternative in line with this impersonal reason when only Ann's interests are at stake, "the distributor treats the [impersonal reason] itself ... as the end of the act of weighing it; and so, by weighing [the impersonal reason], the distributor treats [Ann] as a means to an end: she treats her as a means for the end of satisfying [the impersonal reason]" (Nissan-Rozen, 2017: 5). As such, there is further reason to be sceptical of the claim that in cases where individuals have claims grounded in their (expected) well-being, one ought to sometimes act in accordance with impersonal reasons. My view is that inequality is necessary but not sufficient for a complaint of unfairness. The presence of inequality does not guarantee a complaint of unfairness in distributive cases under conditions of risk. Other conditions are required for a complaint of unfairness to arise.

#### **1.3 Independent Risks and Inequality**

I shall now consider a case involving two individuals with the difference that now *both* individuals' well-being is at stake in Tessa's decision.

*Fully Independent Risk Case*: There are two children, Ann and Bill, who will soon go completely blind through natural causes (utility = 0.65) unless Tessa, a morally-motivated stranger, provides one of four possible treatment alternatives. She can give treatment *A* to Ann and treatment *B* to Bill or vice versa, or she can give both individuals treatment *A* or both treatment *B*. Ann and Bill's outcomes under the treatments are statistically independent of each other. (See Table 1.3.)

|            | <i>S1</i> (0 | .25) | S2 ( | 0.25) | S3 ( | 0.25) | <i>S4</i> (0 | ).25) |
|------------|--------------|------|------|-------|------|-------|--------------|-------|
|            | Ann          | Bill | Ann  | Bill  | Ann  | Bill  | Ann          | Bill  |
| Treatments |              |      |      |       |      |       |              |       |
| Ann: A     | 0.65         | 1    | 1    | 0.65  | 0.65 | 0.65  | 1            | 1     |
| Bill: A    |              |      |      |       |      |       |              |       |
| Ann: B     | 0.8          | 0.8  | 0.8  | 0.8   | 0.8  | 0.8   | 0.8          | 0.8   |
| Bill: B    |              |      |      |       |      |       |              |       |
| Ann: A     | 0.65         | 0.8  | 1    | 0.8   | 0.65 | 0.8   | 1            | 0.8   |
| Bill: B    |              |      |      |       |      |       |              |       |
| Ann: B     | 0.8          | 1    | 0.8  | 0.65  | 0.8  | 0.65  | 0.8          | 1     |
| Bill: A    |              |      |      |       |      |       |              |       |

Table 1.3: Final utilities for Fully Independent Risk Case:

Tessa ought to select *A* for both individuals in this case. The same reasoning that supported the selection of *A* in the preceding cases supports the same

selection in Fully Independent Case. Ann and Bill's futures are in every sense independent from one another, since (i) it is possible to offer either treatment to each individual independently of which treatment is offered to the other, and (ii) under each alternative any risks they face are independent. This implies that the well-being of Ann and Bill is *separate* in the same way that it is in Two-Person Intrapersonal Case. By this I mean their potential futures are not linked together: the well-being value of each individual is not causally dependent on the well-being values of other individuals. For example, if Tessa chooses *A* and Ann ends up with a particular level of well-being, then this has no bearing on what level of well-being Bill ends up with.

It is important to distinguish my use of "separateness" here from what is known in welfare economics as "additive separability". If one tries to order distinct distributions of individual's well-being, one may believe that such orderings are "additively separable", which contains the thought that the moral value of "each person's well-being should be evaluated independently of other people's wellbeing" (Broome, 2015: 220-1). I shall *not* endorse this idea of additive separability. As I shall argue below, in determining the moral value of individuals' well-being, comparisons between their well-being matter just in case their fates are "tied together" in a particular manner. Rather than additive separability, I am instead endorsing an account of separability that is consistent with the separateness of persons, in that it tracks cases in which these futures are thoroughly independent, or unlinked.

In Fully Independent Case (Table 1.3) there is the possibility of inequality between Ann and Bill if Tessa chooses *A* for one or both of them. In Two-Person Intrapersonal Case (Table 1.2), inequality was certain to occur if *A* 

was selected, but this did not form the basis of individual complaints. How should Tessa accommodate the possibility of outcome inequality given that the interests of both individuals are now at stake? I think that, again, the potential inequality if A is selected for one or both of them is no basis for individual claims against selecting A. What is of importance in this case, as in the preceding two cases, is the separateness of Ann and Bill's prospects. The potential futures of Ann are distinct from Bill's potential futures; moreover, nothing that is decided by Tessa about Ann's future affects Bill's fate; nor does anything that happens by chance to Ann affect how Bill ends up, and vice versa. In this regard, the expectably best treatment for Ann leaves Bill's well-being unaffected, and the expectably best treatment for Bill leaves Ann's well-being unaffected. Furthermore, there is a decisive reason to select treatment A for both Ann and Bill given the strength of the prudential justification that can be offered to both, based on the fact that the prudential justification appeals to these unified potential futures of each person.

One could object to the selection of *A* for both individuals by arguing that, while there is no complaint against inequality per se, this is a case in which, if we choose *A* for both, some may end up better off and others worse off than they might as a consequence of our choices, so that there may be competing claims *ex post* between the better and the worse off. And when there are such competing claims, whoever will end up worse off has a claim to an alternative in which they would have ended up better off (so to be given *B* in this case). Such reasoning is suggested by Voorhoeve and Fleurbaey (forthcoming, 10-11) and a definition of this approach is provided by Matthew Rendall, who states: "whenever we distribute benefits and

burdens among more than one person, the parties have competing claims on our solicitude" (2013: 942). I think that this is too inclusive a conception of the conditions under which there are competing claims. For one, it is not the case that when I am deciding who to give a gift to, that each individual has a claim on my solicitude. I can choose who I give the gift to, but no individual has a *claim* on my gift. Second, and more pertinently to our case, the mere fact that, *ex post*, some may end up better off and others worse off than they might have been is insufficient for the existence of competing claims.

An individual has a claim only if their interests are stake. For example, in Two-Person Intrapersonal Case (Table 1.2) there is only a claim on Ann's behalf since Bill's interests are not at stake. In Fully Independent Case (Table 1.3) the interests of both Ann and Bill are at stake, but they are not in competition because they do not conflict *ex ante*; nor, despite the possibility of inequality, can they conflict *ex post*, as I shall now explain.

If *S1* obtains, then it is best for Ann that *B* is selected. Selecting the best for Ann (*B*) does not preclude selecting the best for Bill (*A*), since it is possible to give *B* to Ann and *A* to Bill. This is because of the separateness of Ann's and Bill's prospects. Nothing that is decided about Ann's fate affects the fate of Bill. Similarly, the treatments affect each individual independently. There is therefore no *ex post* conflict of interest if *S1* obtains. Analogous reasoning establishes the same for every other state of the world. There is therefore no conflict of interest *ex post* in any state of the world.

Since there is no conflict of interest, I conclude that Tessa should simply select what is expectably best for each, which is *A*. Though this may leave one or both of them badly off and one of them worse off than

another, there is a prudential justification that can be given to the person(s) who end up in such a position, and ending up badly off (or worse off) was never the result of doing something to that person that was contrary to the badly off person's interests but in the interests of another.

Now consider the following modification to the preceding case:

*Modified Independent Risk Case*: The set-up of this case is the same as before. However, due to technical limitations, Tessa can either provide both with treatment *A* or both with *B*, but cannot offer one of them *A* and the other *B*. (The case is described in Table 1.4.)

Table 1.4: Final utilities for Modified Independent Risk Case:

|          | <i>S1</i> (0 | .25) | S2 ( | 0.25) | S3 ( | 0.25) | S4 (0 | ).25) |
|----------|--------------|------|------|-------|------|-------|-------|-------|
|          | Ann          | Bill | Ann  | Bill  | Ann  | Bill  | Ann   | Bill  |
| Both get | 0.65         | 1    | 1    | 0.65  | 0.65 | 0.65  | 1     | 1     |
| Α        |              |      |      |       |      |       |       |       |
| Both get | 0.8          | 0.8  | 0.8  | 0.8   | 0.8  | 0.8   | 0.8   | 0.8   |
| В        |              |      |      |       |      |       |       |       |

Does this change anything regarding what Tessa ought to do? One may observe that, as visible in the table above, there is a fifty percent chance that there will be a conflict of interest in final utilities, and that there are therefore competing interests *ex post*. Treatment *A* is rationally preferred by both individuals, but this choice is, in states of the world *S1* and *S2*, better than *B* for one person but worse than *B* for another. Suppose, for instance, that *S2* obtains. Giving *B* to both is best for Ann, but giving *A* to both is best for Bill.

Their *ex post* interests therefore conflict, because both must receive the same treatment.

Does this potential *ex post* conflict of interest in Modified Independent Risk Case make a difference to what Tessa ought to do? It does not. Although it is not possible to provide a different treatment to each individual, the treatments affect each individual separately, in the sense that, if one chooses the risky treatment for both, the fact that one is well off doesn't imply (or increase the chance that) the other is badly off.

Moreover, *A* would have been selected individually if it was possible to do so, as in Fully Independent Case (Table 1.3). The mere fact that options are removed that one would not select anyway should not make a difference to what one ought to do. Suppose that an individual is faced with a number of options, and that one alternative is permissibly chosen from these options. Suppose that we now shrink this set of options by removing one of the unchosen alternatives. The permissible alternative ought to remain permissible in this subset. This is the property of "basic contraction consistency" (Sen, 1993: 500). For example, if giving *A* to both is permissible in Fully Independent Risk Case (Table 1.3), then it should also be permissible in Modified Independent Risk Case (Table 1.4), which contains a subset of the alternatives in the former case.

I shall now outline a view to explain the judgments from all cases considered so far. I call this view the "value of chances" view. The view that chances have value is by no means original (Broome, 1990: 98; Wasserman, 1996: 43). The view is not focused only on the distribution of chances between people, but also on the value that a chance has for an individual. This view claims that an individual's possession of a chance to be advantaged has positive

expected value for that individual. Of course, a chance that does not come good does not (I shall assume) contribute to a person's realized well-being. But, following Voorhoeve and Fleurbaey, I hold that while "the well-being value of a chance evaporates once it is clear that this chance is unrealized (...), a chance's contribution to fairness does not evaporate" (2012: 396). For example, having a chance of winning the lottery or of receiving a muchneeded kidney is of expected value to a person. If two individuals both need a kidney and each is given an equal chance at receiving the kidney, then the chance's contribution to fairness does not disappear, even though one person will be without a kidney. The chance's contribution to fairness does not disappear because it remains true that an individual was given a chance after an outcome obtains.<sup>12</sup> A chance has value for a person in virtue of tracking a possibility of being advantaged. I shall bracket the consideration of whether a chance actually *gives* the individual something of value

<sup>&</sup>lt;sup>12</sup> This can be further illustrated by considering the fact that a chances contribution to fairness can remain even if the chance of a benefit disappears:

|        | Alternative A |           | Alternative B |                      |       |
|--------|---------------|-----------|---------------|----------------------|-------|
|        | States (equi  | probable) |               | States (equiprobable |       |
| People | Heads         | Tails     | People        | Heads                | Tails |
| Ann    | Dies and      | Dies and  | Ann           | Dies                 | Lives |
|        | treated       | treated   |               |                      |       |
|        | unfairly      | unfairly  |               |                      |       |
| Bill   | Lives         | Lives     | Bill          | Lives                | Dies  |

In a scenario where it is possible to select between two different coin flips (Alternative A and B above), a chances contribution to fairness can persist in how we evaluate the goodness or badness of outcomes (Broome, 1991: 113). A chances contribution to fairness remains even when the chance of a benefit disappears, since Ann dying and not having a chance of survival (Alternative A) is worse than Ann dying and having a chance of survival (Alternative A) is worse than Ann dying and having a chance of survival (Alternative A) is worse than Ann dying and having a chance of survival (Alternative B). This allows us to have a principled reason for not being indifferent between the two Alternatives: because in Alternative B equal chances of survival are given to both, which is not the case in Alternative A. If a chances contribution to fairness disappeared once outcomes eventuate, then it would not be possible to establish the greater fairness of Alternative B over Alternative A.

(Wasserman, 1996: 30; Otsuka, ms: 9-10) until Chapter 2, and I shall also remain agnostic on the question of what makes a chance genuine.<sup>13</sup>

To illustrate: the prudential justification that is offered to an individual in a single-person case when a gamble is taken for their sake appeals to the *possibility* of that individual being advantaged even if it turns out that *ex post* they are disadvantaged. If they turn out disadvantaged, the morally-motivated stranger can say that because they had a *sufficiently advantageous chance*, the option was chosen for their sake – it was the best way to promote this person's interests at the time, given the knowledge possessed. As well as chances having prospective value for an individual, the *distribution* of chances between people also possesses value: (Broome, 1990; Voorhoeve and Fleurbaey, 2012: 396). If two individuals are equally in need, and equally deserving of a kidney, a fair way of distributing the kidney would be to give each individual an equal chance at receiving the good, rather than an unequal chance. The distribution of chances can therefore facilitate fairness.

Some have questioned whether the distribution of chances can help facilitate fairness (Hyams, 2017; Wasserman, 1996). Hyams has critiqued the view that the distribution of chances can make outcome inequality less unfair than it would have otherwise been. Hyams does so by arguing that what he calls the "egalitarian mixed view" (the view that outcome inequalities are less unfair when they arise from a less unequal *ex ante* distribution of chances than a more unequal *ex ante* distribution) cannot accommodate for an intuition regarding what we ought to do in the following pair of cases . Suppose that Tessa is charged with allocating a medicine to either Ann or

<sup>&</sup>lt;sup>13</sup> For instance, I am not committed to the claim that genuine chances can be merely subjective (epistemic) chances of an advantage, or instead must be objective.

Bill in the first case, which we can call Initial *Ex ante* Choice. Tessa can either hold a lottery in which either Ann or Bill will receive the medicine, and receive a utility level of 10, or a utility level of 0 if they do not receive the medicine, or Tess can crush the pill and divide it between them, giving each a utility of 4 as the crushing reduces the medicine's effectiveness. Suppose now that in Initial *Ex ante* Choice the lottery device is broken, and has been all along, such that Ann will certainly receive the medicine. Hyams claims that intuitively we would want to switch to giving both 4 for sure rather than holding the lottery once we realise that it is broken. Hyams argues that the egalitarian mixed view struggles to account for why we ought to swap to crushed pill in the modified *ex ante* choice lottery, but not swap to crushed pill in the first version of the case.

To illustrate this, Hyams considers a further case where instead of the lottery device being broken all along, there is a "magical transformation" such that the previously equal objective probabilities of Ann and Bill winning are modified so that Ann receives a probability of 1 of receiving the good, as in the broken lottery case. Hyams argues that we ought to switch from the lottery to crushed pill, which will give 4 to each. Hyams claims that the fact "that [Ann] had an objective 0.5 chance of winning [before the transformation] now seems irrelevant" (Ibid.). This is a challenge to the idea that a prior distribution of chances makes a contribution to fairness, since the fact that there were prior chances is irrelevant to determining whether we should switch to crushed pill once the transformation has occurred. In response, there is reason to think that we ought *not* to switch to crushed pill in the magical transformation case once we know that Ann is certain to win the lottery. The reason is that the "magical transformation" can be viewed

simply as the *result* of the lottery that had equal chances. There were objectively equal chances, and then one individual was declared the winner. As such, the magical transformation in the modified case does no more work than the declared outcome of the equal chance lottery. There are therefore clear grounds for thinking that it is consistent to switch to the crushed pill in the broken lottery case and not in the magical transformation case. This is because the prior equal chances in the magical transformation case are no different to the prior equal chances in the Initial *Ex ante* Choice case.

Wasserman (1996) argues that equal chances do not make a further contribution to fairness over and above the contribution to impartiality. For my view, I appeal to the fact that it is possible for the distribution of chances to facilitate fairness, not the stronger claim that this is the only way to facilitate fairness, or the denial of impartiality based reasons for giving equal chances. Further, it is of value for individuals with equally strong claims to a good to receive equal chances of receiving the good, rather than unequal chances, for the reason that the former is fairer than the latter. It is fair that identically situated people should be treated identically (Broome, 1990: 95). In Chapter Two, I consider in more detail why the distribution of chances can help facilitate fairness.

To summarise, there are at least two related ways in which a chance at an advantage can have value:

- A chance can have positive expected prudential value for an individual;
- 2. Equality in the distribution of chances contributes to fairness.

When providing a prudential justification to Ann for treatment *A* in Single-Person Case (Table 1.1) and Two-Person Intrapersonal Case (Table 1.2), there is an appeal to (1). When providing a prudential justification to Ann and Bill in Fully Independent Risk Case (Table 1.3) and Modified Independent Risk Case (Table 1.4) for giving both *A*, there is an appeal to (1) and (2): each individual had a chance to be advantaged, and each individual has this chance equally.

I have argued that there are no complaints which give us reason to go against the recommendation to provide treatment *A* in these cases. The pattern of inequality arising in Two-Person Intrapersonal Case (Table 1.2) was not sufficient grounds for a complaint on behalf of Bill as he did not have a claim against Tessa because his interests were not affected by her action. Against the idea that the potential inequality in Fully Independent Risk Case (Table 1.3), if *A* were given to both should be of concern, I argued that this inequality need not concern us because this is not a case of competing interests (neither *ex ante* nor *ex post*) between a better-off and the worse-off person. I then appealed to both the independence of Ann and Bill's fates to argue that in Modified Independent Risk Case (Table 1.4) Tessa ought to still give *A* to both.

# 1.4 Inversely Correlated Risks and the Complaint that Some are Better Off at the Expense of Others

In this section, I consider how the view I am defending handles conflict of interest cases and arrive at a characterisation of when people have competing claims. I argue that there is an important moral difference between risks that are independent and risks that are inversely correlated. In the case of inversely correlated risks, the complaint of unfairness is not

solely the outcome inequality itself, but rather the fact that one is better off *at the expense of another*. I argue that this complaint ought to be included in the account of claims I am providing. I argue that in cases of independent risks this particular complaint does not arise, and that this is so even if the *same* pattern of inequality arises as it does in an inversely correlated case.

Consider the following:

*Inversely Correlated Case*: Ann and Bill will both soon, through natural causes, go completely blind unless Tessa administers one of two treatments. Treatment *A*, will either, with fifty percent probability, cause Ann to go blind (0.65) and Bill to retain full vision (1), or instead, with fifty percent probability, cause Bill to go blind and Ann to retain full vision. Treatment *B* will restore both Ann and Bill to partial sight for sure (0.8). (See Table 1.5.)

Table 1.5: Final utilities for Inversely Correlated Case:

|             | <i>S1</i> (0.5) | S2 (0.5) |  |  |
|-------------|-----------------|----------|--|--|
|             | Ann Bill        | Ann Bill |  |  |
| Treatment A | 0.65 1          | 1 0.65   |  |  |
| Treatment B | 0.8 0.8         | 0.8 0.8  |  |  |

Which treatment should Tessa choose? There are reasons that pull in different directions. Firstly, there are considerations in favour of *A* due to the presence of a prudential justification to Ann and Bill, since *A* maximizes both Ann's and Bill's expected well-being. Secondly, there are considerations against the selection of *A*, due to the fact that (I shall argue) one will be
better off at the expense of another. On balance, I shall argue that Tessa ought to select *B* in this case.

In this case, there will always be a conflict of interest *ex post*, whereas in Two-Person Intrapersonal Case (Table 1.2), although there was inequality for certain there was no competing claims complaint. There will be a competing claims complaint in Inversely Correlated Case (Table 1.5) on behalf of whoever turns out to be worse off. This speaks in favour of selecting *B*. In structurally analogous cases, others agree with the selection of *B* (Otsuka, 2012: 373-4; Otsuka and Voorhoeve, 2009: 173-4; Voorhoeve and Fleurbaey, 2013: 118-9). One important consideration is that the identity of the individual who would be better off if *A* were selected, and that of who is worse off is not known in this case - all that is known is that one individual will be in each position. This is important because there is a moral distinction between placeholders and persons. As Johann Frick argues, it makes a difference to the type of justification that can be given to each person in this case: "contractualist justification is owed to persons, with determinate identities and interests, not placeholders in a pattern of outcomes" (Frick, 2013: 141). We might think that because we do not know who will be better off and who will be worse off, Tessa could justify the selection of A to both because she does not know the identity of who will be better off. To see how this might work, consider Frick's argument in this case:

"[Ann] and Bill, if they are self-interested [and competent choosers], would not *want* [Tessa] to choose [*B*]. It is in both persons' *ex ante* interest that we take a gamble on their behalf by choosing [*A*]. The question is: to what extent could either [Ann] or Bill complain of

"outcome unfairness" when any outcome inequality under [*A*] results from having forgone an option, in line with their own selfinterest, that would have satisfied both of their claims to a significant extent and produced no inequality? We might think that, by receiving [*A*], [Ann] and Bill "exchanged" their claim to the significant benefit that they could have gotten from [*B*] in return for the chance of getting an even greater good — a gamble that was in both persons' self-interest. It is not clear that, having made this exchange, either [Ann] or Bill is left with any valid complaint of unfairness if [*A*] does not turn out in [their] favour" (Frick, 2013: 144-5).

Contra Frick, I think that the individuals *do* possess a valid complaint of unfairness. For the complaint of unfairness in this case does not arise merely from the pattern of inequality that results, but rather from the fact that, inevitably, someone is benefitted by the other's misfortune. In this case, the fact that it is not possible for both individuals to be simultaneously better off (or worse off) means that it is possible for a compliant to arise on behalf of the worse off: "someone else is better off *at my expense*". I believe that this is a distinct complaint of unfairness that an individual can raise irrespective of any judgments about the pattern of inequality. This distinct complaint of unfairness arises when one is worse off as a *causal flipside* of someone else benefitting. This view also captures the following two ways in which one can be worse off than another: when one is made worse off as a *causal flipside* of benefiting another, and when one is made worse off as a *side-effect* of benefiting another. In both of these cases one is made worse off as a causal flipside of another it not for

another being burdened. To illustrate the potential objectionableness of having one's fate tied to another, suppose that at the flip of a coin, I receive £1000 if the coin lands heads and you receive nothing, or if the coin lands tails, you receive £1000 and I receive nothing. We might think that there is something intrinsically bad about the fact of inequality between two individuals, whatever the result of the coin flip. But we might *also* think that there is something bad about one individual being better off at the expense of another, when one person being better off is – literally, in this instance – the flipside of another's being worse off.<sup>14</sup>

Part of the grounding for this distinct claim of unfairness comes from the separateness of persons. When an individual's potential futures involve another's interests, a claim of unfairness may arise if the individuals' interests are linked in such a way that inequality may arise. The potential complaint of unfairness can be further illustrated by the following example. Following the structure of Single-Person Case (Table 1.1), suppose Ann ends up badly off after Tessa selects treatment *A* on her behalf. If Ann then learnt that the gamble was one in which the potential to gain was *hers* but that it in fact failed to materialize, then Ann may see her position as justified as she had a large enough chance of being better off to make the gamble in her

<sup>&</sup>lt;sup>14</sup> An interesting feature of inversely-correlated risks is the role that consent might play in legitimatising exposure to such risks. We might think that having one's fate tied to another's arbitrarily is unjust, whereas if one chooses to tie one's fate to another's then this is permissible. For example, Ann and Bill may choose to go to a casino where all of their payoffs are dependent (if Ann wins, then Bill loses and vice versa) or to a casino where all their pay-offs are independent of one another. This would be due to the preferences that Ann and Bill have for risk. Whereas if one entered a casino with the intention of gambling alone (with payoffs independent of others') then one may object if one's gamble turned out to be tied up with another. If Ann loses a gamble then it might appear that Ann lost *at the expense of* Bill. But whether we think it is permissible that one's fate is inversely correlated with the fate of another may come down to whether the inverse correlation is voluntary. Here, I consider only risks that are non-voluntary, because neither Ann nor Bill can make a choice for which they are responsible – only Tessa can.

interest. But if Ann were to then learn that the flip-side of her loss was in fact a gain to another person, she may well think that the other person is better off at her expense, since if she were not worse off, he would not be better off. Suppose that in Single-Person Case, Ann ends up badly off. In this case the potential alternative future of Ann's, where she could have been better off evanesces when it fails to materialize for her. Suppose, now, that Ann ends up badly off in Inversely Correlated Case. Instead of Ann's potential future evanescing where she could have been better off (as in Single-Person Case), it instead falls to another individual, Bill. I believe that there is an important moral difference between these two states of affairs. This is because it is not possible in Inversely Correlated Case for *both* Ann and Bill to be simultaneously better off. Only one person can be better off. Whereas, in the independent risk cases it *is* possible for both to be better off. To illustrate further, I shall adopt and briefly expand on Hugh Lazenby's account of the "uniqueness of experience" (Lazenby, 2014: 339-41). This is "the fact that a person experiences only the one life she actually leads, and not the other possible lives that there was a probabilistic chance that she might have had" (Lazenby, 2014: 340). Returning to the coin flip example, suppose that a faraway potential benefactor flips a coin, and if the coin lands heads I receive £1000. If it lands tails I receive nothing. The coin lands tails. It seems that I have gained nothing of concrete value since I only experience what I actually have. What I think *does* matter however, is whether other individuals experience the other possible futures that were open to me at my expense. If it was the case that when the coin landed tails the £1000 went to another person somewhere else, they are better off than me at my expense, since this potential future of mine (having £1000) goes to them, while I get nothing. Although I do not experience the possible life that I could have

lead, the fact that another individual does experience that life while I do not is something that is in itself potentially objectionable.

For this relationship between fates to be objectionable, it must not only be the case that one person gains and another person loses as a causal flip-side of the others' gain, but also that the one who is worse off is worse off than they would otherwise have been. Although it is the case in the coin-flip example above that one is better off as a flip-side of another being worse off, the person who is worse off than the other is not worse off than they would otherwise have been. It would be strange to think that the loser of the gamble (who ends up no different to before the gamble) could complain that now someone else is better off, since had the coin not been flipped at all, their predicament would be the same. Whereas, in the Inversely Correlated Risk Case it is true that the person who is worse off if A is selected is worse off than they would have been had treatment *B* been selected on their behalf. So it is the fact that one is worse off than one would have otherwise been, coupled with the fact that one's being worse off is the causal flip-side of another person's being better off that grounds the distinct complaint of one being better off at the expense of another.

One may object that in cases like Inversely Correlated Risk Case, if an individual turns out to be worse off, the individual who is better off still had a chance to be better off, so although another individual is in fact better off it is not clear that the better off individual is experiencing the potential future of another since they both had the chance to be better off. In response, it is important that it is not possible in this case for *both* individuals to be better off. If an individual fails to receive the advantage, another person will receive it; their loss is always someone else's gain. This is the core of the

complaint of being worse off at the expense of others. By contrast, in Modified Independent Risk Case, each individual has it in their potential future that they could be better off. It is possible for each to be better off simultaneously.

A further challenge to this view is as follows. Suppose Tessa is facing the decision in Single-Person Case, believing that Ann is exposed to an independent risk if treatment A is selected. Tessa is about to select A. A person then comes along and says "Wait! Are you sure that the risk is fully independent? Might it not in fact be inversely correlated?" One might argue that it doesn't make any difference to what Tessa ought to do in this case, and that she would have no reason to find out whether the risk is in fact inversely correlated. Surely, it may be argued, it is enough that this treatment is in Ann's expected best interests. In response, I claim that it does matter whether the risk that Ann is exposed to is either an independent or an inversely correlated risk. It matters because it determines the sort of complaint that might be available to Ann (if at all) if she ends up badly off. Now suppose that treatment A is selected in this case and Ann ends up badly off. On her way home she discovers someone who was in fact a beneficiary of Tessa's decision to select treatment A. Now Ann has grounds for complaint since someone is better off at her expense. In this scenario it may be objected that the fact there is now a beneficiary ought to count in favour of the justifiability of an action and not against it. In response, the fact that there turned out to be a beneficiary is not *in itself* something that counts against selecting treatment A, but it is rather the causal relationship between the benefit and the burden that matters: that one could only have the benefit if another had the burden.

In order to more precisely relate this distinction between different complaints of unfairness to the cases under discussion, consider the following two propositions:

- (i) Ann is better off at the expense of Bill; or Bill is better off at the expense of Ann.
- (ii) Ann is worse off *and* Bill is better off; Bill is worse off *and* Ann is better off.

Proposition (ii) is true in Two-Person Intrapersonal Case and the pair of independent risk cases when there is outcome inequality, but proposition (i) is not true in these cases. Proposition (i) and (ii) are both true in Inversely Correlated Case. Proposition (ii) obtaining is insufficient for a complaint of unfairness, whereas proposition (i) obtaining is sufficient grounds for a complaint in a conflict of interest distributive case.

Our discussion so far leads to the following characterization of when individuals have competing claims:

*Competing Claims*: Competing claims obtain if and only if individuals' interests conflict *ex ante* or, *ex post*, some might be better off at the expense of another. That is to say, either *ex ante*, all alternatives that are best for one individual are worse for another individual, or there is a chance that, *ex post*, all alternatives that are best for one individual involve that individual being better off at the expense of another.

This view can explain why Tessa ought to select *B* in Inversely Correlated Case. In this case there are competing interests *ex post*. A feature of the complaint of some being better off at the expense of others is that the

interests of individuals are in conflict; a benefit to you will be a loss to me, and vice versa. In Inversely Correlated Case there will always be competing claims as the interests of the better off will always conflict with the interests of who is worse off. This is a decisive reason to select *B* in this case. This view also explains why there are no competing claims in Two-Person Intrapersonal Case and the independent risk cases. All of the alternatives that are best for one individual do not involve that individual being better off at the expense of another. There is no conflict of interest *ex post* in these cases.

My view can be contrasted with Otsuka, Voorhoeve and Fleurbaey's account of competing claims.<sup>15</sup> According to these authors, an individual has a claim to some benefit if and only if their interests are at stake, where the strength of this claim is determined by

- i) his potential gain in well-being; and
- ii) his level of well-being relative to others with whom her interests conflict.<sup>16</sup>

As it stands, this view would state that competing claims exist in the independent risk cases discussed in Section 1.3, because there is an *ex post* conflict of interests between the winners and the losers of the gambles.<sup>17</sup>

<sup>&</sup>lt;sup>15</sup> It is also possible to distinguish my view from the Restricted Prioritarianism view outlined by Andrew Williams (2012). In intrapersonal cases, individuals have a claim on the alternative that maximises their expected utility. However, in interpersonal cases individuals have claims to be benefited that become stronger as their absolute position worsens. Because I do not explicitly engage with prioritarianism in this chapter, I leave out a detailed examination of this view. Benjamin Lange (forthcoming), however, examines the differences between Restricted Prioritarianism and the competing claims view of Otsuka, Voorhoeve and Fleurbaey.

<sup>&</sup>lt;sup>16</sup> This specification of the view is from Benjamin Lange (forthcoming: 4). Voorhoeve and Fleurbaey (2012: 397) also provide a similar outline.

<sup>&</sup>lt;sup>17</sup> One might argue that complaints can be discounted by the probability that inequality would occur. For example, in Fully Independent Risk Case (Table 3) inequality occurs in

Although the claims of the winners and the losers of the gambles in these cases are in competition with one another, in the respect that some are better off than others, it is not the case that the interests of each person are in *conflict* with one another in a way specified by my view. Under my proposal, competing claims do not merely obtain when the interests of individuals are in competition with one another, but rather only when some might be better off at the expense of another. To this extent, my view can be distinguished from Otsuka, Voorhoeve and Fleurbaey's view, by specifying when precisely competing claims obtain. Their view, as it stands, allows for more unfairness complaints than I think are warranted. My view limits the scope of cases where competing claims obtain to those cases where some are made better off at the expense of others. It holds that inequality itself does not determine the existence of complaints, but rather that complaints of unfairness arise from how one's interests relate to others.

To further illustrate the implications of the view that it matters whether one is exposed to an independent risk or an inversely-correlated risk, consider the following case inspired by Roger Crisp (2011: 107):

*Many-person Independent Risk Case*: This is the same as Fully Independent Case, except that there are now 500 people.

What one can infer from this case is that it is extremely likely that there will be outcome inequality if Tessa chooses A – indeed, that the pattern of outcome inequality will be similar to the inequality in Inversely Correlated Case, in that roughly half will be well off and half badly off. We might also

half of the states of the world if treatment A is given to both, and as such one's competing claims complaint at ending up worse off than another could be discounted by its improbability. Nevertheless, a competing claims complaint still exists.

think that here the Law of Large Numbers does the work that inverse correlation does in Inversely Correlated Case (Frick, 2013: 146). Crisp argues that in this case, one ought to select *B* for each child for prioritarian reasons (2011: 107) – in other words, applied to this case, it is more important to improve a person's well-being from 0.65 to 0.8 than from 0.8 to 1. Otsuka and Voorhoeve agree with this conclusion, but for egalitarian reasons (2011: 113). Contrary to these authors, however, I have argued that independent risk cases possess an important moral feature. If inequality obtains, individuals are not worse off to another's benefit. Therefore, there are no competing claims. The outcome inequality does not constitute grounds for individuals' complaints in this case. In this case proposition (ii) will be true and proposition (i) will be false. Even if the pattern of inequality turns out to be identical in Independent Risk Case (Table 1.3) and Inversely Correlated Risk Case (Table 1.5), individuals will not be better off at the expense of others in the former, but will be in the latter.<sup>18</sup>

<sup>&</sup>lt;sup>18</sup> It is worth briefly considering the possibility of competing claims complaints occurring in cases where risks are *partially* correlated. I have so far considered cases where risks are either strictly inversely-correlated or independent. "Correlation between A and B's fates in a risky context obtains if and only if learning about person A's outcome alone should cause one to change one's subjective probabilities for B's outcome. For example, suppose that things were such that if we were to learn that A gets 1, then we would adjust our prior beliefs about the chance that B will get 0, then this would be a case of correlation. To illustrate, suppose that one must choose between certainly giving Ann 0.5 and Bill 0.25 and the alternative outlined in the following table:

|      | Equiprobable outcomes |   |   |   |  |  |  |
|------|-----------------------|---|---|---|--|--|--|
|      | 1                     | 2 | 3 | 4 |  |  |  |
| Ann  | 1                     | 0 | 0 | 1 |  |  |  |
| Bill | 0                     | 1 | 0 | 0 |  |  |  |

Before we learn anything about Ann's outcomes, we should believe that the probability of Bill getting 1 is <sup>1</sup>/<sub>4</sub>. Their outcomes are partially correlated here because we know that if Ann gets 0 then there is a <sup>1</sup>/<sub>2</sub> chance that Bill will get 1, and if Ann gets 1 then Bill will have 0 for sure. There is some inverse correlation, and three quarters of the time, someone will be

Of course, in Inversely Correlated Case, it is true that each individual still has a prospect of being the individual who is better off if *A* is chosen. This remains in favour of selecting *A* in this case, which has to be balanced against the unfairness of being worse off to another's benefit. I submit that when risks are inversely correlated, it is permissible to select *B* for some small advantage in each person's expected well-being of *A* over *B*. In other words, if risks are independent, one ought to select *A* for any d > 0, where *d* is the gain in expected well-being *A* has over the egalitarian alternative. However, because of the unfairness of some being better off at the expense of others, and the associated presence of competing claims, it is permissible to select *B* for some sufficiently small d > 0 when risks are inversely correlated.<sup>19</sup>

#### **1.5 Competing Claims and the Separateness of Persons**

In this section I shall consider how the competing claims account handles conflict of interest cases under certainty. I thereby hope to expand the part of the account which sees the badness of inequality in conflict of interest cases being the fact that some are worse off at the expense of others. I then demonstrate how the view reflects a concern for the separateness of persons.

Consider the following case:

*Certainty Case:* Ann will develop partial sight (0.8) and Bill will go wholly blind (0.65). Tessa can either give treatment *A* which moves Ann from her partial sight to full health (1), or provide treatment *B* 

worse off because of another's success. The complaint to which this gives rise can, I propose, be *discounted* by the probability (seventy-five percent) that this situation will arise. As such, the view I am proposing is able to apply to partially correlated risks." <sup>19</sup> I leave it open to judgment for deciding when the value of the prospects each has to be better off can outweigh the complaint that some are better off at the expense of others.

which moves Bill from complete blindness to partial sight (0.8). (See Table 1.6.)

|             | S1       |  |
|-------------|----------|--|
|             | Ann Bill |  |
| Treatment A | 1 0.65   |  |
| Treatment B | 0.8 0.8  |  |

Table 1.6: Final utilities for Certainty Case:

It is clear that Tessa ought to select B – there are competing claims in this case. The interests of Ann and Bill are in conflict, and providing A will only be in the best interests of Ann and contrary to the interests of Bill, who can never be better off than Ann. Ann starts off better off than Bill and selecting A will further benefit Ann at Bill's expense. On the view I am proposing, this makes A especially problematic. In such cases of certainty the competing claims account has maximal force. Both individuals' claims are in conflict both *ex ante* and *ex post*; it is never possible to have a joint satisfaction of claims. There is no greater value of chances that could outweigh the claim of unfairness in this case, unlike in Inversely Correlated Case (Table 1.5). How could Tessa justify the selection of A to Bill when he is much worse off than Ann, this treatment would never be in *his* best interests and he could gain nearly as much as Ann could if he were treated instead?<sup>20</sup>

Throughout, I have appealed to respect for the difference between the unity of the individual and the separateness of persons in order to justify the view I am proposing. According to the "unity of the individual" an individual's

<sup>&</sup>lt;sup>20</sup> Cf. Otsuka and Voorhoeve (2009: 183-4).

life possesses a unity that makes it appropriate to balance benefits and burdens which accrue to her for her sake, but inappropriate to balance benefits to some with only costs to others (Voorhoeve and Fleurbaey, 2012: 381). This helps to justify the selection of the utility maximising treatment in Single-Person Case since the potential benefits and burdens only accrue to Ann. According to the "separateness of persons", individuals' lives have a separateness that renders it inappropriate to balance benefits and burdens that accrue to each person as if they accrued to a single life (Fleurbaey and Voorhoeve, 2012: 382). It is therefore worthwhile to see how the proposed view handles a case that clearly contrasts intra- and interpersonal trade-offs, adapted from Voorhoeve and Fleurbaey (2012: 387):

*Intra- versus Interpersonal Trade-off Case*: In this case Ann and Bill will both soon go completely blind (0.65) unless Tessa intervenes and selects one of the available treatments. There are two scenarios: one interpersonal and one intrapersonal.

*Interpersonal Scenario*: Treatment *A* (interpersonal) will with 50 percent probability either give Ann an increase in well-being leading to full health (1) while ensuring Bill gains partial sight (0.8), or instead, with 50 percent probability Ann will gain partial sight while Bill remains blind. Treatment *B* will ensure that both Ann and Bill have partial sight

*Intrapersonal Scenario*: Treatment *A* (intrapersonal) will ensure that Bill has partial sight, while giving Ann a 50 percent chance of full health and a 50 percent chance of partial sight. Treatment *B* will ensure that both Ann and Bill have partial sight (See Table 1.7.)

#### Table 1.7: Final utilities for Intra- Versus Interpersonal Trade-off

Case:

|                 | <i>S1</i> (0.5) |      | S2 (0.5) |      |
|-----------------|-----------------|------|----------|------|
|                 | Ann             | Bill | Ann      | Bill |
| Treatment A     | 1               | 0.8  | 0.8      | 0.65 |
| (Interpersonal) |                 |      |          |      |
| Treatment B     | 0.8             | 0.8  | 0.8      | 0.8  |
|                 | ·               |      |          |      |
| Treatment A     | 1               | 0.8  | 0.65     | 0.8  |
| (Intrapersonal) |                 |      |          |      |
| Treatment B     | 0.8             | 0.8  | 0.8      | 0.8  |

In Interpersonal Scenario there is no prospect for Bill to ever be better off than Ann. If Tessa selects *A* then only Ann has a chance at being more advantaged, whereas Bill does not. Bill only has a chance of being disadvantaged. Tessa ought to select *B* in this scenario. The proposed view respects the separateness of persons by being sensitive to some only bearing a potential burden in order for others to only receive a potential benefit. The proposed view respects the unity of the individual by requiring the selection of *A* in Intrapersonal Scenario for Ann's sake, since the potential benefits and burdens only fall to her. Treatment *B* is not permissible in Intrapersonal Scenario because there are no competing claims. By contrast, in Interpersonal Scenario there are competing claims because there is a conflict of interest *ex ante*. The competition between these claims favours *B*, because *A* would involve making one person better off at another's expense. In making these contrasting judgments, the view respects the difference between the unity of the individual and the separateness of persons.

#### **1.6 Conclusion**

I have argued that when the interests of only one person are at stake, one ought to act out of a concern for their best expected interests. A pattern of outcomes that contains inequality does not itself constitute grounds for individual complaints of unfairness. I also argued that the *type* of risk that one is exposed to can affect the permissibility of selecting certain treatments. If individuals are exposed to independent risks, there is both a compelling prudential justification that can be offered that appeals to the separate futures of distinct individuals and there is an absence of competing interests, so that one does not have to balance competing claims on behalf of a better off and a worse off person. By contrast, I have also argued that when there are competing interests (either *ex ante* or *ex post*) then one has reason to avoid inequality that would arise because some are worse off to others' benefit. In sum, I have provided a competing claims account that guides how a decision-maker ought to act in conflict of interest cases involving risk and also in cases of certainty. I have shown that this approach respects both the unity of the individual and the separateness of persons.

## Chapter 2

# Fairness and the Satisfaction of Unequal Claims<sup>21</sup>

## 2.1 Introduction

What does fairness require in cases where individuals have unequal claims to a good? Some have argued that what fairness requires depends on the strength of claims of those who have their interests at stake (Broome, 1990; Piller forthcoming). On this view, when claims are very unequal fairness requires giving the good directly to the person with the strongest claim, and when claims are slightly unequal fairness requires the use of a weighted lottery.<sup>22</sup> Others have questioned the contribution to fairness made by weighted lotteries and have argued that fairness requires that the person with the strongest claim ought to receive the good directly (Hooker, 2005: 348-9; Lazenby 2014). Hugh Lazenby, for example, has recently argued that "any contribution to fairness from entering claims into a lottery is lexically posterior to fairness in outcome." (2014: 331).

I shall argue for the inverse of Lazenby's lexical ordering: that in distributive cases with indivisible goods and unequal claims, fairness requires that a weighted lottery be selected over an option that directly guarantees fairness

<sup>21</sup> An earlier version of this chapter was presented at the London Moral and Political Philosophy Worksop at UCL. I am grateful to the audience for helpful comments.
<sup>22</sup> Broome claims that whether it is fairest to give the good to the person with the strongest claim or hold a weighted lottery "depends on a complicated judgment" (1990: 98-99), where one has to weigh the fairness contribution of the lottery against the fairness that can be achieved directly in outcome. Broome suggests, however, that in cases where the difference between the strength of claims is slight, one ought to hold a weighted lottery, and when the difference is very great between claims then it is fairer to satisfy the strongest claim directly in outcome. in outcome. Fairness in outcome is the fairness of the pattern of actual satisfaction of claims.<sup>23</sup> In establishing this conclusion, I outline and defend an account of distributive fairness called *Fairness as Proper Recognition of Claims*. This account provides a justification for lotteries in cases with unequal claims and is appropriately sensitive to changes to the number and strength of claims. In Section 2.2 I outline Lazenby's account and in Section 2.3 I consider three objections. I outline my justification of weighted lotteries in Section 2.4 before responding to potential objections to my view in Section 2.5. I offer my conclusion in Section 2.6. Before proceeding, I wish to emphasize that an option being fairer than another does not imply that it is all-things-considered superior. For example, even if it is fairer to destroy the good in question, considerations of utility might outweigh those of fairness. My focus in this chapter is on what fairness requires, rather than what one ought to do all things considered.<sup>24</sup>

#### 2.2 Fair Distribution and Unequal Claims

## 2.2.1 Claims

A "claim" is a kind of moral reason. This is a type of reason that is held by particular individuals, and their satisfaction is a duty owed to the holder (Lazenby, 2014: 332). Claims are in competition with one another when there is a conflict of interest and it is not possible to fully satisfy each claim in

<sup>&</sup>lt;sup>23</sup> The actual satisfaction of claims occurs when an individual actually receives the good that they have a claim to. The term "good" can include probabilistic goods such as lottery tickets (if that is what individuals have a claim to), or non-probabilistic goods.

<sup>&</sup>lt;sup>24</sup> To this extent, I shall bracket a discussion of the debate between Broome (1998) and Frances Kamm (1993) on "saving the greater number". In this debate, a distinction is drawn between what fairness requires and what goodness requires. In a case where one can either save one person's life or five lives, Broome thinks that although fairness requires that each person is given equal chances, goodness overrides these considerations and as such we ought to save the five.

proportion to its strength. I shall assume that claims and the value of the resources that can satisfy them can be measured on a ratio scale.<sup>25</sup> Claims can be satisfied in one of two ways: either through satisfaction in *outcome*, or through what is called *surrogate* satisfaction. Satisfaction in outcome occurs when a claimant actually receives a good in proportion to their claim. Surrogate satisfaction occurs when a claim is entered into a lottery and an individual is given a *chance* in proportion to the strength of the claim. By doing this, "a sort of partial equality in satisfaction can be achieved … each individual can be given a *chance* of getting the good" (Broome, 1990: 97-8). How a claim can be satisfied in outcome is uncontroversial. More needs to be said about how being given a chance of receiving a good can be a form of surrogate satisfaction for a claim, even if one does not receive that good. I return to this question in Section 2.4 when I defend an account of weighted lotteries.

#### 2.2.2 Unequal Claims and the Direct Receipt of an Indivisible Good

Some have argued that in cases with unequal claims to an indivisible good, the individual with the strongest claim should directly receive the good. Brad Hooker, for example, questions whether there would be any unfairness in the strongest claimant receiving the good directly (2005: 349). Hooker claims that "weighted lotteries are not required by fairness – indeed are contrary to fairness" in cases when people have unequal claims to an

<sup>&</sup>lt;sup>25</sup> On a ratio scale, the difference between two values is meaningful, as is their ratio. For example, if well-being were measured on a ratio scale, then it would make sense to say that the difference between a well-being level of 10 and a well-being level of 20 is as large as the difference between a well-being level of 30 and a well-being level of 40. Moreover, it would make sense to say that someone with a well-being level of 20 has twice the well-being of someone with a well-being level of 10. The comparative strength of claims is also measured on a ratio scale. Suppose that Ann's claim is 1.5 times stronger than Bill's. This can be represented with claims of strength 6 and 4 respectively.

indivisible good (Ibid.). Some, such as Broome (1990: 98-99) and Piller (forthcoming: 18) claim that it would sometimes be fairer to give the good directly to the strongest claimant when there is a large difference between the strongest and weakest claims. Lazenby argues that it is fairer to give the good directly to the person with the strongest claim in unequal claim cases, rather than hold a weighted lottery (2014: 331). I shall focus on Lazenby's account since it provides a thorough explanation for why the individual with the strongest claim should receive the good directly.

#### 2.2.3 Lazenby's Argument

Consider the following case:

*Distributive Case*: There are two equally-situated individuals, Ann and Bill. The value of the good to be divided between them is 100. Ann has a claim of strength 6, and Bill has a claim of strength 4.

Lazenby defends the following ranking of options in order of most fair to least fair:<sup>26</sup>

- 1. Give 60 to Ann and 40 to Bill
- 2. Destroy the good
- 3. Give 100 to Ann
- 4. A weighted lottery 60% chance of 100 to Ann, 40% chance of 100 to Bill

<sup>&</sup>lt;sup>26</sup> Lazenby calls this ranking the "alternative conception", as it is formulated as an alternative to Broome's theory of fairness. Since I am not concerned with comparing the relative merits of Broome's and Lazenby's accounts, but rather the plausibility of Lazenby's account as an independent conception of fairness, I shall simply refer to "Lazenby's account".

- 5. Toss a fair coin giving a 50% chance of 100 to each
- 6. Give 100 to Bill

Underlying his ranking is the following view of fairness:

*Fairness as Lexically-prior Outcome Satisfaction (FLOS)*: fairness in outcome is lexically prior to the fairness provided by lotteries (Lazenby, 2014: 335). In cases with unequal claims and an indivisible good up for distribution, one ought to give the good directly to the individual with the strongest claim (2014: 344).

Supplementing this principle is the view that, following Broome, fairness is a matter of the proportionate satisfaction of claims. Further to this, "whether a claim is satisfied proportionately is a matter of how well it is satisfied relative to other claims" (Lazenby, 2014: 334).<sup>27</sup> Lazenby's explanation for this ranking is as follows. Dividing the good in proportion to the strength of each individual's claim is most fair as this ensures that claims are proportionately satisfied in outcome. Destroying the good ensures that Ann and Bill's claims are more proportionately satisfied in outcome, and is fairer than giving the good directly to Ann. Giving the good directly to Ann is fairer than a weighted lottery, since giving the good directly to Ann has a greater expected fairness in outcome than the lottery. Giving each an equal chance is less fair than a weighted lottery, as equal chances provide less of a proportional surrogate satisfaction of each claim. The least fair option is giving the good straight to Bill, since there is no surrogate satisfaction of claims, and claims are least proportionately satisfied in outcome. The focus of my disagreement is with the ordering of options 3 and 4. I shall argue that

<sup>&</sup>lt;sup>27</sup> Although this view can be attributed to John Broome, Lazenby endorses it for his account.

in *Distributive Case,* a weighted lottery is fairer than giving the good directly to Ann. The rest of the paper is a defence of this claim.

#### 2.2.4 Proportionate Satisfaction

Before outlining my three objections to Lazenby's ordering, it is important to examine the nature of "proportional satisfaction", as this plays a central role in determining Lazenby's ranking. There is some ambiguity in this account about what precisely it means for a claim to be proportionally satisfied. There are at least three possible interpretations. First, proportional satisfaction may refer to a ratio where fairness requires that Ann receives 1.5 times more than Bill. Some support for this interpretation can be found when Lazenby states that "if you and I both have claims to a good but my claim is twice as strong as yours, perfect fairness will only be achieved if my claim receives twice as much satisfaction" (2014: 332). However, what I take Lazenby to mean here is that perfect fairness requires that he receives twice as *much*, rather than twice as much *satisfaction*. It may be the case that Ann needs 3 doses of an expensive drug and Bill needs 2 doses of the drug. Giving both what they need is full satisfaction, but Ann gets something that is more valuable than Bill. The ratio view can explain why option 1 is ranked first: in option 1 Ann gets 1.5 times what Bill receives. However, this conception of proportional satisfaction is unable to explain the ranking of options 2-6. In these options, one of Ann or Bill receives 0, and the other receives 100. But if, for example, Ann is entitled to have 1.5 times what Bill receives, and Bill receives 0 and Ann receives 100, then the ratio of Ann's to Bill's receipts is 100 divided by 0, which is infinite. Meanwhile, if Bill receives 100 and Ann receives 0, the ratio of Ann's to Bill's receipts is zero. And infinity is not "closer" to the supposedly optimal ratio of 1.5 than zero

is. This interpretation of proportionate satisfaction therefore cannot generate Lazenby's ranking.

A second possibility is that proportionate satisfaction is measured relative to the "ideal" fair distribution. This view measures the aggregate absolute expected distance from the ideal distribution. In *Distributive Case*, by hypothesis the ideal distribution is that Ann receives 60 and Bill receives 40. The fairness ranking on this view is as follows:

- 1. Give 60 to Ann and 40 to Bill
- 2. Give 100 to Ann
- A weighted lottery 60% chance of 100 to Ann and 40% chance of 100 to Bill
- 4/5/6. Destroy the good
- 4/5/6. Toss a fair coin giving a 50% chance of 100 to each
- 6. Give 100 to Bill.

Dividing the good in proportion to the strength of claims is therefore the fairest distribution as it conforms to the ideal distribution. Giving the good directly to Ann is the second fairest because this leaves Ann with 40 above the ideal distribution and Bill 40 below. A weighted lottery would be the third fairest, as can be demonstrated by calculating its expected value.<sup>28</sup> There is some difficulty on this view, however, ranking the destruction of

<sup>&</sup>lt;sup>28</sup> This can be calculated by multiplying both the probability that Ann receives the good and that Bill receives the good in the weighted lottery. If Ann receives the good in the lottery then there is a deviation of 80 from the ideal distribution (Ann would have 40 more than she needs, and Bill would have 40 less). If Bill receives the good in the lottery, then there is a deviation of 120 from the ideal distribution. Ann has 60 less than she needs, and Bill has 60 more). Multiplying these values by the probabilities of their occurrence in the lottery gives expected value of 96. This is a further deviation from the score that is given to option 2 (where Ann gets the good directly) since this has a score of 80, because in this option Ann would have 40 more than she needs and Bill will have 40 less than what he needs. Therefore, on this view, option 3 would be less fair than option 2.

the good and a fair coin toss, since they both have the same expected deviation from the ideal distribution, by a score of 100.<sup>29</sup> This is an implausible implication of the view because there are good grounds for thinking that a fair coin toss where one person gets the good should be ranked higher than destroying it. Someone will actually receive the good in the former case, and both are given equal chances.<sup>30</sup> Giving the good directly to Bill is the furthest away from the ideal distribution. Although this view of proportionality yields the aforementioned implausible result, the ranking preserves Lazenby's claim that giving the good directly to Ann is fairer than holding a weighted lottery.

A third interpretation is an ordinal view of proportional satisfaction, where it matters that the person with the strongest claim has their claim satisfied more than weaker claimants. The "proportion" referred to in this account is a loose one, where all that matters is that the person with the largest claim receives more than the person with the smaller claim. This view of proportionate satisfaction is most closely aligned to *FLOS*, because this view mandates that the person with the strongest claim ought to receive the good, no matter how much stronger the claim is than others. In what follows, I shall argue that the second and third interpretations are also unsatisfactory because they support the claim that in *Distributive Case* it is fairer that Ann receives the good over the use of a weighted lottery.

<sup>&</sup>lt;sup>29</sup> If the good is destroyed then Ann is 60 away from the ideal distribution and Bill is 40 away, giving a score of 100. The score of an equal-chance lottery can be calculated by the same method used in fn. 6, yielding the score of 100.

<sup>&</sup>lt;sup>30</sup> One might argue that, on the contrary, destroying the good is fairer because there will be no outcome inequality whereby one person receives the good and another does not. However, this would be grist for my claim that destroying the good and giving each an equal chance are not equally fair.

My preferred account of proportionality is a modification of the first account discussed above. If Ann has a claim that is twice as strong as Bill's, then fairness requires that Ann receives twice as much of the good as Bill. Proportionality here refers to a ratio. The ratio account originally ran into difficulties because it applied the same method to both divisible and indivisible goods. I submit that what fairness requires with respect to proportionality depends on whether there is a perfectly divisible or an indivisible good. For cases where there is a perfectly divisible good, it is possible to divide the good in a ratio proportionate to the claims of Ann and Bill. This leads to both individuals receiving full and proportionate satisfaction. In indivisible good cases, by contrast, it is not possible to proportionally divide the good with respect to a ratio. Indeed, as I explained above, the ratio account delivers implausible results when the ratio is applied to indivisible goods. As such, the ratio ought to apply to the *chances* of receiving the good, since, I shall assume, it is possible to divide these in line with the proportionate strength of claims.

The view I am proposing consists of two stages:

- If those with a claim can all receive some of a good *ex post* then the good ought to be distributed in a ratio to strength of claims between the claimants.
- If it is not possible to divide the good in proportion to the ratio between claims, then chances of receiving the good ought to be divided in proportion to the strength of claims.

To illustrate, suppose Ann has a claim of strength 6 and Bill has a claim of strength 4 and there is a perfectly divisible good of 100, then fairness requires that Ann receives 60 and Bill 40. Whereas, if the good of 100 is indivisible then *chances* of receiving the good ought to be divided in proportion to the strength of claims, so that Ann receives a 60% chance and Bill receives a 40% chance of getting the good. This two-stage application of a ratio view of proportionality avoids the problems that beset the original ratio view considered above. Because the proportionate ratio of claims does not apply to the good in indivisible good cases, but instead to chances, there is no incoherence in saying that "Ann gets twice as much as Bill".<sup>31</sup> I develop this approach in Section 2.4 when I defend my account of weighted lotteries. Before doing so, I will now consider three objections to Lazenby's fairness ranking.

## 2.3 Three Objections

#### 2.3.1 Marginally Stronger Claims Objection

My first objection to Lazenby's ordering is that it has implausible implications in cases where one individual comes to have a slightly stronger claim than another. To illustrate, consider a variation of *Distributive Case* – call it the *Equal Claims Case* — where Ann and Bill now have equally strong claims to the good. It seems that Lazenby's ranking would now involve the following:

- 1. Give 50 to Ann and 50 to Bill
- 2. Toss a fair coin, giving each a 50% chance of 100
- 3. Give 100 to Ann, or give 100 to Bill.<sup>32</sup>

<sup>&</sup>lt;sup>31</sup> As it stands, this view does not provide a comprehensive fairness ranking, including the destruction of the good. The view is a defence of weighted lotteries as well as the claim that the fairness of a weighted lottery is ranked above what fairness can achieve by giving the good straight to the individual with the strongest claim.

<sup>&</sup>lt;sup>32</sup> I bracket for now the option of destroying the good, since its place in the ranking has no effect on the claims that I make.

It would be preferable to give half of the good to each over holding a lottery. Giving each half of the good ensures that both claims are satisfied, whereas a lottery will only satisfy one claim. It would then be preferable to toss a coin, giving each a 50% chance of receiving the good rather than simply giving the good to either Ann or Bill. Given that Ann and Bill are equally situated and both have a claim of the same size, there are no good reasons to give the good directly to either of them, rather than flipping a coin. A reason in favour of flipping a fair coin is that it gives each an equal (and proportionate) chance of receiving the good. This is not true when giving the good directly to Ann or to Bill.

Suppose now that due to unexpected natural causes, Ann's claim becomes slightly stronger, at strength 5.1, while Bill's claim remains at 5.<sup>33</sup> Call this the *Almost Equal Case*. On Lazenby's original ranking, from the perspective of fairness a distributor ought to prefer giving the good directly to Ann. Support for this claim can be found when Lazenby says that his view "holds that we should give the kidney directly to [Ann] because her claim is strongest and that this answer does not depend on how much stronger [Ann's] claim is" (2014: 344). This is because there is greater expected fairness in outcome by giving the good to Ann than by holding a weighted lottery. This implication of the ranking is implausible. It is implausible that the *slight* change in the strength of Ann's claim should morally require us to let Ann have a 100% chance over each having a 50% chance, whereas when Ann's claim is just slightly weaker we are required to give them each a 50% chance rather than giving the good outright to Ann.

<sup>&</sup>lt;sup>33</sup> I am assuming that there is a discernible difference in the size of claims and that proportionate chances can be assigned to such claims. James Kirkpatrick and Nick Eastwood (2015) critique this possibility.

A further issue is the following. Whether or not one has the strongest claim depends on the strength of others' claims. In certain circumstances, this too can be cause for complaint. This complaint can be drawn from *Almost Equal Case* when *FLOS* is followed. When Ann's claim becomes slightly stronger than Bills', Bill can say the following:

"I had a claim of strength 5, and I had a 50% chance of receiving the good. Now, because Ann's claim has increased slightly I have a 0% chance of receiving the good, even though my claim remains the same strength."

It seems unfair that an individual's chance of receiving a good should go from 50% to 0% just because of the slight change in the claim of another. A weighted lottery, by contrast, can avoid this complaint by updating the chances that each individual has of receiving the good to reflect the comparative strength of claims. On this view, Bill will have a slightly diminished chance of the good and Ann will get a slightly larger chance of receiving the good. This, I hold, is a fairer distribution of chances.

Bill's complaint highlights the dependency of chances in distributive cases with indivisible goods. By dependent chances, I am referring to the fact that one person's chances for a good affects the chances that another person has. The interests of Ann and Bill in *Distributive Case* are in conflict with one another, and as such giving the good directly to Ann or holding a weighted lottery are zero-sum.<sup>34</sup> Either Ann receives a 100% chance and Bill gets a 0% chance, or Ann 60% and Bill 40% and so on. Because their interests are

<sup>&</sup>lt;sup>34</sup> This is not true for cases where chances are independent, i.e. where the chance given to one person does not depend on chances that are given to another. A simple example would be a fair coin flip to see if you win a prize or not. The chances of this fair coin flip are not affected by the chances that another person has for a separate prize.

connected in this way, it is fairer to hold a weighted lottery to guarantee a more equitable distribution of chances consistent with the claims of individuals.

#### 2.3.2 Multi-Person Objection

The second objection is that Lazenby's ordering has difficulties in plausibly accounting for cases of more than two individuals involving unequal claims. By way of illustration, consider the following case:

> *Multi-Person Case*: There are five individuals, Ann, Bill, Catherine, David and Eleanor. Ann has a claim of strength 2.2, and the other four individuals have a claim of strength 1.95. There is a single, indivisible good of value 100.

*FLOS* would rank giving the good to Ann lexically higher than a weighted lottery, because this would give a higher expected fairness in outcome over holding a weighted lottery. On the contrary, it seems to me that it is preferable to hold a weighted lottery in which each of the five individuals' chances of winning are proportional to the strength of their claims, so that Ann has a 22% chance of receiving the good, and the other four each have a 19.5% chance of receiving the good. The latter option is a more egalitarian distribution of chances, and more consistent with recognition of each person's claims, than giving it outright to Ann. While it may be the case that a weighted lottery has a lower expected fairness of *outcome*, this is different to having lower expected fairness all things considered.

There is, moreover, stronger reason to hold a weighted lottery the more individuals' claims are at stake. It is worse for a larger number of individuals to receive no chance at receiving a good to which they have a

claim, than it is for a smaller number of individuals to be so deprived.<sup>35</sup> I submit that the explanation is that it matters for fairness not only whether claims are *satisfied* in proportion to their strength, but also whether claims are given what I call "proper recognition". For a claim to be given proper recognition, a claimant's chances of receiving an indivisible good ought to be proportionate to the strength of their claim. One may press the opposite view, and argue that the greater the number of individuals with weaker claims, the more important it is that the strongest claimant receives the good directly. For example, suppose that there is a case where there is one individual with a claim of strength 4, and one hundred thousand individuals each with a claim of strength 0.00006. In this case the expected outcome is for one of the individuals with a very weak claim to win the lottery, as the probability that a very weak claimant will get the good is .6. As such, one might argue that fairness requires that the good ought to go directly to the strongest claimant. But I take the thought here to be that there are other reasons than fairness, such as a reason to maximise expected goodness, which may sometimes overrule considerations of fairness.<sup>36</sup> This influences the thought that we ought to give the good directly to the individual with the strongest claim. In some cases, when one is deliberating about what to do all-things considered, it may be appropriate to give the good directly to the strongest claimant, but this is separate from the claim that that would be what *fairness* requires. Section 2.4 considers in more detail

<sup>&</sup>lt;sup>35</sup> To illustrate, suppose that in a modification to *Multi-Person Case* there are now nine individuals, where Ann has a claim of strength 2 and the other eight have a claim of strength 1. If the good is given directly to Ann then there are eight unsatisfied claims that did not receive a chance of being satisfied. Whereas, I take it to be a worse state of affairs when there are ninety-eight individuals with unsatisfied claims that did not receive a chance. This is because there are more unsatisfied claims in this case than the former. <sup>36</sup> See, for example, Broome (1990: 90; 1998: 956-7) and Kamm (1993: 107-8).

the claim that in such cases of unequal claims fairness requires a weighted lottery.

#### 2.3.3 Relational Unfairness Objection

Leading on from the preceding two objections is a third: *FLOS* is insensitive to the relational complaint that some may have stronger claims than others due to no fault or choice of their own. In *Distributive Case, FLOS* would recommend giving the good directly to Ann over holding a weighted lottery. It seems unfair that just because one happens to have a weaker claim than another that one has no chance at all of receiving the good one has a claim to. It is a mere contingent fact that others happen to exist with stronger or weaker claims. One may respond and claim that this isn't a cause for complaint at all. So long as there was some degree of *ex ante* randomisation with respect to who had a claim of a particular size, then it could be said that Bill *did* have a chance of having the strongest claim. To develop this response further, we can draw on the following from Kornhauser and Sager:

*Principle of Convolution*: "However badly biased the [chance device] may be, the equiprobability of the initial fair draw of the [lots] produces equiprobability overall in the candidates' chances of getting the prize" (1988: 486).

To illustrate, suppose that there is a loaded die that always comes up 6. There is a prize up for grabs, and six people. If the allocation of numbers to people itself gives each an equal chance of receiving a particular number, then the use of a biased die is not unfair, since each person had an equal chance of being allocated to 6. Following this, we might think that Ann and Bill had an equal chance of ending up as the person with the strongest claim.

As such, it would not be unfair to simply give the good directly to Ann in *Distributive Case*. However, it seems unlikely that one can guarantee that prior claims have been randomised with the degree of precision required to make such a process fair.<sup>37</sup> Although it may be true that one does not know what the size of one's claim will be in a particular case, this does not mean that one's chances of having a particular-sized claim are equal. Given that we often do not know the likelihood of particular individuals having particular sized claims, it is reasonable to assume that individuals did not have an equal chance at ending up with the strongest claim.

## 2.4 An Account of Weighted Lotteries

## 2.4.1 The Argument

This section outlines a distinct view of fairness for conflict of interest cases with unequal claims that provides a justification for weighted lotteries, and provides a more plausible rationale for how one should act in *Almost Equal Case* and *Multi-Person Case*. My argument for the use of weighted lotteries in distributive cases is the following:

- 1. Fairness requires that affected claims be given proper recognition.
- In cases with unequal claims to indivisible goods, giving claims proper recognition requires the assignment of chances in proportion to the strength of claims.
- 3. A weighted lottery is necessary and sufficient for giving chances proportionate to the strength of claims.

<sup>&</sup>lt;sup>37</sup> Ben Saunders provides a critique of this claim: "there is no sense in which [Bill] – as a concrete individual – has any chance of being in the [most advantaged position]. He had a chance only if we appeal to something like the "original position", in which his identity was unknown and he could have been one of [Ann or Bill]" (2009: 282).

4. Therefore, fairness requires a weighted lottery in cases with indivisible goods and unequal claims.

The important premises in my argument are 1 and 2. Regarding premise 3, there is no other method of assigning proportionate chances to individuals that is not some form of weighted lottery, where chances are assigned in light of the differing strengths of claims. I shall now turn to a defence of the important premises.

#### 2.4.2 Premise One

The first premise states that fairness requires that affected claims be given proper recognition. There are at least two important ways that a decisionmaker can take into account the claims of individuals. First, the claims of individuals can be *weighed*. For this "the right thing to do, and the right way to distribute a good is determined by the balance of reasons, whether claims or not" (Broome, 1990: 93). Broome argues that claims cannot be satisfied simply by weighing them against other reasons or other claims. This is because claims are held by particular, separate individuals, and so one does not satisfy these claims merely by weighing them against other claims and reasons.

Second, claims can be given *proper recognition*. When giving a claim proper recognition, it is important to do more than merely take notice of the claim and compare its strength to the claims of others. It matters that a claim is given appropriate consideration. The starting point for this view is the claim that the pattern of claims and their (expected) satisfaction do not exhaust the considerations that are relevant for fairness. Claims and their satisfaction are important, but so is the manner in which claims are treated. It will be helpful

here to consider Stephen Darwall's account of "recognition respect". On this view, "to say that persons as such are entitled to respect is to say that they are entitled to have other persons take seriously and weigh appropriately the fact that they are persons in deliberating about what to do" (1977: 38). I propose that something similar is true with respect to claims. One ought to take seriously the fact that people have claims to a good.<sup>38</sup> A variety of different objects can be the object of recognition respect (Dillon, 2014). The following case from Kamm (1993: 114-19)<sup>39</sup> helps to further draw out what it means for claims to be given proper recognition. Suppose there are two individuals, A and B, who are stranded on separate islands. We can only save one from drowning at little cost to ourselves. We ought to flip a fair coin in order to decide who gets saved, since there are no good reasons for directly favouring one person over the other. Suppose now that C is added alongside B. If we continue to flip a coin, then we would be ignoring C's presence, since we will be acting in exactly the same way as if C was not there.40

To give claims proper recognition is to appropriately take them into account. Part of what this means is that a change in either the number or intensity of claims should make a difference to what a decision-maker ought to do. It should also make a *continuous* difference, so a slight change in the strength

<sup>&</sup>lt;sup>38</sup> Support for this idea is provided by Kamm's account of fairness, whereby fairness requires us to give proper respect to individuals as individuals. For Kamm, fairness requires that we maintain our ties to individuals as individuals (Kamm, 1993: 88; Broome, 1998: 957). In the case at the end of this paragraph, Kamm argues that "we retain our ties to each of A and B when we give each of them as much of a chance as is consistent with retaining our tie to the other", and when C is added alongside B, "[i]f we then save the greater number [(B+C)], this means severing our tie to A before giving him an equal chance" (Kamm, 1993: 88). When it is possible to either save A or B, Kamm argues that tossing a fair coin "allows the [decision-maker] to remain attached to each as an individual" (Ibid).

<sup>&</sup>lt;sup>39</sup> T. M. Scanlon (1998: 232-3) also discusses this case.

<sup>&</sup>lt;sup>40</sup> This scenario is discussed in more detail by Frances Kamm (1993), Scanlon (1998) and Otsuka (2006: 114).

of a claim should warrant a slight change to what one ought to do. For instance, in *Almost Equal Case* when Ann's claim increases from 5 to 5.1 whilst Bill's stays at 5, *FLOS* states that Ann should now receive the good rather than hold a weighted lottery. From a small change in the strength of a claim there is a large change in the chances of the individuals. Instead, a weighted lottery would be sensitive to the small changes in comparative claim size.

One might argue that one can give proper recognition to claims by simply recognising the fact that individuals have claims of different strengths, and then accepting the fact that a decision-maker's action should be based on the fact that an individual's weaker claim is trumped by another's stronger claim. However, this view fails to respect the separateness of persons as it resolves interpersonal conflicts in a way that fails to respect the moral relevance of interpersonal boundaries. The view nominally takes into account the claims of different individuals, but only so far as the claims are sorted in order of strength and the person who happens to have the strongest is given the good directly. Parallels can be drawn with the argument against utilitarianism on the same grounds:

> "Utilitarianism [...] confronts the [...] objection when it recommends giving a resource to a well-off agent over a badly-off agent on the sole ground that the well-off agent will derive slightly more benefit from the resources than the badly-off agent" (Hyams, 2015: 219).

For FLOS the "sole ground" for giving a resource to one person than another is the fact that they happen to have a stronger claim than another. By doing so, it ignores intuitions about the moral relevance of interpersonal

boundaries. The view I am proposing respects the separateness of persons because it takes seriously the distinction between individuals. By giving claims proper recognition, one respects the distinctness of each particular individual's claim. Claims aren't simply overridden on the account I am proposing.

## 2.4.3 Premise Two

Further support for the claim that lotteries give proper recognition to claims can be found in justifications of the fairness of lotteries with equal chances. David Wasserman, for instance, argues that "lotteries are fair if they *respect* the claimants' equal entitlements to the scarce good" (1996: 48, italics added). Although Wasserman discusses only equal chance lotteries, there is no good reason for thinking that weighted lotteries cannot also respect people's unequal claims. By holding a weighted lottery in *Distributive Case* the distributor respects the claims of Ann and Bill by apportioning the chances to reflect their comparative strengths. Peter Stone (2007) builds his account of the justness of lotteries on cases where individuals have equal claims to a good, and such equal-chance lotteries are "uniquely appropriate for the just indeterminacies in allocation" (2007: 280). I wish to challenge the idea that equal-chance lotteries are *uniquely* appropriate for indeterminacies in allocation, rather than lotteries *per se.*<sup>41</sup>

<sup>&</sup>lt;sup>41</sup> Peter Stone appears to recommend a similar proposal to Lazenby in deciding what one is required to do in cases with unequal claims. Stone argues that "an agent charged with allocating [a] good in a case [with unequal claims] should rank order the claimants according to the strength of their claims. [...] Stronger claims should gain access to the good before weaker claims do" (2007: 278). On this view, lotteries are only appropriate for cases where claims are equal.

In cases where the claims of individuals are equal and the good is indivisible it is generally thought that one ought to hold a lottery at equal odds. Part of the justification for this I take to be the following principle:

> *Principle of Proportionality*: In conflict of interest cases with an indivisible good and more than one individual, chances of receiving the good for each individual ought to be proportional to the strength of their respective claims.

In other words, if individuals have equal claims to a good then it would be unfair for one individual to get higher chances of receiving the good than another. The *Principle of Proportionality*, I submit, applies not only to cases with equal claims, where the fairness of lotteries has typically been endorsed, but also to cases of unequal claims. For cases of unequal claims, chances ought to be proportionate to the claims of individuals. One may resist this move and argue that the *Principle of Proportionality* should only apply to cases where claims are equal. It may be argued that a lottery at equal odds functions merely as a "tie-breaker" between equally-situated individuals and that for cases with unequal claims one ought to give the good to the individual with the strongest claim, or use some other rubric.<sup>42</sup> In response, there is no good reason for the *Principle of Proportionality* to only

<sup>&</sup>lt;sup>42</sup> In support of this claim, one may argue that a lottery merely functions as a tie-breaker in cases of equal claims due to the impartiality that the lottery offers. Each individual will receive equal chances of receiving the good, and as such the distributor fails to be partial towards one of the claimants. However, the defender of a lottery in such cases must be after something stronger than a merely impartial procedure. This is because it is possible to use certain impartial procedures for breaking a tie that do not give individuals an equiprobable chance of receiving the good (Sher, 1980: 207-12). For example, giving the good to the person with the largest pancreas (Kornhauser and Sager, 1988: 490). If one then thinks that a lottery is preferable to such impartial procedures, then it must be because of the lotteries allocation of probabilities to the claimants. If the allocation of probabilities bears the value of the lottery for such equal-claims cases, then it is difficult to see why this value only occurs in cases with equal claims.
apply to cases of equal claims. In the same way that equal chances reflect how individuals are equally situated with respect to the size of each other's claims, unequal chances ought to reflect the claims' unequal strength. A reason for thinking that proportionate chances ought to be given to individuals with unequal claims is that giving chances to an individual tracks the value of the claim that an individual has, thereby signifying the claim's importance.

The following principle underlies the view that changes in the strength of claims ought to make a continuous difference to what a decision-maker ought to do:

> *Fairness as Proper Recognition of Claims*: Fairness requires that when individuals have a claim to a good, and it is not possible to proportionately satisfy all claims in outcome, these individuals be given the highest chances that correspond to the ratio of respective claims.

Under this principle it would be unfair to give the good directly to the individual with the strongest claim. Because there are competing claims, and because fairness requires giving proper weight to the claims of the individuals who have their interests at stake, a distributor ought to hold a weighted lottery. It doesn't just matter that individuals have proportionate chances, but that they are the highest proportionate chances that are possible.<sup>43</sup> Giving claims proper recognition requires the application of the

<sup>&</sup>lt;sup>43</sup> For example, suppose that Ann has a claim that is twice as strong as Bill's claim. it would be would be proportionate to give Ann a 2% chance of getting a good and Bill a 1% chance of receiving the good, with a 97% chance of the good being wasted. Whereas, given the choice between the previous option and the following, it ought to be preferred that an option is selected where Ann gets a 67% chance and Bill gets a 33% chance (rounded). If Bill

*Principle of Proportionality* to the ratio of claims. If one does not give proportionate chances in such cases, or if one simply gives the good directly to one of the claimants, then proper recognition is not being given to the claims. In a case with a single individual, and no competing interests of others, this individual ought to directly receive the good as this maximises their chances of receiving the good, and does not conflict with any chances given to others.

### 2.4.3.1 Broome and Surrogate Satisfaction

I shall now distinguish my view from that of Broome. According to him, when it is not possible to satisfy claims in proportion to their relative strengths, a claim may receive "surrogate satisfaction" by being entered into a lottery. How exactly a claim receives surrogate satisfaction is not immediately clear (Lazenby, 2014: 337-8). One interpretation is that "lotteries give something of value to the claimants even if they do not receive the scarce good" (Wasserman, 1996: 47). One might think that this thing of value is then "deducted" from the individual's claim, making it weaker than it was before. To illustrate what this means, suppose that Ann and Bob initially have equal claims to a good and Ann receives a 50% chance to receive this good and Bill does not – for example, a coin is tossed and Ann receives the good if it lands "heads", while the good remains temporarily undistributed if it lands "tails". Suppose further that her chance comes to nought. Then, so the "deduction" account goes, her claim is thereby weaker than Bill's, even though she has ended up with no actual good (Lazenby, 2014: 337). This is a problematic implication of the view that chances provide surrogate

gets a higher chance than 33% then this will come at the expense of Ann's chances of receiving the good.

satisfaction on the same metric as the claim itself, because although Ann had a chance of receiving the good, she does not have any of the good. There is no actual satisfaction of the claim. One might instead posit that there are grounds for thinking that Bill should receive the next chance, separate from the thought that Ann has part of her claim satisfied by having received a chance.

Another strategy is to claim that surrogate satisfaction constitutes an "improvement in fairness without satisfaction" (Ibid). On this view claims are not diminished by being entered into lotteries. Lazenby, however, argues that this view gets the wrong result in an amended version of the above case. Suppose that Ann receives a chance of a good that fails to materialise. Bill receives nothing. There is now a lottery ticket that can be distributed between them. On the "improvement in fairness view" there should be a fair lottery to decide who gets the ticket, since there is no change in Ann's claim after having the previous chances. Lazenby argues that the ticket should go directly to Bill since he has not had a chance previously, whereas Ann has. I maintain the right answer depends both on whether there is a divisible or indivisible good, and whether claims are in competition when a good is distributed. When there is an indivisible good, such as the lottery ticket, there ought to be a lottery to decide who gets the ticket, since in the case of an indivisible good there is a conflict of interests where not all claims can be satisfied simultaneously. This is not the case with divisible goods. The fact that Ann had an unfulfilled chance of receiving a good in the past (when only *she* stood to gain and nobody else did, and giving her this chance did not, at that time, diminish anyone else's chances, because there was as yet no conflict of interests) ought not to diminish her claim when it is in

competition with another. To this extent, I submit that it makes a moral difference whether one received a (chance of a) good when there were no competing claims on the good, and when one receives a (chance of a) good when there *are* competing claims with others.

I assert that it is a mistake to see an individual as *receiving* something when their claim is entered into a lottery.<sup>44</sup> It is not helpful to see a claim being satisfied in outcome when an individual receives a chance. I propose a different approach. The "improvement in fairness" view can be developed into an account of what it means to give proper recognition to a person's claim. This is a more deontological approach which focuses not on whether a lottery can partly or fully satisfy a claim (either directly in outcome, or by depreciating the strength of the claim), but on whether a claim was appropriately taken into account in an agent's deliberation.<sup>45</sup> *Fairness as Proper Recognition of Claims* is an instance of this deontological approach.

## 2.4.4 Fairness as Proper Recognition versus FLOS

To bring out the differences between the two accounts, consider the following case:<sup>46</sup>

Additional People Case: At  $t_1$  Ann has a claim of strength 4 for an indivisible good. At  $t_2$  Bill materialises alongside

<sup>&</sup>lt;sup>44</sup> The language of an individual "receiving" something when their claim is entered into a lottery refers to what Wasserman calls the "distributive view" of lottery fairness. On this view, "lotteries treat equally-entitled claimants fairly by giving each claimant some intangible good in equal measure" (1996: 30). Although individuals trivially receive a *chance* of getting the good, I submit that they do not receive anything that satisfies the claim in outcome.

<sup>&</sup>lt;sup>45</sup> Other accounts also reject the claim that lotteries provide some sort of surrogate satisfaction for claims. See, for example, Peter Stone (2007) and David Wasserman (1996).
<sup>46</sup> There is an important disanalogy between the "rescuing the greater number" case and the distributive cases under consideration in this paper. The good of survival that is at stake in the rescuing cases can be available to more than one person, whereas in the distributive cases under discussion the good can only go to one single individual.

Ann with a competing claim of strength 3.5 for the indivisible good. A few moments later at  $t_3$  Clara materialises alongside Ann and Bill with a competing claim of strength 2.5, leading to three individuals at  $t_3$  with unequal claims.

What is the fairest action to perform at each stage? According to *FLOS*, fairness requires that the option of giving the good directly to Ann ought to be ranked first in each stage. I disagree. *FLOS* ignores morally significant differences among these stages. The account of weighted lotteries that I provide gives a more plausible recommendation for this case. At *t1* Ann should receive the good directly. At *t2* the fairest action would be to hold a weighted lottery with chances in proportion to the strength of the claims of Ann and Bill. At *t3* the chances in the weighted lottery ought to be modified in order to accommodate the extra claim of Clara.<sup>47</sup> Although the actions at *t2* and *t3* do not have a greater expected fairness in *outcome*, since they do not guarantee that the individual with the strongest claim gets the good, they are nevertheless fairer all things considered. This is because, on my view, the claim of each additional person is taken into account and suitably modifies what one ought to do.

One may argue that it is unfair that Ann fails to be guaranteed the good at *t*2 and *t*3 since at these times she has the strongest claim. Ann's chances of receiving the good in *t*1 were 100%, and they drop to 40% in *t*3. In *t*1 there was no conflict of interest. Only Ann's interests were at stake and so it was

<sup>&</sup>lt;sup>47</sup> At *t*2 the chances of Ann and Bill would be 53.5% and 46.5% respectively, to reflect the fact that Ann's claim is 1.15 times bigger than Bills. In *t*3 the chances of Ann, Bill and Clara will be 40%, 35% and 25% respectively, to reflect the fact that Ann's claim is 1.15 times bigger than Bill's and 1.6 times bigger than Clara's.

appropriate to give her the good, whatever the strength of her claim. In *t*2 and *t*3 there is a conflict of interest. Only one person can receive the good, and it would be inappropriate to give the good directly to Ann. Other people also have claims to the good, and this is something that ought to be taken into account in conflict of interest cases. Fairness isn't simply the satisfaction of claims proportionately in outcome, especially in cases when it is not possible for all claims to receive satisfaction.

#### 2.4.5 Summary

My proposal can be summarised as follows. In a distributive case with claims to a divisible good, the good ought to be distributed in proportion to the strength of claims. In a distributive case with an indivisible good, with multiple people and unequally-sized claims, it is fairer to provide a weighted lottery over giving the good directly to the individual with the strongest claim. Given my account of weighted lotteries, option 3 and 4 ought to be reversed in Lazenby's ranking. The account I have given is sensitive to changes in the size of claims and the number of people, and also gives proper recognition to the claims of each affected person. This account provides a more plausible rationale for identifying the fairest option in *Almost Equal Case, Multi-Person Case* and *Additional People Case*.

## 2.5 Objections and Replies

## 2.5.1 Does Fairness Require a Weighted Lottery?

One may question whether fairness *requires* proper recognition and not, like Broome argues, that "claims should be satisfied in proportion to their strength" (1990: 95). On my account, fairness requires that claims are given proper recognition. What follows from this is that in cases with unequal

claims to indivisible goods, fairness requires a weighted lottery. This is not the case on Broome's view. For Broome, whether it is fairest to hold a weighted lottery or give the good directly to the strongest claimant "depends on a complicated judgment" (Ibid: 98-99) between the fairness that can be achieved directly in outcome versus a lottery that provides surrogate satisfaction of claims. Broome endorses weighted lotteries in cases where the difference between the strength of claims is slight (Broome, 1990: 99; Piller, forthcoming: 19). If fairness requires that claims are satisfied in proportion to their strength, then for cases with unequal claims fairness does not always require a weighted lottery, because this aim can sometimes be achieved by giving the good directly to the person with the strongest claim.

In response, *Fairness as Proper Recognition of Claims* guarantees proportionality through a commitment to proper recognition of claims. Claims should be satisfied in proportion to their strength in cases with divisible goods, but proportionality ought to apply to chances in indivisible good cases. A key difference between the two views is that Broome's view allows for claims to sometimes be overridden when the difference between claim strength is great. This is because there are two values of fairness at play: the fairness that can be achieved directly in outcome, and the fairness contribution of chances. Proportionate satisfaction of claims applies to both, and the values are to be balanced against each other. *Fairness as Proper Recognition of Claims*, by contrast, gives a determinate answer to what fairness requires, and plausibly incorporates a consideration of proportionality, even though it is not the starting point of the account.

#### 2.5.2 "But the Claim is Stronger!"

One might argue that as a matter of fairness, stronger claims just ought to be satisfied over weaker claims. For example, Stone argues that weighted lotteries do

"violence to the whole idea of impartially considering the strength of claims. Any grounds that an agent could have for assigning a higher weight to the chance that x will get the good than the chance y will get the good should count as grounds for simply giving x, and not y, the good outright" (2007: n287).

I agree that *x* having a stronger claim than *y* counts as a *ground* for simply giving the good to *x*, but it does not follow that the strength of claims *alone* determines what action is most *fair* in a distributive case with an indivisible good. When other individuals have claims to an indivisible good, fairness demands that our distributive actions are responsive to their claims too. This would not be achieved if one simply gave the good directly to the individual with the stronger claim. A related worry is that one will override the claims of those with weaker claims. This is what happens if the good is given directly to the person with the stronger claim, even though additional people are added who also have a claim, such as in *Multi-Person Case* and *Additional People Case*. A weighted lottery avoids this worry.

## 2.5.3 Very Unequal Claims

One general objection to the use of weighted lotteries is that it runs counter to our intuitions in cases where there are very unequal claims. To illustrate, consider the following case: *Very Unequal Claims*: There are two equally-situated individuals, Ann and Bill. Ann has a claim of strength 9, and Bill has a claim of strength 1.

In this case Ann is entitled to the good nine times more than Bill. Her claim is far stronger than Bill's. Why shouldn't the good just go straight to Ann? This thought is echoed by Brad Hooker, who questions the plausibility of holding a weighted lottery when there are very unequal claims:

> "Suppose your claims to some indivisible good are very much weightier than mine. Is there any unfairness in your getting the indivisible good rather than my getting it?" (Hooker, 2005: 349).

Giving the good directly to Ann would have greater expected outcome fairness than holding a weighted lottery. And intuitively, one might think that because Ann's claim is so strong and Bills is so weak by comparison, we should give the good straight to Ann. This sort of case supports Lazenby's claim that what fairness can achieve directly in outcome is superior to what fairness can be achieved by holding a weighted lottery. However, I submit that fairness requires a weighted lottery in this case. By holding a weighted lottery there is still a large expected fairness in outcome. There is a 90% chance that the person with the larger claim gets the good. But crucially, the weighted lottery also gives proper recognition to Bill's claim by giving him a chance of receiving the good that is proportional to the size of his claim. This plurality of factors ought to be taken into account in a plausible account of distributive fairness.

One may respond and argue that there is less fairness-based reason to perform a weighted lottery the more extreme the difference between claims. Christian Piller (forthcoming: 13-19) considers this possibility when defending Broome's theory of fairness. Piller states that there are two "layers" of fairness. The first is with respect to outcomes, where it is more unfair that the person with the weaker claim receives the good in a weighted lottery compared to if the person with the stronger claim receives the good. There is outcome unfairness even if the person with the stronger claim receives the good because the person with the weakest claim will receive nothing, whilst they receive something. Piller states that "the bigger the difference between the stronger and the weaker claim, the bigger the [outcome] unfairness will be if the weaker claim wins" (Ibid: 17). The second aspect of fairness applies to how claims are treated in the procedure of distribution. A weighted lottery can be said to be fair if it gives chances proportional to the size of claims as this treats claimants fairly. These two aspects of fairness have to be traded off against one another. Proportionate lotteries are fair with respect to the *ex ante* treatment of claims, but there is a chance that there will be outcome unfairness *ex post* where the person with the weakest claim receives the good over the person with the strongest claim.

Piller states that:

"Proportional lotteries are fair in a way. They are unfair in another way. This further unfairness increases as the difference in the claim strength increases. Thus, weighted lotteries are overall fair only when people have claims of similar strength" (Ibid: 19).

There are two problems with this approach. Firstly, Piller's statement of the fairness of weighted lotteries struggles with cases with more than two people. Recall Multi-Person Case where there was one person with a claim of strength 2.2 and four with a claim of strength 1.95. The claims are of similar strength, but if a weighted lottery is run then there is a 78% chance that a person with the weakest claim gets the good over the person with the strongest claim. It is more likely that someone with a weaker claim will get the good, but the difference in claim strength is only slight. Because failing to satisfy the strongest claim leads to greater unfairness than failing to satisfy the weakest claim, then there should be reason to think, on Piller's view, that one should give the good directly to the person with the strongest claim. This is because there is a 78% chance that there will be greater outcome unfairness than if the person with a claim of strength 2 received the good directly. However, there is a reason *against* giving the good to the person with the strongest claim because the claims are of similar strength. The preceding objection arises because of the two "layers" of fairness which can be traded off with one another. Because *Fairness as Proper Recognition of Claims* does not hold this view, it is not vulnerable to the objection, since the account provides a determinate answer to what one ought to do in indivisible goods cases.

Secondly, the view that weighted lotteries are fair overall only when people have claims of similar strength can be challenged by considering the following example:

*Unequal Claim Case*: Ann and Bill have unequal claims to a good with the value of 10.

Table 2.1: Expected Outcome Unfairness of Holding a WeightedLottery versus Giving the Good Directly to the Person with theStrongest Claim.

| Ann's    | Bill's   | Expected        | Expected        | Expected              |
|----------|----------|-----------------|-----------------|-----------------------|
| claim    | claim    | outcome         | outcome         | outcome               |
| strength | strength | unfairness of a | unfairness of   | unfairness            |
|          |          | weighted        | giving good     | "cost" of             |
|          |          | lottery         | directly to Ann | lottery <sup>48</sup> |
|          |          |                 |                 |                       |
| 9.5      | 0.5      | 1.9             | 1               | .9                    |
| 9        | 1        | 3.6             | 2               | 1.6                   |
| 8.5      | 1.5      | 5.1             | 3               | 2.1                   |
| 8        | 2        | 6.4             | 4               | 2.4                   |
| 7.5      | 2.5      | 7.5             | 5               | 2.5                   |
| 7        | 3        | 8.4             | 6               | 2.4                   |
| 6.5      | 3.5      | 9.1             | 7               | 2.1                   |
| 6        | 4        | 9.6             | 8               | 1.6                   |
| 5.5      | 4.5      | 9.9             | 9               | .9                    |
| 5        | 5        | 10              | 10              | 0                     |

Table 2.1 compares the expected outcome unfairness of holding a weighted and the expected outcome unfairness of giving the good directly to Ann. This is calculated for ten different scenarios where the claim strengths are of

<sup>&</sup>lt;sup>48</sup> The higher the number, the larger the expected unfairness cost of holding a weighted lottery over giving the good directly to Ann.

differing magnitudes.<sup>49</sup> It is then possible to calculate the "cost" of using a weighted lottery by taking the expected outcome unfairness of holding a weighted lottery and deducting the expected outcome fairness of giving the good directly to Ann. This is the loss in expected fairness when one holds a weighted lottery over giving the good directly to Ann. The first observation to note is that there is a lower expected outcome unfairness when one gives the good directly to Ann in all cases apart from when Ann's and Bill's claims are equally strong. In this scenario the expected unfairness is the same for giving the good directly to Ann as it is for holding a lottery. A reason in favour of holding a lottery in this case is that the expected unfairness cost of the lottery is zero. The alternative of giving the good directly to Ann, although having the same expected outcome unfairness as holding the lottery, would be fairer for reasons *additional* to the expected outcome unfairness. For example, it is unfair that Bill has no chance of receiving the good given that his claim is exactly the same strength as Ann's.

There is, of course, greater expected outcome fairness in giving the good directly to the person with the strongest claim, but I have argued in this chapter that fairness demands that claims be given proper recognition, and as such the fairest alternative is not always the one that will have the

<sup>&</sup>lt;sup>49</sup> The expected outcome unfairness of holding a weighted lottery is calculated as follows. First, I introduce a distance measure between what both Ann and Bill receive under a particular lottery and what an ideal distribution of the 10 would be. The distance measure that I adopt is the "sum of absolute differences". I use this particular measure because it captures Piller's claim that the larger the distance between claims, the larger the unfairness if the weaker claim wins. If Bill has a weak claim, and does not receive any of the good, then this is reflected in a low score. Whereas, if Ann has a much larger claim than Bill (and Bill then receives the good), a higher score is given to reflect the greater unfairness of Bill receiving the good. An illustration of how this is calculated with the scenario where Ann has a claim of strength 7 and Bill a claim of strength 3 is as follows. Outcome 1: with 70% chance, Ann gets 10 and Bill gets nothing. The sum of absolute differences would be (10-7) + (0-3) = 6. In outcome 2, with 30% chance Ann gets 0 and Bill gets 10. The calculation of the absolute difference is as follows: (7-0) + (3+10) = 14. The expected outcome unfairness is then the following: 0.7\*6 + 0.3\*14 = 8.4.

greatest expected fairness in outcome. This is because, following Broome, I take fairness to be about how one *treats* claims to a good (Piller, 2017: 214). A weighted lottery treats claims in a proportionate way, unlike the rule of giving the good to the strongest claim. There may be outcome inequality as a result of the weighted lottery, whereby the person with the weaker claim gets a good and the stronger claim does not, but this does not always unfair, where unfairness is about how claims to a good were *treated*.

A second observation is that the marginal outcome fairness cost of holding a weighted lottery decreases as the comparative claim strength moves away in either direction from the case where Ann has a claim of strength 7.5 and Bill a claim of 2.5. As such, the expected outcome unfairness of holding a weighted lottery decreases as the difference between the strength of claims gets very large and very small. Therefore, holding a weighted lottery with very unequal claims may have lower or equal expected outcome unfairness cost to a lottery that is held with moderately unequal claims. If one's motivation for holding a weighted lottery when claims are close in strength is that there is a low fairness cost to the lottery, then one should not be resistant to the use of a weighted lottery with very unequal claims if the expected fairness cost is the same.

Broome and Piller may object that Table 2.1 fails to reflect the *severity* of the unfairness if Bill receives the good when his claim is very small. Table 2.2, below, demonstrates the expected unfairness of holding a weighted lottery where the severity of the unfairness if Bill gets the good is weighted by the distance between the claims of Ann and Bill. This is to reflect Piller's claim that it is more unfair if a person receives the good when they have a much smaller claim, than it is when they have a nearly equal claim to the good.

Table 2.2: Expected Outcome Unfairness of Lotteries Conditionalon the Person with the Weaker Claim Receiving the Good.

| Claim Strengths |      | Distance | Expected          |
|-----------------|------|----------|-------------------|
| Ann             | Bill | between  | unfairness of the |
|                 |      | claims   | lottery           |
| 9.5             | 0.5  | 9        | 14                |
| 9               | 1    | 8        | 17                |
| 8.5             | 1.5  | 7        | 19                |
| 8               | 2    | 6        | 20                |
| 7.5             | 2.5  | 5        | 20                |
| 7               | 3    | 4        | 19                |
| 6.5             | 3.5  | 3        | 17                |
| 6               | 4    | 2        | 14                |
| 5.5             | 4.5  | 1        | 9                 |
| 5               | 5    | 0        | 10                |

To reflect the greater unfairness of the weaker claimant receiving the indivisible good in the lottery, the value of the good is multiplied by the distance between the claims of the stronger and the weaker claimant. This provides a measure of the severity of the outcome unfairness *if* the weaker claimant receives the good in the lottery. The larger the number, the greater the outcome unfairness. The expected value of the lotteries is calculated by multiplying the probability of the strongest claimant receiving the good by the value of the good and then adding it to the probability of the weaker

claimant multiplied by the weighted value of the good.<sup>50</sup> Even when the severity of the unfairness of Bill receiving the good when his claim is very weak is included in the calculation of the expected unfairness of the lotteries, there is as much expected unfairness as cases where claims are of nearly equal strength. Table 2.1 and 2.2 both illustrate that there is greater expected unfairness when a weighted lottery is held when claims are around 7.5 and 2.5 in comparative strength. When claims are either very unequal or almost equal, there is less expected unfairness. The small chance of the very weak claim being satisfied discounts the large unfairness. As such, if one holds, as Broome and Piller do, that weighted lotteries ought to be held when claims are slightly unequal (because there is lesser expected outcome unfairness), then weighted lotteries ought to also be favoured when claims are very unequal. I have demonstrated why this ought to be the case through the use of two different measures of expected outcome unfairness.

## 2.6 Conclusion

This chapter has argued that in distributive cases with indivisible goods, unequal claims, and no other morally relevant factors, it is fairer to hold a weighted lottery than give the good directly to the individual with the largest claim. This conclusion stands in contrast to Lazenby's claim that one ought to rank giving the good to the individual with the strongest claim above holding a weighted lottery. I defended an account of fairness that is

<sup>&</sup>lt;sup>50</sup> To illustrate, in the case where Ann has a claim of 8.5 and Bill has a claim of 1.5, the expected value is calculated as follows: .85 \* 10 + .15 \* 70 = 19. The value of the good that Bill receives is multiplied by the distance between Ann and Bill to reflect the severity of unfairness, where it is less unfair that Bill gets the good when the claims are roughly equal and vastly more unfair if Bill gets the good when the distance between claims is very unequal.

sensitive to both changes in the size of claims and the number of claimants, and which adequately respects the claim of each affected person.

## Chapter 3

# Ambiguity and Fairness<sup>51</sup>

## 3.1 Introduction

A decision-maker is in a *risky* situation when they can assign at least subjective probabilities to possible states of the world and the associated outcomes of their actions. They are in an *ambiguous* situation when they cannot assign such subjective probabilities (Bradley and Drechsler, 2014:1231). The ethics of distribution under risk has been extensively studied (see, e.g. Diamond 1967; Broome 1984 1990; Otsuka and Voorhoeve 2009; Voorhoeve and Fleurbaey 2012, 2016; Bovens 2015; Hyams 2015 and Chapters 1 and 2 of this thesis). By contrast, the ethics of distribution under ambiguity has been much less discussed. This lacuna is important, because in many real-world cases, a decision-maker does not have a basis for assigning probabilities (Gilboa et al. 2009). For example, when evaluating climate policies, one often cannot obtain precise objective information about the probability of particular future scenarios; nor can one be expected to assign precise subjective probabilities to them (IPCC 2014a: 9).

This chapter considers what the implications of such ambiguity are when making a distributive decision. It argues against the use of a version of the *ex ante* Pareto principle applied to ambiguous prospects, or "Pareto under

<sup>&</sup>lt;sup>51</sup> This chapter is joint work with Alex Voorhoeve. I am lead author of the material presented here. I presented earlier versions of this chapter at the Warwick Graduate Conference in Political and Legal Theory, the VI Meetings on Ethics and Political Philosophy at the University of Minho (Portugal), the Society of Applied Philosophy Conference at the University of Edinburgh, the Workshop on Ethics of Social Risk at the University of Montreal, and the 14<sup>th</sup> Pavia Graduate Conference in Political Philosophy in Italy. I am grateful to the audiences at these conferences for their very helpful feedback.

Ambiguity", for short. It rejects this principle on two grounds. One is familiar from discussions of the *ex ante* Pareto principle under risk, namely that it conflicts with egalitarian concerns (see, e.g. Fleurbaey and Voorhoeve 2013 and Chapter 1 of this thesis). The other is novel: we show that prospects that are ambiguous at the individual level may be far less so at the social level. An ambiguity-averse distributor, who rightly considers the latter level, may therefore correctly favour a set of ambiguous individual prospects that eliminates social ambiguity, even when one would not prefer this set of individual prospects on any individual's behalf.

The argument proceeds as follows. Section 3.2 introduces a first set of cases and contrasts the implications of Pareto under Ambiguity with those of an egalitarian view. Section 3.3 considers the distinction between ambiguity at the individual and the social level, and argues that this provides a further reason to reject Pareto under Ambiguity. It also shows how the egalitarian and social-level ambiguity considerations are independent, and can either reinforce each other or work in opposing directions. Section 3.4 concludes.

### **3.2 Ellsberg Cases and Distributive Fairness**

The two distributive examples that are the focus of this paper are modifications of the two cases used by Daniel Ellsberg to formulate his eponymous paradox, which kick-started the literature on decision-making under ambiguity (Ellsberg 1961). Ellsberg's examples involve a selfinterested decision-maker choosing between different gambles, with varying degrees of probabilistic information, each with a prospective cash prize. He had both an empirical purpose (to argue that people would respond differently to ambiguous gambles than to risky gambles, and in particular, that many of them would be averse to ambiguity, in a sense to be defined

below) and a normative purpose (to argue that such a differential response was rational). The examples we will construct also involve a morallymotivated decision-maker selecting between different gambles containing varying degrees of probabilistic information. However, each gamble will lead to a distribution of an indivisible good to *others*. Moreover, we will be concerned with the following normative question: supposing that some aversion to ambiguity is rational (as Ellsberg argued), how should a decision-maker act under conditions of ambiguity?

### 3.2.1 Pareto under Ambiguity and Fairness

Suppose that a morally-motivated stranger, Tim, comes across two children, Ann and Bill. Both children are in urgent need of a kidney. Unless Tim acts, both Ann and Bill will die. There is only one kidney available, and it cannot be split between Ann and Bill. Tim must allocate the kidney on the basis of a random draw of a ball from an urn with 100 balls which are either red or black. There are three such urns to choose from. The first, *Risky Equal Proportion* urn, is known to have precisely 50 red and 50 black balls, so that the chance of either colour being drawn is equal. The second, *Ambiguous Urn*, has an unknown number of red and black balls, so that the chance of either colour being drawn. The third, *Risky Unequal Urn*, is known to have precisely 45 red and 55 black balls, so that the chance of red being drawn is 45% and of black being drawn is 55%. Ann wins if (and only if) red is drawn, Bill if (and only if) black is drawn. These urns are represented in Table 3.1. The number 1 represents the utility of receiving a kidney, and the number 0 represents the utility of not receiving a kidney.

| Risky, Equal Proportion Urn |              |                   |                 |  |  |
|-----------------------------|--------------|-------------------|-----------------|--|--|
|                             |              | Numbe             | Number of Balls |  |  |
| Act                         | Person       | 50                | 50              |  |  |
|                             |              | Red               | Black           |  |  |
| Α                           | Ann          | 1                 | 0               |  |  |
|                             | Bill         | 0                 | 1               |  |  |
|                             | Ambig        | guous Urn         | 1               |  |  |
|                             |              | Numbe             | Number of Balls |  |  |
|                             |              | 1                 | 100             |  |  |
| Act                         | Person       | Red               | Black           |  |  |
| В                           | Ann          | 1                 | 0               |  |  |
|                             | Bill         | 0                 | 1               |  |  |
|                             | Risky, Unequ | al Proportion Urr | 1               |  |  |
|                             |              | Numbe             | Number of Balls |  |  |
| Act                         | Person       | 45                | 55              |  |  |
|                             |              | Red               | Black           |  |  |
| С                           | Ann          | 1                 | 0               |  |  |
|                             | Bill         | 0                 | 1               |  |  |

## Table 3.1: Final utilities for Three Urn Case

In evaluating these urns, one relevant issue is how one should value them on each individual's behalf. As in Chapter 1, we shall assume that when evaluating risky prospects, one ought to respect the Von Neumann-Morgenstern axioms and therefore maximize expected well-being. With respect to *ambiguous* prospects, however, we will depart from standard decision theory by assuming that it is rationally permissible to have a

preference for known over unknown (ambiguous) probabilities on an individual's behalf. This means, for example, that on Ann's behalf, one may rationally prefer act A (a draw from the Risky, Equal Proportion Urn), which offers her a 50% chance of receiving the kidney, to act *B* (a draw from the Ambiguous Urn), which offers her anywhere from a 0% to a 100% chance of a kidney. Indeed, we shall assume that Tim takes the rationally permissible view that, when considering each person's prospects alone, prudence requires that one be moderately averse to ambiguity, so that on Ann's behalf, one should even favour act *C* (a draw the Risky, Unequal Urn), with its known 45% chance for her of getting a kidney, over *B* (a draw from the Ambiguous Urn). Though the rational permissibility of ambiguity aversion is disputed (see, for instance, Al-Najjar and Weinstein 2009), these assumptions seem to us plausible. In the absence of probabilistic knowledge, and when the stakes are as high as in this case, it seems sensible to be conservative. When evaluating the Ambiguous Urn on Ann's behalf, such conservativeness involves giving somewhat greater weight to the less favourable possibilities (which involve there being fewer than 50 red balls in the urn) and somewhat lesser weight to the more favourable possibilities (which involve there being more than 50 red balls). Such conservatism or pessimism has long been defended as a rational response to ambiguity (see, e.g. Hurwicz 1951, Ellsberg 1961, Gilboa and Schmeidler 1989, Binmore 2009, Gilboa et al. 2009). Indeed, empirically, when faced with gambles of this kind involving moderate probability gains, the predominant finding is that individuals are at least moderately ambiguity averse (for a recent survey of the literature, see Trautmann and van de Kuilen 2015, especially p. 96).

In sum, because of this moderate ambiguity aversion, Tim's preferences on Ann's behalf are *A* (the Risky, Equal Proportion Urn) > *C* (the Risky, Unequal Proportion Urn) > *B* (the Ambiguous Urn). By similar reasoning, the preferences on Bill's behalf are: C > A > B. This is because *C* gives Bill the highest precise chance of receiving a kidney, *A* gives him the second-highest precise chance, and *B* gives him an unknown chance.<sup>52</sup>

What should Tim, our decision-maker, do? One could offer at least a partial answer to this question by drawing on these two sets of rankings from each person's prudential perspective. For note that on both Ann's and Bill's behalf, *B* is strictly dispreferred to the other alternatives. And Tim should, one might think, follow this unanimous preference. In other words, one might draw on the following principle, which is the analogue under ambiguity of the *ex ante* Pareto principle under risk:

Pareto under Ambiguity:

<sup>&</sup>lt;sup>52</sup> It is possible to turn ambiguous probabilities into subjective probabilities by simply assigning equal probability to each of the outcomes in the three acts. But while this is certainly possible, one must ask why, in the absence of any information at all about the proportion of red and black balls in the urn, it is rationally required to do so. Isn't any such assignment of probabilities entirely arbitrary in the absence of such information? And can rationality require such an arbitrary assignment rather than a different decision procedure that simply keeps in mind the full range of possible probability distributions and responds somewhat conservatively to that. While such an approach is controversial, it is defended by many leading decision-theorists, and appears to be applied by a large proportion of respondents in surveys. This chapter therefore takes the view that it is worthwhile exploring what would follow if such an ambiguity-averse approach were rationally permissible and adopted by an egalitarian decision-maker. We add that if one were to assign equal subjective probabilities by adopting the Principle of Insufficient Reason, one gets the unpalatable result that one ought to be indifferent between the Ambiguous Urn and the Equal Proportion Urn. This is unpalatable for the reason that we think there are good reasons for preferring the Equal Proportion Urn to the Ambiguous Urn because we know for sure that each individual has an equal chance of receiving the kidney, whereas we do not know that in the Ambiguous Urn. In fact, there is reason to think that it would be unlikely that the proportion of balls would be equal..

(a, 'universal strict preference') If, for every person, a first alternative provides prospects that are prudentially superior to the prospects provided by a second alternative, then the first alternative should be strictly preferred to the second.

In what follows, we will also draw on the following extensions of this Pareto under Ambiguity principle, each of which has its analogue in familiar versions of the *ex ante* Pareto principle under risk:

(b, 'weak preference for all and strict preference for some') If, for every person, a first alternative provides prospects that are prudentially at least as good as the prospects provided by a second alternative and the first alternative provides prospects that are strictly superior to the second for at least one person, then the first alternative should be strictly preferred to the second;

(c, 'universal weak preference') If, for every person, a first alternative provides prospects that are prudentially at least as good as the prospects provided by a second alternative, then the first alternative should be weakly preferred to the second;

(d, 'universal indifference') If, for every person, a first alternative provides prospects that are prudentially precisely as good as the prospects provided by a second alternative, then the decision-maker should be indifferent between the two alternatives.

Pareto under Ambiguity would, in sum, demand that Tim prefer *C* to *B* and *A* to *B*. (It would be silent about how to rank *C* and *A*, and permit a choice between them to be made on other grounds.)

However, we will now argue that Tim ought not to choose fully in line with Pareto under Ambiguity. Instead, he ought to have following preference ordering over the three alternatives and choose accordingly:

Consistently with Pareto under Ambiguity (although not demanded by it), A ought to be preferred to *C* on egalitarian grounds, because *A* gives each individual an equal chance of receiving the kidney whereas C gives Bill a greater chance than Ann of receiving the kidney. In line with Pareto under Ambiguity, we submit that Tim ought to prefer *A* to *B*. But this is not merely (as Pareto under Ambiguity has it) because one would prefer A to B on each person's behalf, but also on grounds of equality. For A gives each individual an equal chance of receiving the kidney for sure, whereas, even though B treats Ann and Bill symmetrically, it is unknown in *B* whether Ann and Bill are given an equal chance. In fact, it is very likely that the proportion of balls, once revealed, is skewed in favour of one of the individuals.<sup>53</sup> Finally, and contrary to Pareto under Ambiguity, there is reason for Tim to select B over *C*. For even though there is uncertainty about the proportion of the balls in *B*, each individual is treated equally (by being given equally valuable prospects) and therefore fairly by Tim, whereas if Tim selected *C*, he would be knowingly favouring Bill over Ann. Moreover, such equal treatment, while it comes at the cost of the ambiguity-weighted value of one person's

<sup>&</sup>lt;sup>53</sup> The phrase "very likely" does not involve assigning a precise probability to the event that there will turn out to be a different number of red and black balls in the urn. Instead, it is merely based on the idea that among the one hundred and one different possible combinations of balls that may be in the urn, only one is an equal proportion of red and black balls. This is, in my view, enough on which to base a *range* statement, say, a "more than 90% chance" that the proportion of the balls will turn out to be skewed in one person's favour. But this is not to say that any precise chance is assigned to any event; nor could one conclude from it that there is an *equal* chance of Ann or Bill winning.

(Ann's) prospects, does not come at the cost of the expected goodness of the distribution of final well-being. For while *individual* prospects are ambiguous under *B*, the anonymized (and therefore, impartially considered) *social* distribution of final well-being is known and unaffected by the choice of *B* over *C*. What is best for each individual, from the perspective of their interests alone, ought therefore not to always determine what one ought to do considering both fairness and expected goodness.<sup>54</sup>

In sum, the proposed reason to depart from Pareto under Ambiguity is fairness. But there is more to be said about Pareto under Ambiguity. Notably, we have, as yet, barely explored the possible relationship between ambiguity at the individual level and at the social level. Nor have we considered reasons relating to inequalities in final well-being (as opposed to prospects). We turn to these topics in the next section.

# 3.3 Inequality in Final Well-being and Individual versus Social-Level Ambiguity

Suppose that Tim must choose between four ways of using a random draw from a single urn to distribute transplant kidneys. The urn contains 90 balls; 30 red balls and 60 white or black balls in an unknown proportion. This Single-Urn Case is outlined in Table 3.2.

<sup>&</sup>lt;sup>54</sup> However, one may press the following objection. In *B* the proportion of balls could be wildly skewed in the favour of one of the individuals. One may think that it is preferable to be somewhat unfair to Ann for certain in *C* rather than select *B* which may be *even more* unfair in the distribution of objective chances. What if, for example, there were 49 red balls and 51 black balls in *C*? Should Tim still prefer *B* to *C*? In response, Tim ought to still prefer the act with more ambiguity, because even though it may be the case that when the proportion of balls is eventually revealed the proportion of balls is unequal, Tim does not know when making the decision *who* it is that is favoured. Whereas if Tim chooses *C* he *knows* that he will be favouring Bill over Ann. If Tim cares about fairness, then he ought to select the act which has greater ambiguity when it is fairly distributed, even if he could instead select an act that had less ambiguity but known unfairness in chances.

|     |        | Number of Balls |       |       |
|-----|--------|-----------------|-------|-------|
|     |        | 30              | 60    |       |
| Act | Person | Red             | Black | White |
| D   | Ann    | 0               | 1     | 1     |
|     | Bill   | 1               | 0     | 1     |
| E   | Ann    | 1               | 1     | 0     |
|     | Bill   | 1               | 0     | 1     |
| F   | Ann    | 1               | 1     | 0     |
|     | Bill   | 1               | 1     | 0     |
| G   | Ann    | 0               | 1     | 1     |
|     | Bill   | 0               | 1     | 1     |

## Table 3.2: Final utilities for Single-Urn Case

As before, we will assume the following rational, ambiguity-averse preferences over these acts on behalf of Ann:

$$G \sim D > E \sim F.$$

And the following preferences on behalf of Bill:

$$G > D \sim E \sim F.$$

The reasons for these preferences are as follows. Acts *D* and *G* both give Ann a precise 2/3 chance of receiving a kidney. This is because if a black or a white ball is drawn from the urn, Ann will receive a kidney (60 out of the 90 balls in the urn). Act *E* (in which Ann gets the kidney if red or black is drawn) gives Ann a chance of receiving a kidney that ranges from 1/3 to 1 (inclusive). This is because, in the worst case, there are no black balls and

Ann has only a 1/3 chance; in the best case, there are 60 black balls and Ann gets a kidney for sure. Similar reasoning establishes that *F* also gives Ann a chance of winning that ranges from 1/3 to 1. Due to ambiguity aversion, when he is considering her interests alone, Tim will therefore prefer an option that exposes her to less ambiguity and will therefore prefer *D* or *G* to *E* or *F*.

As for Bill, *G* is the only option which gives him a known chance of receiving a kidney, namely 2/3. In contrast, *D*, *E* and *F* each expose Bill to an ambiguous prospect with a chance of a kidney ranging from 1/3 to 1. This is why Tim will prefer *G* to all of the others on Bill's behalf.

From the above individual orderings, it follows that Pareto under Ambiguity yields the following social ordering:

$$G > D > E \sim F$$

We shall now explore whether Tim should indeed adopt this ordering. In deliberating about what Tim's preferences ought to be, we should consider the aforementioned distinction between ambiguity at the individual level and ambiguity at the social level. Ambiguity at the individual level refers to ambiguity contained within an individual's prospects, such as the ambiguity that Ann and Bill are exposed to in prospect *E*. This is distinct from ambiguity at the social level. When a decision-maker evaluates an act in terms of the moral value of the distributions of final well-being that this act makes possible, ambiguity may or may not be present. For example, in terms of the distribution of final well-being *E* yields no ambiguity from the social perspective, because Tim knows that there is a 1/3 chance that two individuals will get a kidney and a 2/3 chance that only one individual will

get a kidney. Though Tim doesn't know who will get the kidney in the latter scenario, because he is impartial, insofar as he is considering the distribution of final well-being, he is indifferent whether Ann or Bill receives the kidney, and so indifferent to the events over which he lacks precise probabilistic information (that is, whether black or white is drawn). This is not true, however, for *D*. If Tim chooses *D*, it is not known what the chances are of the two possible patterns of final well-being. There is a chance ranging from 0 to 2/3 that both will receive a kidney (if a white ball is drawn). There is a chance ranging from 1/3 to 1 that precisely one person will receive a kidney. *D* therefore contains ambiguity at both the individual and social level, as it contains an ambiguous prospect for Bill and it is not known precisely how likely it is that particular anonymised patterns of final well-being will occur.

We will assume that when dealing with such social ambiguity, Tim rightly adopts the same degree of moderate ambiguity aversion, or conservatism, that he has on each individual's behalf. That is, insofar as he considers the anonymized distributions of final well-being to which they may give rise, he rightly favours the socially risky *E* over the socially ambiguous *D*.

We will now argue that, considering both each individual's perspective and the social perspective, Tim ought to adopt the following preference ordering:

G > F > D and G > E > D, with all relative rankings of F and E being permissible.

This ordering is the result of applying the following four criteria:

 (i) Reducing ambiguity in individual prospects, in the sense that, other things being equal, less ambiguity yields a more valuable prospect.

- (ii) Reducing ambiguity at the social level, in the sense that, other things equal, less ambiguity about the anonymized distribution of final well-being is better.
- (iii) Reducing inequality in the value or individual prospects, in the sense that, holding other things constant, less inequality is better.
- (iv) Reducing equality in final well-being, in the sense that, other things equal, less inequality is better.

The first two criteria are obvious, given the assumption of rational ambiguity aversion. The subsequent two egalitarian criteria can be justified as a natural extension of familiar egalitarian views. Just as under risk, equal chances can help mitigate the unfairness (and therefore badness) of outcome inequality, we propose that the equal ambiguity can help mitigate the badness of outcome inequality. This implies that insofar as the burden of ambiguity is present at the individual level, it is better that it is distributed equally. Of course, outcomes matter too; and more equal outcomes are therefore fairer. In sum, in line with egalitarian views which value both equality of chances and outcomes (Arneson 1997, Temkin 2001, Voorhoeve and Fleurbaey 2012, Hyams 2015), the view holds that one must therefore rank acts by taking account both the distribution of risk-or-ambiguity weighted prospects and the possible distributions of final well-being.

Applying these criteria, *G* is clearly best. First, there is no ambiguity about Ann and Bill's prospects. *G* therefore offers them prospects which are prudentially at least as good as any other alternative that Tim might choose. Second, there is no social ambiguity: it is known that there is a 2/3 chance that both will end up with a kidney and a 1/3 chance that both will go without. Third, the prospects for Ann and Bill are *ex ante* of equal value with

this option, and each is given equal chances for sure. Finally, there is no inequality in outcomes: either both will end up with no kidney or both will end up with a kidney.

While *G* is uniformly best, in ranking the remaining alternatives, our four criteria sometimes pull in opposite directions. Consider F. First, if F is selected, there is ambiguity from the individual perspective as the precise proportion of black and white balls is unknown. Both Ann and Bill have a chance of receiving a kidney ranging from 1/3 to 1. There is also ambiguity from Tim's point of view, as he knows only that the chance of both people receiving a kidney ranges from 1/3 to 1 and the chance of neither receiving a kidney ranges from 0 to 2/3. Both forms of ambiguity count against *F*; indeed, the fact that *F* is ambiguous at the social level is especially problematic in the light of the fact that *F* allows for an unknown probability of a disaster, in which neither gets a kidney. Standard ambiguity-averse decision criteria are conservative, and hence give special weight to the worst possible probability distribution over outcomes (see e.g. the Hurwicz criterion discussed in Binmore 2009 and Binmore et al. 2012 and the criterion proposed in Gilboa and Schmeidler 1989). This means giving special weight to the possibility that the chance of a white ball being drawn is in fact 2/3, and consequently to the possibility that this disaster occurs.

On the other hand, *F* does ensure full equality. For the prospects it gives Ann and Bill are *ex ante* subjectively and objectively equally valuable. Moreover, *F* gives Ann and Bill equal (though unknown) objective chances, as each will receive a kidney if a red ball is drawn, and each will receive a kidney if a black ball is drawn. So whatever the number of black balls turns out to be, Ann and Bill will have an equal objective chance of receiving a

kidney. Finally, whatever happens, there will be no outcome inequality: either both get a kidney or neither gets one. *F* is therefore perfectly fair, both *ex ante* and *ex post*.<sup>55</sup>

Next, consider *E*. If Tim chooses *E*, both Ann and Bill have ambiguous prospects. However, there is no ambiguity at the social level: Tim knows that the chance of both receiving a kidney is 1/3 and of one receiving a kidney is 2/3. As regards equality: the prospects of Ann and Bill are equally valuable, although it is very unlikely that they each possess objectively equal chances. Finally, there is likely to be inequality in final well-being. In sum, *E* represents a mixed bag: individual but no social ambiguity, some equality in prospects but possibly not in outcomes. How one ranks it compared to *F* is therefore a matter of how averse to ambiguity at the social level one is (the more ambiguity averse one is, the better *E* looks compared to *F*) and how averse to inequality in final well-being one is (the more inequality averse one is, the worse *E* looks). Striking an overall balance is a matter of judgment, with different rankings being reasonable.

Finally, consider *D*. If *D* is selected, then Bill, but not Ann, will face an ambiguous prospect. Ann's chances of receiving a kidney will be exactly 2/3, while Bill's will range from 1/3 to 1. Moreover, there will be ambiguity at the social level. For the chances of precisely one person receiving a kidney will range from 1/3 to 1, and the chance of both receiving a kidney will range from 0 to 2/3. Moreover, Bill will have less valuable prospects than Ann; Ann and Bill are also very unlikely to face equal objective chances. Finally, there is a chance of between 1/3 and 1 of inequality obtaining. In sum, *D* is

<sup>&</sup>lt;sup>55</sup> For further support for this claim, see Chapter 1 of this thesis. For the fairness of no one receiving a good, see Section 2.1.4 in fn. 27.

strictly worse than *E* and *F* in terms of both equality in prospects and outcomes. It is also worse than *E*, and on a par with *F*, in terms of ambiguity at the social level. It is only somewhat better than *E* and *F* with regard to individual ambiguity. On balance, it therefore seems that *D* ought to be Tim's least preferred alternative.

The resulting ranking disagrees with Pareto under Ambiguity. The most important point of disagreement is that it ranks *D* below *F* and *E*. It does so despite the fact that *D* offers Ann more valuable prospects than either *F* or *E* does, and offers Bill prospects that are just as valuable as either *F* or *E* does. There are two rationales for this departure from what one would weakly prefer on each person's behalf. First, *D* involves more inequality than *F*. In this respect, our argument to favour *F* contrary to Pareto under Ambiguity is in line with well-known egalitarian reasons for rejecting *ex ante* Pareto offered in the context of risk (see, e.g. Fleurbaey and Voorhoeve 2013). However, the second rationale is particular to the context of ambiguity: the strongest reason for favouring *E* over *D*, contrary to Pareto under Ambiguity, is not equality, but the absence (in *E*) as opposed to the presence (in *D*) of ambiguity at the social level. One can put it this way: even if Tim were an ambiguity-averse *utilitarian*, who cared not at all about inequality, then he would have reason to favour *E* over *D*, because *E* yields a distribution with one kidney with a probability of 2/3 and of two kidneys with a probability of 1/3, whereas *D* yields a distribution of one kidney with a chance ranging from 1/3 to 1 and a chance of two kidneys ranging from 0 to 2/3. Since ambiguity aversion at the social level implies giving special weight in decisions to the worst possible probability distribution (that is, for *D*, to a 100% chance of only one kidney being distributed), *D* counts as

worse than *E* for an ambiguity-averse decision-maker even in the absence of inequality aversion. This reason for rejecting an anticipatory Pareto principle is novel and particular to the context of ambiguity.

It is noteworthy, then, that *even a utilitarian would, if they were ambiguity averse, rank E over D, despite the fact that E levels down Ann's prospects.* In our view, this finding weakens the force of the levelling down objection in prospects that is sometimes raised against egalitarians (see, e.g. McCarthy 2015). For it shows that such levelling down is merely a by-product of moral considerations that arise not at the level of the individual, but at the social level. Equality is only one such consideration; the way people's prospects together generate knowledge (or the lack thereof) of the chances of various distributions of final well-being is another.

## **3.4 Conclusion**

The central conclusion of this chapter is that Pareto under Ambiguity should be rejected. It is worth emphasising the two distinct grounds we have offered for rejecting Pareto under Ambiguity. The first, familiar one is inequality aversion. The second, novel one is ambiguity aversion at the social level. These grounds are logically independent. Moreover, either one of these grounds is sufficient for rejecting Pareto under Ambiguity, but together they provide a more powerful case.

Though we have focused on stylized cases, the form of ambiguity-averse egalitarianism articulated here can be used to evaluate public policy alternatives. For example, many public policy evaluations (climate policy being an obvious example) do not contain precise probabilistic information with respect to potential events.<sup>56</sup> The egalitarian view can take account of this lack of precise probabilistic information when evaluating the fairness of policy alternatives. Another potentially useful feature of the egalitarian approach is that it is sensitive to distributional concerns in different policy alternatives. For instance, it may be known that one alternative leads to a higher precise chance of European countries benefiting from a mitigation policy, with the rest of the world having ambiguous chances of benefitting, whereas another alternative exposes all regions to equally ambiguous chances of benefitting. Even if the former involves less individual-level ambiguity, the latter may be deemed more desirable overall, because it is fairer.

In sum, Pareto under Ambiguity ought to be rejected. The ambiguity-averse egalitarian position outlined in this chapter provides a more plausible rationale for deciding how to act in distributive cases under conditions of ambiguity.

<sup>&</sup>lt;sup>56</sup> Throughout the latest Intergovernmental Panel on Climate Change report, a special uncertainty framework is used to communicate the likelihood of potential events. For example, one such statement is the following: "It is *likely* that land temperatures over Africa will rise faster than the global land average, particularly in the more arid regions" (IPCC WG II, 2014b: 1202). The term '*likely*' refers to a probability interval of 66%-100%.

## Chapter 4

## When is it Permissible to Impose a Risk of Harm?<sup>57</sup>

## 4.1 Introduction

Risk is a pervasive feature of modern societies. Many activities impose risks of harm on others.<sup>58</sup> Some risky activities such as driving a car or taking a commercial flight are generally seen as permissible, whereas other risky activities such as playing Russian roulette on an unsuspecting victim or drunk driving are not. This chapter examines when it is permissible to impose risks of harm on others.<sup>59</sup>

Section 4.2 examines four candidate views of what it is that gives risking in and of itself moral significance (rather than the outcomes of risky actions). By "moral significance" I refer to what it is about an imposition of risk *itself* that calls for justification. The first view states that risks of harm are not *themselves* harms, and therefore, if imposing a risk of harm were to possess moral significance it is not because of its harmfulness (Perry, 1995, 2003). A second view states that the potential impermissibility of imposing a risk of harm arises from the fact that a risky act may itself wrong the victim. A third view is that an imposition of a risk of harm's impermissibility stems from the fact that it curtails the victim's autonomy (Oberdiek, 2009, 2012). A fourth view uses a preference-based account to argue that an imposition of a

<sup>&</sup>lt;sup>57</sup> An earlier version of this chapter was presented at the VII Meetings on Ethics and Political Philosophy at the University of Minho in Portugal. I thank the audience for their helpful comments.

<sup>&</sup>lt;sup>58</sup> In this chapter I shall, for the sake of brevity, often refer to activity of imposing risks of harm as "risking".

<sup>&</sup>lt;sup>59</sup> I will not be here be concerned with the separate yet important question of whether there is a *right* against having risks of harm imposed. Sven Ove Hansson (2003), Madeleine Hayenhjelm & Jonathan Wolff (2011), Sune Holm (2016), and David McCarthy (1997) are notable examples of authors who discuss this particular question
risk of harm is *itself* a harm and thereby possesses moral significance (Finkelstein, 2003). I argue that each of these approaches fails to adequately capture the moral significance of risking.

Section 4.3 then outlines and defends a novel account of the moral significance of imposing a risk of harm, called the *Insecurity of Interests Account*. It is argued that what gives an imposition of a risk of harm its moral significance is that it renders the victim's interests less secure. This constitutes a distinct form of harm, and grounds the potential impermissibility of imposing a risk of harm. Section 4.4 concludes.

# 4.2 Accounts of the Moral Significance of Risking

This section critiques four views of the moral significance of risking. I argue that each view fails to provide a plausible account.

# 4.2.1 No Harm Account

The first runs as follows:

*No Harm Account*: Impositions of a risk of harm are not themselves harms. If they possess moral significance it is because of some other factor than the mere risk.

This view has initial plausibility because risks of harm *themselves* need not have any impact on the victim. This can be illustrated with the following example:

*Drunk Driver*: A drunk driver careers down a road, narrowly avoiding a pedestrian. The pedestrian is unaware of the actions of the drunk driver.

The drunk driver imposes a risk of harm on the pedestrian. However, the harm that the risky activity forebodes fails to materialise. He walks home unscathed and does not suffer from feelings of shock or fear. This is an example of an imposition of a "pure" risk of harm.<sup>60</sup> This can be contrasted with an "impure" risk of harm, where the risk does yield a harm. Before proceeding, I shall stipulate a definition of harm. A common account of harming states that an individual is harmed when her interests are set back (Feinberg, 1987: 35).<sup>61</sup> Feinberg makes the clarification that the sorts of interests that matter for whether an individual is harmed are the "important interests", which are "presumably of a kind shared by nearly all his fellows, in the necessary means to his more ultimate goals, whatever the latter may be" (Ibid: 38). For example, an important interest of mine is not being physically injured. If a person assaults me, they thereby thwart my interest in not being injured. The important feature in this conception of harm is that it can allow for material as well as non-material setbacks to a person's interests.<sup>62</sup> So although there appears to be no causal impact on the material interests of the victim, because they walk home unscathed, there may be an impact on the pedestrian's non-material interests.

<sup>&</sup>lt;sup>60</sup> Judith Thomson (1986: 173) coins this terminology. I make the extra stipulation that the victims of the risk imposition are not aware of the risk or the risky act. This is to avoid *other* possible harms that might arise from individuals being aware of the risk imposition. Robert Nozick (1974: 65-6) considers this possibility.

<sup>&</sup>lt;sup>61</sup> This definition is endorsed by the main figures in the risk and harm literature, such as Finkelstein (2003: 971), Oberdiek (2012: 351-3) and Perry (2003: 1285-7).

<sup>&</sup>lt;sup>62</sup> The terminology of "material" harm and "nonmaterial" harm is borrowed from John Oberdiek (2012: 342), who uses the former term to refer to the harms that ripen from a risky act, such as being hit by a car, and the latter term to refer to set backs to interests that have a non-physical impact on the victim, such as a set-back to a person's autonomy.

Can a risk of harm itself be a harm? Stephen Perry is the most notable defender of the *No Harm Account*.<sup>63</sup> According to Perry, there is no such thing as "risk damage", where risk damage is the harm caused only by a risk of harm itself (1995: 322).<sup>64</sup> In establishing this claim, Perry considers the English legal case of *Hotson v*. East Berkshire Area Health Authority. In this case, the plaintiff suffered an injury that initially went undiagnosed by the health authorities. The injury then developed into a much worse condition. According to the facts of the case, there was a 0.75 chance that the plaintiff's condition would have deteriorated even after an initial diagnosis. However, when the injury was eventually diagnosed later on by the health authorities it turned out that the chance of the condition's deterioration was certain, rather than 0.75. The plaintiff argued that this increase in the risk of deterioration was itself an injury that ought to be compensated. Perry claims that given an objective account of probability as relative frequency, it is true that if there were one hundred people with the initial injury, seventy-five would go on to suffer the deterioration anyway, and twenty-five would not. As such, Perry argues, it is true that each of the one hundred will end up having either a zero percent chance or a one hundred percent chance of having their condition deteriorate. Either a person is harmed or a person is not harmed. In presenting his view, Perry assumes that determinism is true. This is important, because if this is assumed then it becomes difficult to see that the plaintiff suffers risk damage as a distinct form of harm. As Perry

<sup>&</sup>lt;sup>63</sup> Perry considers the possibility that a risk of harm may itself be a wrong, even if it does not itself constitute a distinct form of harm (2001: 76).

<sup>&</sup>lt;sup>64</sup> Perry uses the term "damage" synonymously with "harm" and "injury" (1995: n.1). One concern with this nomenclature is that "damage" gives the impression of *physical* setbacks to interests, which would preclude the possibility of "risk damage" being something non-material. Further support this claim can be found when Perry states that "harm involves interference with some aspect of well-being, and such interference can ordinarily be expected to have objective existence in the physical world" (1995: 332).

claims, "if the Plaintiff's injury was treatable, then the defendant caused him *physical* damage; there is no reason to say that it also caused him another, separate harm that takes the form of risk damage" (1995: 334). Whereas, if the injury was, in fact, untreatable then the defendant didn't cause the plaintiff any damage. Perry argues that we do not need to know which particular partition of the reference class a particular person will be a member of:

"If the causal processes involved are deterministic then a distinction can be drawn *in principle* between the two categories of case, whether we have the knowledge to do so in practice or not, and so long as such a distinction is possible in principle, it makes no sense to claim that the plaintiff suffered, at the hands of the defendant, a peculiar, non-physical injury of the kind I have labelled risk damage" (Ibid: 334).

The important factor here is that because it is possible in principle to know whether a person will fall into a particular category, there can be no distinct form of "damage" over and above the damage that actually occurs. Risk damage does not occur, Perry claims, merely as a result of our ignorance about how things will actually turn out.

Before engaging with the view per se, we should note that Perry's case is one where the risk is caused by *allowing* the condition to develop on its own, rather than the risk arising from the *actions* of the health authorities. In Perry's example, the risk of harm arises naturally from the health conditions of the plaintiff. By contrast, the cases we are interested in involve *impositions* of a risk of harm. This is morally important, because we do not typically

think of naturally occurring risks, such as the chance of an individual's medical condition worsening if untreated, as something that can be imposed by another without their intervention.<sup>65</sup> There is a moral asymmetry between doing and allowing harm. It is typically worse to act so as to bring about a harm than it is to not act and allow a harm of the same magnitude to occur.<sup>66</sup> In what follows, I shall therefore focus on cases of imposing harm and not on Perry's own case.

One objection to Perry's view is that it merely assumes risk away. It takes a purely *ex post* view, looking only at final burdens and who will bear them. From this point of view, there are only those who are materially harmed and those who are not. However, from an *ex ante* point of view, the identities of the particular individuals who will be harmed are not known.<sup>67</sup> This is relevant, because although there may be some objective fact of the matter about who will or will not be harmed that is presently unknown, there may still be an epistemic risk where it is *believed* that an agent is imposing a risk of harm on others.

There are grounds for thinking that this ignorance about how the world will turn out does, *contra* Perry, possess moral significance. Individuals do not act from an omniscient perspective, even though it is determined what will

<sup>66</sup> Some authors object to this equivalence (Rachels: 2001). However, I shall assume for the sake of this chapter that it is worse for an agent to actively impose a risk of harm on another, than it is to fail to prevent a risk of harm of the same magnitude befalling another. The moral difference between doing/allowing is also considered in Chapter 5.

<sup>&</sup>lt;sup>65</sup> Perry stipulates that the augmented probability of the plaintiff contracting the illness is an "agent-imposed risk" (1995: 331), but this stretches what is intuitively thought of as an imposition of risk. Indeed, there might be less reason to think of the plaintiff being *harmed* if there is no active imposition of risk, as there is in *Drunk Driver*.

<sup>&</sup>lt;sup>67</sup> The moral significance of this distinction between *ex ante* and *ex post* points of view has been considered by authors such as Johann Frick (2013), Sophia Reibetanz (1998), and Fleurbaey and Voorhoeve (2013). This lack of information also forms part of the "prudential justification" for selecting a utility maximizing treatment on a person's behalf in Chapter 1.

happen. Parallels can be drawn with the evidence-relative sense of wrongness. Suppose that, completely unbeknownst to you, your phone has been tampered with such that when you phone your friend, the phone signal detonates a bomb that severely injures your friend's neighbour. Judith Thomson (1990: 231-233) argues that you have infringed a claim of your friend's neighbour. Thomson endorses an objective sense of "ought", where what one ought to do is determined by facts that the agent need not be aware of. So, for instance, you ought not to phone your friend, even though the reasons why you ought not to phone your friend are not available to you. T. M. Scanlon, however, argues that the point of view expressed by this view of "ought" "involves a kind of omniscience that prevents it from being identified with the point of view of a deliberating agent" (2008: 48). Something similar is true regarding one's ability to know facts about the outcome of a risky action where the outcome is determined but unknown.

There is also a second objection to this view. Consider the following case:

*Risky Rifles*: There are one hundred people sitting separately in a park. Alice (who is concealed in foliage) has one hundred rifles, each simultaneously aimed at a separate one of the one hundred people, and linked up to a central trigger. There are seventy-five bullets. At the moment of pulling the trigger, a random process distributes the bullets between the rifles, leaving seventyfive rifles with a bullet and twenty-five rifles without a bullet. Alice does not know which rifles contain bullets. Alice pulls the trigger.

In this case, seventy-five people will be shot and killed by Alice, and twentyfive will not be shot. Assuming determinism is true, it is the case that before firing the rifles each individual will either be in the reference class of "shot" or "not shot". Those who were lucky enough to have an empty rifle pointing at them are materially unscathed. According to Perry's view, there is no harm of the risk itself in either case. I submit that there are grounds on which one of the twenty-five can claim to be harmed by the actions of Alice.<sup>68</sup> For each of the twenty-five it was true that they could have been one of those who were killed. In *Risky Rifles*, each individual is exposed to an objective 0.75 probability of being shot. This is something that remains true even after the state of the world has been settled and individuals have been divided into groups of "shot" and "not shot".<sup>69</sup> To the extent that an individual has an interest in not having arbitrary interference with their more basic interests of bodily integrity and autonomy, say, having a risk of harm imposed runs contrary to this interest.<sup>70</sup> This is true regardless of whether a harm befalls an individual.

One may respond that although the process is a *random* one, it is still subject to deterministic laws, and as such it is possible in principle to determine whether a person will fall into a particular reference class. There can still be cause for complaint, however, given the openness of counterfactuals. A process is counterfactually open when "supposing that we initiate it, there is

<sup>&</sup>lt;sup>68</sup> Of course, those in the 75 also have grounds for complaint.

<sup>&</sup>lt;sup>69</sup> Perry notes that objective probabilities can exist in a deterministic universe. He states that "even if the preceding state of the universe determines the outcome of any given coin flip, the statement that there is a probability of 0.5 that the coin will come up heads expresses a fact about the physical world. As a practical matter, of course, we cannot predict whether a coin will come up heads or tails" (1995: 324).

<sup>&</sup>lt;sup>70</sup> One further explanation might be that the victims' basic interests in dignity have been violated by the gratuitous and intentional imposition of a risk of harm without their consent, thereby harming them because of the interest in guarding one's dignity (Placani, forthcoming).

no fact of the matter about what its outcome would have been if we had not initiated it" (Hare, 2012: 380). To illustrate, the following counterfactual is true: "If I had flipped the coin then it would have landed either heads or tails". But this is consistent with neither of these counterfactuals being true: "If I had flipped the coin then it would have landed heads" or "if I had flipped the coin then it would have landed tails".<sup>71</sup> Suppose we make two modifications to *Risky Rifles*. First, the random process that distributes the bullets between the rifles is in constant flux, such that if the trigger is pressed at t1 it is unlikely to generate the same distribution of bullets to rifles at t2 and the same for t3 and so on. Second, the precise time at which the random process distributes the bullets to the rifles is extremely sensitive to the precise pressure applied to the trigger. Suppose now that Ann decides not to pull the trigger, then given these facts about the case, the counterfactual that: "If Alice pulled the trigger then 75 people will be shot, and 25 people won't be shot" is true, but the counterfactual that: "If Alice pulled the trigger then these 75 particular named individuals, A, B, C ... would be shot, and these particular named individuals W, X, Y ... would not be shot" is false. To this extent, there is no fact of the matter about who would be shot and who would not be shot had the trigger been pulled. Because there is no fact of the matter about who will be shot, each individual has a complaint at the moment before Alice applies pressure to the trigger that *it could be them* who would be shot.

<sup>&</sup>lt;sup>71</sup> For this to be the case, the antecedent conditions have to be very specific. Hare gives the example of the following counterfactual: "If I had flipped the coin while its center of gravity was between 1.48318 and 1.48319 meters from the floor, applying between 2.899 and 2.900 Newtons of force to its upper edge at an angle of...then it would have landed heads." (2012: 382).

There would, then, be cause for complaint from each of the one hundred people as Alice begins to apply pressure to the trigger. This is because which particular determinate outcome will occur depends very sensitively on the timing of the trigger. Because it is not settled at which particular moment the guns will fire, the one hundred have a complaint that it could be them who is shot, when Alice's finger begins putting pressure on the trigger up until the moment that the trigger is activated. It would be implausible to suggest that the complaint evaporates when the gun actually fires a split second or so after pressure is applied to the trigger, and it becomes known who in fact is determined to live and who to die. Perry's view therefore doesn't seem to be able to do justice to these complaints.

Perry's overall position is, I submit, weakened when one considers what he writes about indeterministic causal processes. If *Risky Rifles* were amended so that there was an indeterministic rather than a random process that distributed the bullets to the rifles, then it would not be possible in principle to determine who will in fact be shot and who will not. Perry states that for such indeterministic causal processes, "there seems to be a true detrimental shift in position that is simply not present in the deterministic case" (1995: 337). In the indeterministic case, it is easier to see the complaint that individuals may raise at being subjected to the risk, since the chances of harm will be objective and in principle unresolvable in advance. But, by appealing to the fact of open counterfactuals, I have argued that the position of each person in the indeterministic version of *Risky Rifles* is sufficiently similar to the position of each person in the deterministic version up until Alice activates the trigger.

For the reasons given above, the *No Harm Account* is unsuccessful at demonstrating that there is no moral significance to risk itself.

# 4.2.2 The Wronging Account

Even if it is granted that there is no particular harm that a risk itself causes, there seems to be something objectionable about the risk imposer's actions. One strategy is to say that risk impositions are not themselves harms but rather *wrongs*. This view can be stated as follows:

*The Wronging Account*: Impositions of a risk of harm are morally significant when and because they wrong the victim.

One author who supports this view is Rahul Kumar who states, in relation to *Drunk Driver*, that:

"there is nothing suspect about the claim that one has been wronged by the drunk driver (expressed, perhaps, as resentment to him or anger directed towards him), simply in virtue of his having, without justification, taken your life in his hands by exposing you, even briefly, to so serious a risk" (2003: 103).

Kumar states that a person is wronged when another's conduct intentionally or negligently flouts particular requirements which flow from a person's standing as someone to whom justification is owed (2015: 30).<sup>72</sup> At first

<sup>&</sup>lt;sup>72</sup> These particular requirements include regulating one's conduct such that it "respect[s] the value of persons, as being capable of assessing reasons" with the ability to govern their lives accordingly (Kumar, 2015: 30). This account of what it takes to wrong someone is influenced by T. M. Scanlon (1998: 271). I bracket a discussion in this chapter of how contractualism more generally deals with cases of risk imposition. For authors who have engaged with this

glance, this does give a principled explanation of what happens in *Drunk Driver*. The drunk driver wrongs the pedestrian by taking his life into the driver's unreliable hands. However, John Oberdiek (2012: 355) objects to this approach to *Drunk Driver* because, he states, it is an account of culpability or blameworthiness, rather than an account of permissibility: it "revolves around the reasoning of the person imposing the risk, not the risk imposition itself" (Ibid). Although it true that it is a *fact* about the situation that the drunk driver is taking the pedestrian's life into his hands, the *Wronging Account* sees it as morally important that the drunk driver is acting recklessly, without a concern for the well-being of those around him. This is, as Oberdiek suggests, a problem for the *Wronging Account*. By focussing on the reasoning and intentions of the person imposing the risk, the account, so I shall now demonstrate, is insensitive to whether a wrongful act actually imposes risk.

To illustrate, consider the following case:

*Russian Roulette*: David plays Russian roulette on Eleanor without her knowledge as she walks down the street. There are six chambers and one bullet. It turns out that when he pulls the trigger there is no bullet in the chamber. Eleanor never finds out about the risk of harm imposed on her.

Although David does not physically or emotionally harm Eleanor, David can be said to wrong her. He takes her life into his hands. He flouts the requirements which flow from a person's standing to whom justification is

question, see Barbara Fried (2012), Aaron James (2012), Rahul Kumar (2015), and James Lenman (2008).

owed, since he acted in a way that was unjustifiable to the victim. But suppose now that the case is modified in the following way:

> *Deactivated Russian Roulette*: Unbeknownst to David, the Russian roulette gun that he picks up is a deactivated gun. There are six chambers and one bullet. It is not possible for the gun to fire. David aims the gun at Eleanor, believing it to be a regular six shooter, and pulls the trigger.

In this case, it is not possible for David to kill Eleanor. Eleanor is unaware of David's actions. According to the wronging account, David wrongs Eleanor. David has the same dispositions as in *Russian Roulette*, and is exhibiting exactly the same behaviour and intentions. However, it is the case that David could *never* have shot Eleanor due to the fact that the gun could not fire. The objective chance of Eleanor befalling harm was zero. As such, the *Wronging Account* delivers the verdict that Eleanor is wronged in both Russian roulette cases, regardless of whether an objective risk is imposed.

One might respond that the morally relevant feature is that the risk imposer reasonably *believes* that he is imposing a risk on the victim. David is imposing an epistemic risk of harm on Eleanor, even though he is not imposing an objective risk of harm on her. However, it is hard to see how a purely epistemic risk, by itself, could possess moral significance. In *Deactivated Russian Roulette* there is an epistemic imposition of a risk of harm given David's justified belief, grounded in the evidence available to him, that the gun works. What David does is clearly wrong, and he wrongs Eleanor through his attempt to take her life into his hands. But does the epistemic risk *itself*, taken aside from David's wrongful dispositions to endanger Eleanor's life, possess moral significance? It is implausible to think

that this pure epistemic risk can itself carry moral significance. This is because there is no possibility for Eleanor's interests to be negatively affected by the action. Eleanor walks down the street unscathed and unaware of David's action.

This brings out an important distinction between a subjective and an objective perspective which we encountered above in our discussion of Thomson's view. These perspectives can be taken on both wronging and harming. On the subjective perspective, whether David *wrongs* Eleanor is down to his (reasonable) beliefs about how he treats others. By contrast, according to the objective perspective, David can wrong another even if he doesn't believe that he is wronging her. I submit that the subjective reading is more plausible, because it focuses on the deliberations and attitudes of the acting individual, and thereby on how he dispositionally relates to the victim.<sup>73</sup>

According to the subjective perspective on *harming*, a person is harmed when they (reasonably, given the information available) believe that they are harmed (if one takes the perspective of the subject of the harm) or (if one takes the perspective of the agent) when the person acting (reasonably) believes that they are imposing a risk of harm. According to the objective perspective, whether a person is harmed is belief-independent. On this view, for example, my interests could be set back even though I am not aware of any such set back. The objective perspective about harming is the most

<sup>&</sup>lt;sup>73</sup> Feinberg, for example, builds the notion of intentionality or recklessness into an account of what it means for one to wrong another: "one person, *A*, can be said to *wrong* another, *B*, when he *treats him unjustly*. More precisely the injustice occurs when *A*'s act or omission has as its *intention* to produce an adverse effect on *B*'s interests, or is negligent or reckless in respect to the risk of such an effect" (1987: 108, emphasis added).

plausible one to take.<sup>74</sup> I could falsely believe that I have been harmed, even though none of my interests have been set back<sup>75</sup>, and another's belief that I have been harmed is not necessary or sufficient for me actually being harmed. The discussion of the Russian roulette cases demonstrates that the criterion that is relevant for determining whether an individual is wronged is distinct from the criterion that establishes whether a person has been harmed. David wrongs Eleanor, but he does not harm her.

The *Wronging Account* is insensitive to the presence of objective risk. David wrongs whether or not there is an objective risk that he might kill Eleanor. Furthermore, there is reason to think that epistemic risks *themselves* do not carry moral significance. The reason that the *Wronging Account* claims that Eleanor is wronged is due to the reasoning and intentions of David, and not due to the risk itself. The acts are identical in *Russian Roulette* and *Deactivated Russian Roulette*. As such, this account cannot explain the moral significance of the act itself of imposing objective risk.

#### 4.2.3 The Autonomy Account

Oberdiek (2009, 2012) provides the following autonomy-based account to explain the moral significance of risking:

*The Autonomy Account*: Impositions of a risk of harm are morally significant when and because they diminish the victim's autonomy.

<sup>&</sup>lt;sup>74</sup> Authors in the literature who agree with this include Finkelstein (2003: 973) Perry (1995: 333-4).

<sup>&</sup>lt;sup>75</sup> This may indeed serve as a harm if the thoughts *themselves* lead one to be anxious, angry etc. at the belief that one has been harmed. But this is a separate form of harm to whether or not one's interests have been set back.

On this view, impositions of risk are potentially impermissible because they foreclose previously safe options available to individuals. As Oberdiek states: "It is the bare curtailment of autonomy that risking can involve that calls for justification, grounding the moral significance of risking as potentially impermissible" (Oberdiek, 2012: 353). At the heart of this account is a particular conception of autonomy: that one ought to be able to plot one's own life and to have a range of acceptable options from which to choose in doing so.<sup>76</sup> To illustrate his view, Oberdiek uses an analogy of laying traps. A risk imposition curtails the autonomy of the victim just like a trap laid near the victim curtails the number of safe options available to them. For example, by playing genuine Russian roulette on you while you walk down the street, I am making the option for you of walking down the street less safe than it was before. I have made the set of options that are objectively acceptable smaller. I have thereby set back your autonomy interest in having your actions be ones you would wish to pursue with full knowledge of the situation. I also set back your interests in having what happens to you be determined by you and in leading your own, self-directed life. This diminution of autonomy is what gives risking its moral significance.

However, the *Autonomy Account* struggles with the following modification to *Drunk Driver*:

*Computer Assisted Drunk Driver*: This case has the same structure as *Drunk Driver*, with the modification that now, unknowingly, the driver has stepped into a computer

<sup>&</sup>lt;sup>76</sup> The conception of autonomy that is used is Joseph Raz's (1988: 410-11), where autonomy requires the availability of an adequate range of options.

assisted, rather than regular, car. The software in the car makes it the case that it will avoid all collisions. The software doesn't need to be activated, however, as the drunk driver's path takes the exact same route as *Drunk Driver*, narrowly avoiding the pedestrian.

In this case, the act itself of driving drunk is exactly the same as *Drunk Driver*, but it is not the case that the pedestrian has a safe option removed that would not otherwise have been removed. There was no possibility for the drunk driver to hit her. No safe options were – or could have been – removed. The walk along the pavement turned out to be just as safe for the pedestrian in *Computer Assisted Drunk Driver* as it was in *Drunk Driver*. The autonomy of the pedestrian in *Computer Assisted Drunk Driver* has obviously not been diminished. However, Oberdiek argues that in *Drunk Driver* "[t]he harm is imposed by virtue of the removal of safe options available to those in the drunk driver's vicinity, and that it is an aspect of the *act itself* of driving drunk." (2012: 355, emphasis added). But this reasoning supports the implausible conclusion that a safe option of a pedestrian can be removed even if there was no chance that they would be hit, since the *act itself* of driving drunk is exactly similar in both drunk driver cases.

A second objection to the autonomy account is that the potential impermissibility of impositions of risk is explained by something more fundamental than a curtailment of autonomy. For instance, suppose that while Bill is in a temporary coma, I stand beside him, point a Russian roulette gun at him, and pull the trigger. No bullet is fired. What explains the potential impermissibility of what I have just done? On Oberdiek's account, the potential impermissibility of the risk imposition arises from the

fact that it has curtailed Bill's autonomy. He has fewer safe options than before I aimed the gun at him. But Bill has no present capacity for exercising autonomous choice. And, while, at the moment of pulling the trigger, there was a chance that his future autonomy would be curtailed by his being killed, in fact, no such curtailment took place. Nonetheless, my act was wrong. I conclude that the curtailment of autonomy is not a necessary condition for the moral significance of risking and that the *Autonomy Account* doesn't give the intuitively correct explanation of what it is that makes my act impermissible. I have taken Bill's life into my own hands, and I have now made it vastly more likely that he will come to serious harm than if I hadn't aimed the gun at him.<sup>77</sup> Cashing out this state as one where the victim now has diminished autonomy is a second-order explanation. It may be possible to explain what is potentially impermissible about my actions with reference to autonomy, but this does not mean that it is the risk impositions' effects on autonomy that grounds its moral significance.

A third objection stems from a recent development of Oberdiek's view (ms.). Oberdiek adopts what is called the evidence-relative reasonable person perspective on risk.<sup>78</sup> On this view, what matters is the judgment that is formed about the nature of the risk that is being imposed, where this judgment must reasonably cohere with the evidence available. Oberdiek

<sup>&</sup>lt;sup>77</sup> I think that the reasoning that I have just supplied carries over to cases where Bill does have the capacity for autonomous choice. For example, if he is walking down the street and I aim the gun at him and pull the trigger, the moral significance of the risk I impose on him more plausibly derives from the fact that I have taken his life into my own hands. This remains true whether or not he has the capacity for autonomous choice.

<sup>&</sup>lt;sup>78</sup> This terminology was coined by Derek Parfit (2011: 150-1), who made the following distinctions. First, an act is *"wrong* in the *fact-relative* sense just when this act would be wrong in the ordinary sense if we knew all of the morally relevant facts". Second, an act is *"wrong* in the *belief-relative* sense just when this act would be wrong in the ordinary sense if our beliefs about these facts were true". Thirdly, an act is *"wrong* in the *evidence-relative* sense just when this act would be wrong in the *available* evidence gives us decisive reasons to believe, and these beliefs were true".

states that "the characterization [of the risk] must be *mutually* justifiable – justifiable to the patient as well as to the agent whose conduct is at issue. This simply means, though, that the risk imposition must be given the gravest characterization compatible with the agent's evidence-relative reasonable person perspective" (ms.: 45). Oberdiek also notes that this particular approach "captures the sense in which risks don't have a freestanding ontology but are rather the product of reasoning" (Ibid: 38, fn 54). Oberdiek plausibly states that a fact-relative interpretation of risk would be too demanding, since from the risk-imposing agent's perspective they cannot be expected to know the precise likelihood that harm will arise from their action.

I shall now argue, however, that if risk is to be characterised in the way that Oberdiek suggests it should, then it is difficult to see how risk impositions could negatively affect the victim's interests.<sup>79</sup>

To illustrate, recall *Deactivated Russian Roulette*. Does David impose a risk of harm on Eleanor? If risk is to be interpreted as purely epistemic – a feature of the evidence-relative beliefs of the risk-imposer – then David does impose a risk of harm on her. But suppose now that Eleanor in fact knows with certainty that David's gun is deactivated and therefore harmless. Eleanor then walks down the street with the knowledge that if David attempts to shoot her, he will fail. David believes that he is imposing a risk of death on Eleanor, and Eleanor knows that David is imposing zero chance of death on her. Is David still imposing a risk of harm? It would be implausible to claim that he is, since there is no objective chance of harm befalling Eleanor, and

<sup>&</sup>lt;sup>79</sup> A similar point is raised by Perry (1995: 332-3).

she knows this. Furthermore, the fact that the victim knows this defeats Oberdiek's claim that David curtails the autonomy of Eleanor.

For a risk of harm to *impact* on a person's non-material interests, it must be able to have an effect on them. If an imposition of a risk of harm is merely the evidence-relative beliefs of the imposer, then the risk itself need not directly impact on the victim's interests.<sup>80</sup> Oberdiek's example of impositions of risk being analogous to traps making a path unsafe gives the impression that a risk imposition has the power to *affect* the options of the victim. For risks to be harms, then there must be a discernible effect on the victim's interests. On a purely evidence-relative conception of risk, this cannot be the case. To tie this back to the distinction made earlier, whether a person is harmed ought to be determined from the aforementioned objective perspective.

In summary, the *Autonomy Account* fails as a plausible account of the moral significance of risking. A risk imposition of harm may set back an individual's autonomy, but such a set-back is not necessary for risking's moral significance.

# 4.2.4 The Preference Account

The fourth candidate view is a preference-based account for the moral significance of risking, most notably defended by Claire Finkelstein (2003). Finkelstein argues that a risk of harm *just is* itself a harm. This grounds the moral significance of risk impositions, because harming is itself morally significant and therefore calls for justification. Finkelstein adopts the view

<sup>&</sup>lt;sup>80</sup> It is still, nevertheless, deeply wrong to attempt to play Russian roulette on someone, but the moral significance of this act does not rest on there being a chance that he could be shot, rather it is on the evidence-relative beliefs of the risk imposer and the way he treats the victim.

that harm is the setback to a legitimate interest and states that "a common way in which the term 'interest' is used is as a synonym for 'preference'" (2003: 972). Finkelstein therefore appeals to a preference satisfaction conception of well-being. Finkelstein claims that individuals prefer that risks of harm are not imposed on them, and because whatever one disprefers sets back one's interests, the risk imposition itself constitutes a harm.<sup>81</sup> Call this:

> *Risk Is Harm*: "Exposing someone to a risk of harm itself harms him. That is, exposure to risk entails a reduction of an agent's welfare, regardless of whether the risk eventuates in outcome harm" (2003: 967).

On this view, a risk imposition reduces an agent's welfare whether or not the agent is aware of the imposition. A person is worse off than they would be had the risk not been imposed. To illustrate this point, Finkelstein considers the following example:

> *Defective Plane*: An airline is negligent in maintaining its planes. Unbeknownst to Cara one of two engines on the plane she is on quits mid-flight. Cara only finds out about this after the flight.

Finkelstein argues that Cara has been harmed. This is because the increased risk of harm that Cara was exposed to from the defective engine is itself a harm. Finkelstein claims that Cara has been harmed compared to a similarly situated passenger on a flight without engine failure (Ibid: 971). What might be doing the work here is the fact that Cara finds out about the defective

<sup>&</sup>lt;sup>81</sup> Finkelstein (2003: 967-970) treats symmetrically the increase in welfare that chances of a benefit give to individuals, and cases where an individual has a chance of receiving a burden. Oberdiek (2012: 346) also emphasises Finkelstein's use of a preference-based account of well-being.

engine only later. By finding out later, Cara's informed preference would be to not board the flight, and as such Cara is harmed by the fact that she did in fact board the flight. But what if we suppose that Cara never did find out about the defective engine – would she still have been harmed on the preference-based account? It is plausible to at least think that Cara has been *wronged* by the airline for its negligent maintenance of the aircraft. The preference account has difficulties accounting for why it is that Cara has been harmed if this is the case.

There is reason to think that the preference-based account cannot satisfactorily answer this question. Consider the well-known point that what we actually prefer might also end up being bad for us (Griffin: 1988: 11). A rational preference account can avoid the pitfalls of an actual preference view, but is vulnerable to a different set of objections. For example, Oberdiek argues that rational preferences are fully informed preferences; preferences formed as if one has all the available information pertaining to what one would prefer. From this perspective, he argues, "risk just disappears" (2012: 346). This is because if preferences were formed with all available information, one's preferences would reflect knowledge of how the world is and will actually be, and will therefore be devoid of risk. Further, it is not clear how precisely a person's life can go better if a preference is satisfied without one being aware of the fact (Parfit, 1983: 495-6). Our intuitions about whether Cara suffers a loss in well-being might be influenced by the fact that she also stands to benefit from the flight. To eliminate such considerations of benefit, consider the following alternative:

*Polynesian Islanders*: A plane flies over an isolated inhabited Polynesian island, thereby imposing a very

small risk of death of the islanders through falling debris from the plane. The islanders do not stand to (or wish to) benefit from any of the benefits from this particular flight, or from air travel more generally.<sup>82</sup>

Suppose, too, that the probability that Cara befalls a material harm is the same in *Defective Plane* as it is for an islander befalling material harm in *Polynesian Islanders*. I submit that it is implausible to think that the islanders suffer a loss in welfare because air travel imposes a risk of harm on them, and this is something that they disprefer. There is no apparent objective harm. One may counter that their interests are being set back by the fact that they do not stand to benefit from either the particular flight or the practice of air travel generally. Further, one may argue that these islanders are worse off than similarly situated islanders who are not under a flight path. The islanders are obviously not materially harmed by the plane flying above them. But it is difficult to see how the islanders could be non-materially harmed by the flight either. This is because there is a failure of the risk imposition to have a causal impact on the interests of the islanders.

In response, Finkelstein could argue that a risk imposition of harm sets back a second-order interest; an imposition of a risk of harm frustrates an interest we have in not being exposed to risks of harm. Because we prefer to not be exposed to risks of harm, a risk imposition sets back this second-order interest. However, Perry claims that "risking cannot be regarded as adversely affecting any interest that has a strong or plausible claim to be in

<sup>&</sup>lt;sup>82</sup> There has been debate considering the justifiability of risky activities that impose risks of harm on some people without any corresponding benefits. See, for example, Elizabeth Ashford (2003: 298), Rahul Kumar (2015: 34-35), Veronique Munoz-Darde (2013), and Michael Otsuka (2015: 79-80).

the set of core or primary interests" (2003: 1306). For Perry, second-order interests are defined recursively, with an example being the interest a person has that others do not subject them to a risk of physical injury. It is an interest that others do not attempt to set back an important interest. Perry says that if one acts contrary to the aforementioned interest then his "strong intuition" is that you do not harm them. This is because the further away we get from the core interests that constitute our well-being, the less likely we are to say that an infringement of this interest itself constitutes a harm. This is a problem for the preference-based account because harmful risks are defined recursively to the harm that the risk threatens. As such, the preference-based account does not give a satisfactory response to what it is that harms the Polynesian islanders. If the harm is that the islanders prefer not to have a risk of harm imposed on them, then the account is vulnerable to Perry's critique.

The preference account is found wanting as an account of the moral significance of risk impositions of harm. There are problems with grounding the account in the preference account of well-being. It is also implausible to think that an individual is harmed when they are exposed to risk merely because this is something that they disprefer.

#### 4.2.5 Summary

In summary, the four accounts fail to pinpoint to the moral significance of risking. The *Wronging Account* is insensitive to whether an objective risk of harm is imposed, and the *Autonomy Account* and the *Risk Is Harm* view both attempt to ground the moral significance of risk in risks alleged harmfulness. I have argued that these attempts fail. A successful account of the moral significance of risk has to recognise the fact that risks themselves

have no causal impact on the victim's material interests, and that the risk, if it harms, must harm the victim "objectively" – independent of the beliefs of the imposer or the victim. My view, which I will now develop, meets these desiderata. This view is grounded in an account of what it means to have secure interests.

#### 4.3 Risk Impositions and the Insecurity of Interests

#### 4.3.1 The Insecurity Account

This section defends the following view:

*Insecurity of Interests Account*: An imposition of a risk of harm is pro tanto impermissible when and because it renders the victim's interests objectively less secure than they were before the imposition.

The starting point for the account is the fact that mere risk impositions have no causal impact on the victim. Nonetheless, they can still harm the victim. Things can be bad for a person even if they are not aware of those things. For instance, Thomas Nagel writes: "the discovery of betrayal makes us unhappy because it is bad to be betrayed – not that betrayal is bad because its discovery makes us unhappy" (1979: 5). Risk impositions, I claim, operate in much the same way. It is bad for me that someone, unbeknownst to me, aims a Russian roulette gun at me and pulls the trigger, but we do not have to spell this out with an account of preference.

To unpack this claim, I shall adopt the following terminology from Judith Thomson: "if it would be good for X to get a thing Z, and Y makes it probable that X will get a thing Z, then Y gives X an *advantage*" (1990: 244). Correspondingly, Y *disadvantages* X by making it probable that X gets Z,

where Z is a something that will be bad for X. The size of the advantage or disadvantage is partly determined by the amount by which Y increases the probability that X will get Z and partly by the magnitude of the good or ill that Z would bring (Thomson, 1990: 170). Higher probabilities will confer a greater (dis)advantage, and lower probabilities will confer a smaller (dis)advantage, other things being equal. Being put at a disadvantage is a worsening of the victim's situation. It is better for the victim not to be subjected to a worsening of their situation.

An individual is disadvantaged by a risk imposition in virtue of it rendering their interests insecure. I shall make the following argument:

- The moral significance of an imposition of a risk of harm is that it renders the victim's interests less secure than they were before the imposition.
- 2. Rendering an individual's interests less secure is itself a harm.
- 3. An action that harms an individual is pro tanto impermissible.
- 4. Therefore, an imposition of a risk of harm is pro tanto impermissible when and because it renders the victim's interests less secure than they were before the imposition.

The core premises are 1 and 2. I shall assume the truth of premise 3.<sup>83</sup> I will now address premises one and two in turn.

<sup>&</sup>lt;sup>83</sup> One might argue that harming a person may be permissible if there are some other sufficiently weighty reasons for the harm, such as in order to avoid a greater quantity of harm. But my point here is that there is merely a "case to answer" for the permissibility of an act if it harms another.

#### 4.3.2 Premise One

Once it is settled what an individual's important interests are<sup>84</sup>, individuals have a further interest in the *security* of those interests. Pure risks themselves do not have a causal impact on the victim's core first-order interests since the risk does not ripen into the harm that it portends. But pure risks can make it the case that some interests of an individual are made less secure by being threatened by the risk imposing act. Recall the following case:

> *Russian Roulette*: David plays Russian roulette on Eleanor without her knowledge as she walks down the street. There are six chambers and one bullet. It turns out that when he pulls the trigger there is no bullet in the chamber. Eleanor never finds out about the risk of harm imposed on her.

Eleanor's situation is worsened by David's playing Russian roulette. It is bad for her that the gun is aimed at her, even without her knowing. Eleanor's situation is worsened by David because the gun being aimed at her makes it the case that her interest in bodily integrity is being rendered less secure than before the gun was aimed at her.<sup>85</sup> This is true even though there isn't a causal impact on this core first-order interest. Rather, there is an impact on Eleanor's interest in her interests being secure. It is important to distinguish this type of explanation from the one given by the *Autonomy Account*. According to that account, the moral significance of risking is derived from

<sup>&</sup>lt;sup>84</sup> Joel Feinberg (1987: 38) leaves it open what the comprehensive list of core interests would be, but he provides examples of such interests: the integrity and normal functioning of one's body, at least minimal income and financial security, physical health and vigour, etc. <sup>85</sup> Of course, there is also something objectionable about David's conduct, but as I have argued in Section 4.2.2, it is not his conduct *itself* that explains the moral significance of imposing risks of harm.

the fact an individual has fewer acceptable options. This explanation surely applies to *Russian Roulette*, as David makes the option of walking down the street less safe for Eleanor. Eleanor's option of walking down the street is less safe than before the gun was aimed at her, but this is not what is fundamentally significant about David's actions. The risk itself worsened Eleanor's situation and threatened her interest in bodily integrity. It may be that as a *consequence* of this that Eleanor's autonomy is diminished, but this is not what, at base, is morally significant about the risk, which is that it renders the interests of Eleanor less secure. The risk itself interferes with this second-order interest. This sets-back her second-order interest, even though there isn't a causal impact on the first-order interest.

#### 4.3.3 Premise Two

The interest in having security of one's core interests is a second-order interest that is of great importance to an individual. Having one's interests made insecure without one's knowledge is bad for that person. This interest can be set back by impositions which threaten the security of an individual's particular, first-order interests. A set back to important interests is itself a harm. For example, imposing a risk on someone such that there is a chance that their necessary medical supplies are destroyed makes the victim's interest in the maintenance of their physical health insecure. This harms the person even if turns out that they never need the medical supplies and are ever aware of the imposition of risk. Another example is if I put a person's life savings at the mercy of a fair coin flip where either the savings are destroyed if the coin lands heads, or left as they are if the coin lands tails. To the extent that that person has interest in financial security, I make that

interest less secure than it was before my intervention. In both of these cases, my action interferes with the victim's second-order interest.

Seth Lazar has recently argued that security is a "robustly demanding good" (2017: 8). This is because our enjoyment of the good depends not only on how our actual life goes, but also how our life goes in the counterfactual scenarios where we may not be so lucky. Lazar argues that "we are insecure to the extent that others make our avoidance of wrongful harm depend on luck" (Ibid). Furthermore, Lazar claims that "the more you depend on luck, the less control you have over your life, and so the less autonomous you are. Autonomy is non-instrumentally valuable, so its constituent parts – such as control over whether your *most important interests* are satisfied – are non-instrumentally valuable too" (Ibid; italics added). To this extent, the security of interests is something of high importance for an individual. There is good reason to think that, given the importance that the security of our interests has for how our life goes overall, there is a distinct locus of an interest harms us because of the value of the core interests such as autonomy that it protects.

#### 4.3.4 Objections and Replies

A first objection is that the insecurity of interests account is merely redescribing risk. Risk impositions – by definition – make an individual's interests less secure. In response, the insecurity of interests account provides a unique conceptualisation of how a probability of harm can *itself* set back a second-order interest without having a causal impact on the victim's firstorder, core interests of bodily integrity, autonomy etc. The account identifies a particular non-material interest that impositions of a risk of harm set back. As such, it is not a mere re-description of risk. One may object that

identifying this second-order interest leads to an implausible inflation of harms. Risks are imposed on others all of the time, and as such this secondorder interest is constantly being set-back. The *Insecurity Account*, however, provides a framework for establishing what it is about impositions of risks of harm that call for justification. It may be the case that a great many of the risks arising from transportation, nuclear power, and other risky social practices are *justified* in light of the benefits they bring, even though they render individuals' interests less secure. The insecurity account provides a framework for when an imposition of a risk of harm is pro tanto impermissible.

Second, one may object to premise two and claim that set-backs to secondorder interests do not harm a person. This is the same objection that was levelled against the preference account. In response, the interest in security of one's interests is not defined recursively in the way that the interest that others do not attempt to set back an important interest is. This is because the second-order interest introduces a new object of concern that individuals plausibly have an important stake in: that their interests not only be promoted, but that they also be secure. The Insecurity Account therefore avoids the objection. This is not the case for the preference-based account because the interests that are set back by impositions of risk are defined recursively. To support the idea that this second-order interest is very important, recall Feinberg's criteria for "important interests": they are "presumably of a kind shared by nearly all his fellows, in the necessary means to his more ultimate goals, whatever the latter may be" (1987: 35). The second-order interest in the security of one's interests is a strong candidate for inclusion in the set of a person's core interests. A person's life

goes better when, not only do they avoid set-backs to their interests, or when their interests are furthered, but when these interests are secure; when there is the absence of a *chance* that the interests will be set-back.<sup>86</sup>

One might deny that one's life necessarily goes better if one's interests aren't rendered insecure. If someone plays Russian roulette on Eleanor without her ever finding out, and she is unscathed, she is still able to pursue her life's projects. But, I submit, even if it turns out that she was not physically harmed by such a risk, we can still say that she was made worse off. Because at the moment the trigger is pulled, we do not have epistemic access to whether she will be shot; her second order interest can be deemed insecure. And even looking back with knowledge of what happened, one could sensibly say that her life was not secure and therefore that her interest in such security was thwarted. It turned out that she was fortunate not to be killed. This claim can be strengthened by imagining that Eleanor has Russian roulette played on her every day without her knowledge. As her life unfolds we can say that her interests are being made insecure by the evil game since there is a possibility that she will be killed, yet we do not know which day, if at all, this will happen.

A third objection is that the *Insecurity of Interests Account* is vulnerable to Perry's critique of the idea that risk itself does damage. Recall that Perry

<sup>&</sup>lt;sup>86</sup> Shelly Kagan argues that "it is one thing for a person to be well-off, and another thing for that person's life to go well" (1994: 318). According to Kagan, changes to well-being involve changes in either the body or the mind. Whereas, changes in how one's life goes need not have any effects on the mind or body of an individual. For example, a person may be enjoying decent levels of well-being, but it so happens that they are being deceived by everyone around them. Their spouse cheated on them, their relatives only pretended to like them in order to use their wealth etc. Kagan argues that the deceived person's *life* is not going well, but this can be distinguished from the person's well-being, which is unaffected by the deception. Oberdiek also draws on this wider sense of a person's "normative life", which are one's interests that go beyond mere facts about one's body and mental states (2012: 351). It is in this wider sense of a person's life that impositions of risk have an effect.

claimed that because determinism is true it is possible in principle to determine whether one will in fact come to be harmed as a result of a risk, or whether one will not. As such, there is no harm to the risk *itself*. In response, I note that one's interests can be made insecure even though risky processes are governed by deterministic laws. To illustrate, consider the following case:

> *Poison Assassin*: Jane slips a potentially lethal slow-acting poison into Kevin's tea. The poison may have no effect on Kevin, or it may eventually kill him. Which particular result obtains will be revealed in one years' time when the poison has its final effect on Kevin.

The poison and its effects on Kevin are subject to deterministic laws. However, I submit that Kevin has been harmed due to his interests being made less secure than they were before the poison was administered. The act itself of putting the poison in the tea made Kevin's interests less secure. This is because the risk *itself* constitutes an interference with the secondorder interest, even if there is no outcome harm from the risk. This is true even though he will either be killed by the poison, or not killed. Jane has made it the case that Kevin's interest in the integrity and normal functioning of his body is threatened. As such, Jane harms Kevin even though it is not settled yet whether he will in fact be killed by the poison.

A fourth objection to the *Insecurity of Interests Account* is that there is no apparent distinction between itself and the *Wronging Account* with the verdicts it gives. The accounts differ in the feature that they pick out for determining what it is about a risk imposition that makes it impermissible. I have argued that the "wronging criterion" is unsuccessful, and as such the

accounts differ in their success at explaining cases. The two views come apart in cases that highlight the presence or absence of objective risk. I argued that the wronging account is insensitive to the presence of objective risk. For example, in *Deactivated Russian Roulette* the wronging account states that the risk imposition of harm is potentially impermissible in the same way that it is in *Russian Roulette*. Whereas, the insecurity of interests account would draw a distinction between the two cases due to the fact there is no set-back to a second-order interests in the *Deactivated* case, but there is such a set-back in *Russian Roulette*.

#### 4.4 Conclusion

The *Insecurity of Interests Account* plausibly explains why risk impositions of harm may themselves be potentially impermissible. For example, in *Drunk Driver* the driver sets back the pedestrian's interest in the security of his first-order interests, even though there is no causal impact on the latter. The *Autonomy Account* and the preference-based account both suggest that risk impositions of harm are themselves harms because they respectively set back one particular interest (in leading one's own life) and a preference. The *Wronging Account* plausibly accounts for the objectionable nature of the reasoning and intentions of risk imposing individuals in the various cases considered in the chapter, but it fails to provide an account of how risk imposing acts themselves can have moral significance. The *Insecurity of Interests Account* provides a plausible rationale for how an individual can be harmed when there is no causal impact on their first-order interests.<sup>87</sup> This

<sup>&</sup>lt;sup>87</sup> Although this chapter has focused on the moral significance of imposing risks of harm, there is scope for a development of the theory to account for what is termed "social risk". Johann Frick (2015: 178-9) coins the term social risk to describe a set of risky actions that have the following features: (i) affect a large number of people; (ii) give burdens to some individuals which are greater than the benefits to others; (iii) rarely affect the same person twice, and as a result, one cannot assume that over time everyone will come to benefit from the action; and (iv) the risky action is intuitively permissible. One suggestion for how the

second-order interest is a plausible candidate for inclusion within the set of an individual's core interests.

Parallels can be drawn with the discussion of the value of chances in Chapter 1. Chances of a benefit can have positive value for an individual. The *Insecurity of Interests Account* explains why it is that impositions of a risk of harm can themselves be of negative value for an individual in virtue of setting back an interest in having security of one's interests. To illustrate, consider the following modification to *Poison Assassin*. Jane slips a potentially lethal poison into Kevin's tea, and thereby renders his interests less secure. This chance has *negative* value for Kevin. Now suppose that another individual, Liam, makes it such that there is a probability (let's assume it's the same probability that the poison will take effect on Kevin) that Kevin will have access to an antidote for the poison. Here Liam is rendering Kevin's interest more secure. This chance has *positive* value for Kevin. As such, there is reason to think that there is symmetry to the value and disvalue of chances.

*Insecurity Account* can be adapted to social risk is through an account of when an individual can accept the burden of insecure interests for the sake of greater (expected) benefits in return. Authors who appeal to expected benefit in a similar way include Frick (2015), Sven Ove Hansson (2003) and Oberdiek (2009).

# Chapter 5

# Providing Aid and Foreseeing Harm<sup>88</sup>

# 5.1 Introduction

Sometimes, in order to aid individuals, one must do things that will foreseeably enable evildoers to better pursue their harmful aims. This chapter addresses the following question. In the pursuit of the greater (expected) good, when is it permissible to foreseeably contribute to the harm of innocent others through such enabling of the agency of evildoers? I provide what I call the *Moral Purity Account* to both explain when it is permissible to provide aid in such cases, and determine when aiders bear moral responsibility for the harms imposed by another.

The chapter consists of three sections. I begin in Section 5.2 with two realworld motivating cases from Rwanda and Somalia, where the provision of aid led to the foreseen, but unintended, imposition of harm on some through the agency of militant groups. I then outline a more general case, called *Aid Case*, where the provision of aid enables an evildoer to impose harms that would not have otherwise been imposed.

In Section 5.3, I examine Jennifer Rubenstein's (2015) recent conceptualisation of the predicament facing aid workers in *Aid Case*, after considering two candidate views. Rubenstein offers what she calls the "spattered hands approach". According to this view, aid workers sometimes have a responsibility to grudgingly accept contributing to injustices

<sup>&</sup>lt;sup>88</sup> An earlier version of this chapter was presented at the Brave New World conference at Manchester University. I thank the audience for their helpful comments. I also wish to thank Luc Bovens for very helpful comments on a previous draft.

perpetrated primarily by others. When they do so, they bear a degree of moral responsibility for these harmful actions.

In Section 5.4, I outline and defend the *Moral Purity Account*. I argue that this account is superior to Rubenstein's approach. Against Rubenstein, I argue that when suitably motivated, aiders are not blameworthy for the actions of evildoers. Instead, I argue that although aid workers may be knowingly causally responsible for the circumstances that lead to an evildoer imposing harm, this does not entail that they are *morally* responsible. I argue that agent regret is a more suitable reaction to foreseen causal responsibility for the harmful effects of the agency of another.

## 5.2 Motivation

Consider the following two real-world cases:

*Somalia Famine*: The 2011-2012 East Africa drought caused a famine that lead to the deaths of around 250,000 people. Parts of the most severely affected regions were controlled by the Al-Shabaab militant group, widely considered a terrorist organisation. Most aid agencies were expelled or prevented from operating in the areas under Al-Shabaab control. Often, remaining aid agencies were permitted to operate in the controlled areas only if they paid the Al-Shabaab militants. Al-Shabaab desired to co-opt and materially and politically benefit from the provision of aid. Aid provision helped to alleviate the

suffering of civilians, whilst also strengthening the hand of Al-Shabaab.<sup>89</sup>

*Rwandan Genocide*: In 1994, Hutu-dominated Rwandan Armed Forces (Forces Armées Rwandaises or FAR) fled with around 2 million predominantly Hutu civilians to humanitarian camps in neighbouring Tanzania, Burundi, and what was then Zaire. These civilians included many genuine refugees. "Ex"-FAR members siphoned off vast quantities of aid intended for civilians and used the civilian presence as an opportunity to militarily regroup. The provision of aid ended up sustaining the perpetrators of genocide, and helping them reorganise.<sup>90</sup>

In both cases, individuals (who we can call "rescuers") attempt to save some from severe misfortune.<sup>91</sup> A foreseen but unintended consequence of their attempt is that it facilitates the harmful aims of others. The provision of aid is inextricably tied up with facilitating the potential wrongdoing of others.<sup>92</sup> There are two ways in which the actions of rescuers can contribute to the harmful projects of others. The actions may help *sustain* the harmful aims of

<sup>&</sup>lt;sup>89</sup> Abdi Aynte & Ashley Jackson, 2013, 'Al-Shabaab engagement with Aid Agencies', *Overseas Development Institute*, Policy Brief 53. See also: Ashley Jackson, 'A deadly dilemma: how Al-Shabaab came to dictate the terms of humanitarian aid in Somalia', *Overseas Development Institute*, 2013, <u>http://www.odi.org/comment/8066-al-shabaab-somalia-</u> <u>negotiations</u>

<sup>&</sup>lt;sup>90</sup> Rubenstein, Jennifer, 2015, *Between Samaritans and States: The Political Ethics of Humanitarian INGOs*, Oxford University Press: Oxford, pp. 88-89. Further details of this case can be found in "Rwandan Refugee Camps in Zaire and Tanzania 1994-1995", *Medicins San Frontieres*, Laurence Binet et al.

<sup>&</sup>lt;sup>91</sup> For the purposes of argument, I shall assume that the individuals in need would not be aided by other individuals if the rescuers failed to provide aid.

<sup>&</sup>lt;sup>92</sup> In reality, at least some of those who stand to gain from the provision of aid are not threatened by an aggressor, and at least some of those who are threatened by an aggressor do not stand to gain from the provision of aid. It might make a difference to the permissibility of intervening if it would be to the expected advantage of each individual to run the "gamble" of being the one who is killed by the aggressor, in order to have the greater chance of being one of those who is saved.
others. For example, Al-Shabaab's hand was strengthened through the provision of aid. This helped Al-Shabaab sustain their activities. The actions of rescuers may also facilitate *new* harmful threats; threats that would not have existed had the rescuers not provided aid. This is evident in *Rwandan Genocide*, where the humanitarian camps provided by rescuers enabled the "ex-FAR" members to militarily regroup so that they could carry out new threats. For the purposes of this chapter, I shall assume that the harmful acts of evildoers are *new* acts; acts that would not have occurred were it not for the acts of rescuers.<sup>93</sup>

In order to highlight the structure of the motivating real-world cases, I shall refer to the following stylised case throughout the paper:

> *Aid Case*: Rescuers are deciding whether to provide essential aid to a group of 100 individuals. A foreseeable consequence of the rescuers providing aid is that a villainous aggressor is enabled to establish new threats which will kill 20 separate individuals. If the rescuers decide not to provide aid, then those who require aid are left worse off than they would have otherwise been. If the rescuers provide aid then 100 will be saved, but the villainous aggressor will kill 20 individuals, who did not have a chance of being aided.

<sup>&</sup>lt;sup>93</sup> It may be the case that *new* threats arise through the sustaining of a pre-existing threat. For example, by facilitating a terrorist group one may be helping it *sustain* its threats, but also, in virtue of doing this, the terrorist group is allowed to pursue *new* harmful threats. These threats would not necessarily have occurred were it not for the actions of the aider. I therefore submit that it is preferable to focus on new threats.

This case is one way of capturing core elements of the motivating cases.<sup>94</sup> In this case there are competing claims between those who wish to be saved, and those who wish to avoid harm. Jennifer Rubenstein considers and rejects several different attempts to find principles appropriate to the *Aid Case*.<sup>95</sup> I shall consider each in turn before moving to Rubenstein's positive proposal.<sup>96</sup>

#### 5.3 Theorising the Aid Case

## 5.3.1 Complicity

According to the complicity view, "INGOs<sup>97</sup> are *complicit* in the injustices to which they knowingly contribute, even if they are furiously opposed to those injustices" (Rubenstein, 2015: 98). I think that the structure of the rescuers' intentions is important for assessing the level of complicity. Being "complicit in injustice" has the ring of sharing the aims of the injustice.

<sup>&</sup>lt;sup>94</sup> *Aid Case* is stylised to the extent that there are a simplifying assumptions about both the contributory factors involved and the numbers of individuals. Although this itself is a limitation of *Aid Case*, it does not detract from the overall analysis of permissibility and blameworthiness.

<sup>&</sup>lt;sup>95</sup> More particularly, Rubenstein is interested in "stay or go questions", where rescuers must decide whether to stay and continue providing aid or to leave. *Aid Case* is structurally similar to this predicament, absence the fact that rescuers are already providing aid. However, this does not affect the discussion of permissibility and blame. One may claim that it does make a difference because of the distinction between withdrawing and withholding aid. There may be a difference in the negative effects of aid provision in cases where aid is withdrawn (such as a vacuum in the provision of basic services), to when aid is merely withheld. I shall bracket a consideration of these negative effects, and focus instead on the consequences caused directly by both the aid workers and the villainous aggressor.
<sup>96</sup> Rubenstein considers the "do no harm" approach, but I omit a discussion of this in order to focus on the other approaches in more detail.

<sup>&</sup>lt;sup>97</sup> Rubenstein uses the term 'INGO' (International Non-Governmental Organisation), whereas I use the term 'rescuers'. These terms can be treated synonymously. Rubenstein calls the situation in question the "INGO Predicament", and outlines it as follows: "INGOs (1) have mostly good intentions, (2) some good effects, yet (3) contribute knowingly to injustices perpetrated primarily by others" (Rubenstein, 2015: 104). I take the INGO Predicament and the situation faced by rescuers as described in *Aid Case* to be sufficiently similar. In *Aid Case* the rescuers are motivated by the desire to aid others, they do indeed aid others, but a side-effect of their aiding is that there is a chance that a villainous aggressor will be enabled to harm others.

Rubenstein points out that this would wrongly characterise the intentions of INGOs negatively (Ibid: 99). I shall assume in this chapter that the rescuers do not share the aims of the evildoers. Rubenstein argues that a further weakness of this approach is that it leaves out of consideration the good that INGOs do. I think that Rubenstein is correct to point out the limitations of this approach, although the approach highlights the fact that although that INGOs aim to do good, they sometimes *knowingly* contribute to injustices committed by others.<sup>98</sup> This is a point worth highlighting, because *Aid Cases* deal with instances where the evil action is dependent on the beneficial action being performed.

Lepora and Goodin provide a framework for measuring complicity. On their view an agent is complicit in the wrongdoing of another if that agent makes a contribution to their wrongdoing. This is different from sharing the aims of the wrongdoer, or being a co-principal agent with the wrongdoer. The minimum conditions for an agent being complicit with the wrong another in her wrongful actions is that there is (a) a *contribution* to her wrongful actions, and (b) *knowledge* that you are contributing to her doing wrong (2013: 82). On this view, acting under the knowledge that one's causally contributory act will foreseeably contribute to the wrongful action is enough to make one complicit in that action.

According to Lepora and Goodin the level of pro tanto blameworthiness for a complicit act is a function of the following four variables: the moral badness of the principal wrongdoing, the responsibility for the contributory

<sup>&</sup>lt;sup>98</sup> There are different types of complicity that may be relevant. Firstly, it is possible for rescuers to merely be *aware* that if they provide aid then others will cause harm. Secondly, rescuers may be complicit in harmful actions by actively *collaborating* with the villainous aggressor.

act, the extent of the contribution, and the extent of a shared purpose with the principal wrongdoer (2013: 99).<sup>99</sup> It is left open how to precisely weigh and measure the variables. The higher the values of the variables, the higher the level of blameworthiness of the contributory agent for their act.

On this view, rescuers in *Aid Case* would be complicit in the wrongdoing of the villainous aggressor because their provision of aid is a causal contribution to the wrongdoing carried out by the aggressor, and they also foresee that this wrongdoing will occur as a result. Even though the rescuers do not share the purposes of the villainous aggressor, they would also possess a degree of pro tanto blameworthiness for their contribution to the wrongdoing. This is because the actions of the rescuers are voluntary, and there is knowledge both of their own causal contribution the wrongdoing and the severity of the harm that the aggressor carries out.<sup>100</sup> For now it is sufficient to signpost this view, as I argue in Section 5.4 that rescuers are *not* blameworthy either for their own actions or those of the villainous aggressor.

### 5.3.2 The Doctrine of Double Effect

The Doctrine of Double Effect (DDE) states that "the pursuit of some good tends to be less acceptable where a resulting harm is intended as a means

<sup>&</sup>lt;sup>99</sup> Not every variable has to be satisfied for an individual to be pro tanto blameworthy. Lepora and Goodin write that "not all complicit secondary agents necessarily share the wrongful purposes of the principal. [...] Sharing the wrongful purposes of the principle wrongdoer [...] is something *else* for which the complicit agent bears blame, in addition to whatever blame she bears for that causal contribution" (2013: 111-2).

<sup>&</sup>lt;sup>100</sup> More explicitly, Lepora and Goodin argue that an agent is complicit and bears more or less pro tanto blame for contributing to a principal wrongdoing if there is knowledge of their own contribution and of the magnitude of the wrongdoing the principal wrongdoer, as well as voluntary causal contribution to the wrongdoing (Ibid: 110).

than where it is merely foreseen" (Quinn, 1989: 335).<sup>101</sup> Part of the doctrine is a proportionality constraint – that the good end must be proportionate to the foreseen negative effect. Rubenstein argues that the DDE is helpful because of the focus on proportionality. It would prohibit cases where rescuers had good intentions, saved individuals, but a foreseen consequence was the death of many more individuals. Rubenstein outlines two problems with the DDE for assessing actions in *Aid Case* (Ibid: 100). First, most applications of the DDE are to cases where harm is direct but unintended, not cases contributing knowingly but unintentionally to harms perpetrated primarily by *others*.<sup>102</sup> According to Rubenstein, because the INGOs' contributions to harm are unintended side effects to achieving a good outcome "they therefore almost always fall on the 'possibly morally permitted' side of the DDE's main distinction" (Ibid).<sup>103</sup> It is not clear why this would be a problem, but perhaps Rubenstein has in the mind the possibility that so long as harm is not intended by rescuers, and the prospective foreseen harm does not outweigh the prospective good, then any intervention meeting these conditions would be judged permissible by the DDE. This is seen as problematic by Rubenstein, since the DDE would prescribe that INGOs ought to stay so long as the benefits outweigh the costs. As such, Rubenstein states that the DDE "does not offer much leverage" (Ibid.) on the Aid Case. Although it is unclear precisely what this means, I take Rubenstein to mean

<sup>102</sup> There is an interesting question here about the applicability of the DDE for cases where the foreseen harm of one agent's action is the intentional imposition of harm by another. If the agent's original action creates the circumstances which allow another agent to impose a harm does it make sense to call this resulting harm *foreseen* or is it instead *intended*? I address this question in Section 5.4 when I develop the *Moral Purity Account*. <sup>103</sup> It is not made explicit whether Rubenstein thinks that this is a problem, however she argues that the DDE is unhelpful for determining whether an INGO ought to "stay" in or "leave" conflict situations because the harmful actions that their aiding facilitates are side

<sup>&</sup>lt;sup>101</sup> It is difficult to give a unique definition of the DDE, since there are different formulations (Liao, 2012: 704). I settle for the Quinn (1989) definition here.

effects, and therefore permitted other things being equal.

that the DDE does not adequately apply due to the exclusive focus on permissibility.

The second main purported limitation to the DDE is that it does not provide any account of the "felt experience" of the moral conflicts faced by actors. The DDE does not remark on the "felt experience" of the individual acting, that rescuers may feel guilt or remorse for doing what they do because it facilitates harm to others. This is unsurprising, given that the DDE is about the moral permissibility of actions, rather than the appropriateness of particular emotional responses. Nonetheless, I maintain that this is an important desideratum for a plausible account that conceptualises Aid Case. A plausible account will explain not only the permissibility of intervening, but the reasonable moral emotions of those who are providing aid where harms are likely occur as a result of their intervention. Due to the two problems discussed above, Rubenstein takes the DDE to be unhelpful for theorising *Aid Case*, and as such attempts to look elsewhere. By contrast, I do not believe that Rubenstein's argument works due to the unconvincing critique of the DDE. I shall therefore appeal to the DDE in developing the Moral Purity Account in Section 5.4. First, however, I shall consider Rubenstein's alternative.

#### 5.3.3 Rubenstein's Account

Rubenstein offers a positive account called the "spattered hands" account. On this account, "INGOs sometimes have a responsibility to *allow their hands to be spattered* – that is, they sometimes have a responsibility to grudgingly accept contributing to injustices perpetrated primarily by others" (2015:

104).<sup>104</sup> The account is named "spattered hands" because the hands of the INGOs are only dirty *indirectly*, primarily through the actions of others. This is in contrast to the theory of "dirty hands" which focuses on the *direct* dirtying of one's hands when engaged with a morally testing scenario.

Rubenstein considers and rejects a "dirty hands" account of the situation as falling short at offering a plausible conceptualisation of *Aid Case*. Political rulers get their hands dirty when they do the right thing in utilitarian terms, but are also guilty of a moral wrong (Walzer, 1973: 161). For example, allowing the torture of a suspected terrorist in a ticking time-bomb scenario in order to save thousands of innocent people. Dirty hands captures the felt experience of rescuers, unlike the DDE. Rubenstein argues, however, that to say that rescuers have dirty hands is "melodramatic" (2015: 103). The negative effects of the actions of rescuers in *Aid Cases*, she claims, are typically less violent and less severe than dirty hands cases. Rubenstein notes another problem that the dirty hands approach faces: the harms that occur in *Aid Case* are typically *indirect*, unlike in traditional dirty hands cases where the harms typically perpetrated by the actor themselves. I think that this is an important point, since in *Aid Cases* it is the agency of a villainous aggressor that commits evil acts, and not the rescuers themselves.

Rubenstein's account suggests that an INGO may continue to provide aid to individuals so long as:

<sup>&</sup>lt;sup>104</sup> I shall bracket a discussion of whether and when rescuers have a *responsibility* to intervene and provide aid. It is not made clear how the permissibility of intervening comes apart from the responsibility to intervene, in Rubenstein's account. My focus is instead on when it is permissible to intervene, and who is blameworthy for the potential harms that may arise.

(a) The INGO's work does more good than harm, and (b) the resultant harms are only begrudgingly foreseen by the INGO (Ibid: 108).<sup>105</sup>

As it stands, the account is quite permissive. So long as more good than harm is done, and the harms are only begrudgingly foreseen, rescuers are permitted to provide aid. This account is structurally similar to the DDE, but with the addition of the claim that rescuers may sometimes "spatter their hands" by contributing to injustices caused primarily by others. Rubenstein suggests that it is not merely the case that rescuers have a permission to provide aid when the benefits of aiding outweigh the negative effects, but that rescuers may sometimes have a *duty* to get their hands spattered (2015: 106). According to this view, rescuers cannot discontinue the provision of aid in order to "keep their hands clean", when the benefits of their aid outweigh the costs. To this extent, Rubenstein's account remarks on the permissibility of intervening in *Aid Case*, as well as providing grounds for when it is obligatory for rescuers to provide aid. I bracket a discussion of when it is obligatory to provide aid, and my focus is instead on when it is permissible.

On the "spattered hands" account, rescuers' hands are spattered by the blood of others but "INGOs' good intentions do not shield them from *moral responsibility* for the predictable effects of their actions, even if those effects are indirect" (Rubenstein, 2015: 105, emphasis added). This is the core feature of the "spattered hands" account. Rubenstein does not provide an

<sup>&</sup>lt;sup>105</sup> To give further support to this, Rubenstein states that "INGOs should focus primarily on the likely consequences of their actions. Rather than simply "doing no harm", they should, as the DDE suggests, seek to implement a principle of "proportionality … if the negative effects of staying outweigh the benefits, INGOs have a duty to withdraw" (2015: 105-6).

account of moral responsibility to guide when precisely rescuers are morally responsible for the predictable effects of their actions. On one account, what is meant by moral responsibility is that a person X is responsible for a given action when "it is appropriate to take it as a basis of moral appraisal of that person" (Scanlon, 1998: 248). On this view, a person is morally responsible for X, if X can be properly attributed to her (Ibid: 277). Nothing is implied, on this view, about what an appropriate appraisal should be. I don't think that this is what Rubenstein has in mind, here. It would be trivially true that the rescuer's actions are a basis of moral appraisal. Because Rubenstein's account is centred on "spattered hands", where moral dirt "clings" to rescuers in some cases, I think that by moral responsibility Rubenstein has in mind *blameworthiness*. By way of a working definition, a person is morally responsible for X when he or she is blameworthy or praiseworthy for X.<sup>106</sup> To say that a person is liable to blame is to say that this person is liable to a negative emotional response, such as resentment or indignation (Rosen, 2004: 295-6).<sup>107</sup> Henceforth, I shall concentrate on blameworthiness.

Rubenstein argues that rescuers do not bear moral responsibility for all the things a villainous aggressor may do, but rather they are responsible for their own contributions (Ibid: 114). On this view, rescuers are pro tanto blameworthy for the negative effects of their own actions. This has two potential readings. Firstly, that the rescuers are responsible for the negative effects of their actions insofar as these actions contribute to the

<sup>&</sup>lt;sup>106</sup> It is beyond the scope of this chapter to outline and defend a particular view of moral responsibility against competing views.

<sup>&</sup>lt;sup>107</sup> T. M. Scanlon provides an alternative view of blame which states that to blame a person for an action, "is to take that action to indicate something about the person that impairs one's relationship in a way that reflects this impairment" (Scanlon, 2008: 122-3). For the purposes of this chapter, however, I shall assume that to be blameworthy is to be liable to moral sanctions, such as indignation or resentment.

circumstances which lead to the villainous aggressor harming others. Or, secondly, that the rescuers bear moral responsibility for a proportionate share of the negative effects from the villainous aggressor's actions.<sup>108</sup> The account I offer critiques the claim that rescuers are blameworthy in any degree for the negative foreseen effects of their morally-motivated, permissible aiding. I argue that although rescuers may be causally responsible for foreseeable harms, this does not entail that they are blameworthy for these harms.

It is worth considering the possibility that on Rubenstein's account, one can act permissibly and yet also be blameworthy for performing the action. Suppose that I perform an action, X, and foresee that Y is a direct consequence of this action. I will be blameworthy for Y if it turns out that Y is appropriately related to X; that is, that it was foreseeable that Y was a direct consequence of my performing X. For example, if I provide aid to a group of individuals then I am blameworthy or praiseworthy for that act, but if an individual then uses this act as a means for carrying out his own evil plan, then this is, I submit, something that I cannot be blameworthy for. If Y is a consequence of X that is foreseeably caused by a *separate* agent, then I cannot be blameworthy for Y.<sup>109</sup> If the above is correct, then serious doubt is cast on the interpretation of Rubenstein that rescuers are morally

<sup>&</sup>lt;sup>108</sup> A third potential reading is that cases such as *Aid Case* are tragic ones, such that a rescuer is blameworthy *whatever* she does. If she fails to act then she is blameworthy for the deaths of the 100 and if she does act, she is blameworthy for the deaths of the 20. But, as I will argue in the remainder of the chapter, rescuers are not blameworthy for consequences which they are not the direct cause of. I argue that a rescuer would *not* be blameworthy if they decided to save the 20.

<sup>&</sup>lt;sup>109</sup> Risk can be introduced if it is assumed that although the harmful actions of another are foreseeable, they may not be certain. The aggressor could decide not to cause harm, even though the rescuers predicted that they would. This could increase the permissibility of providing aid in *Aid Case* since there is the possibility that no harmful effects foreseeably arise from the provision of aid.

responsible in the sense of being answerable for their actions foreseen consequences. This is because rescuers cannot be answerable for the evil actions that originate from a separate agency, since they are not the rescuers actions. I defend this view in the next section.

Before outlining the *Moral Purity Account* I shall first outline some desiderata – in light of the examination of Rubenstein's approach – for an account to adequately explain the permissibility of acting in *Aid Case*. The desiderata are as follows:

- 1. The separate agency of both the rescuer and the villainous aggressor.
- 2. The fact that the rescuers have good intentions and merely foresee that harms will occur as a result of the actions of others.
- 3. The "felt experience" of the rescuers.

The *Moral Purity Account*, I submit, adequately meets the three desiderata and is superior to the "spattered hands" approach.

#### 5.4 The Moral Purity Account

The *Moral Purity Account* has two parts. The first part outlines when it is morally permissible for rescuers to intervene in *Aid Case*, and the second part outlines when rescuers in such cases are blameworthy for the harms caused by the agency of another. In the course of building the account, I shall consider four different permutations of *Aid Case*.

#### Consider the first:

*Infected Rescuer*: One hundred individuals are stranded on a large island, and are about to die from a serious illness unless a rescuer provides a treatment to each. Behind the large island is a small island with twenty healthy individuals. The rescuer has a choice of going on to the large island to provide medicines to each of the hundred, or to refrain. The rescuer carries a virus that is harmless to him and those on the large island, but lethal to those on the small island. As a foreseen but unintended consequence of the rescuer saving the one hundred, the virus will spread to the small island, killing all twenty inhabitants.<sup>110</sup>

Would it be morally permissible for the rescuer to save the one hundred individuals? I claim that it would be. The intention of the rescuer is to save the one hundred individuals from their fatal illness. The deaths of the twenty individuals are a foreseen side-effect of his aiding the one hundred.<sup>111</sup> There are competing interests in this case. The twenty do not want to be killed, and will be made worse off than they would have been

<sup>&</sup>lt;sup>110</sup> This case is inspired by UN peacekeepers in Haiti who inadvertently spread Cholera. This case is different in the respect that the peacekeepers' harmful effects were both unforeseen and arose from alleged negligence (through an improperly managed human waste disposal site). The case I am considering assumes that there is no negligence on the part of the rescuers.

<sup>&</sup>lt;sup>111</sup> I argue that although the number of people saved has to be greater than the number of people who are foreseeably killed, the number saved ought to be sufficiently greater. For instance, it would not be permissible to foreseeably kill 24 individuals in order to save 25, but it would be permissible (I assert) to foreseeably kill 20 in order to save 40. It is difficult to say precisely where the boundary lies between a permissible ratio of saved individuals to foreseeably killed individuals and an impermissible ratio. An explanation for this thought is that there are stronger moral constraints against harming people than there are against not aiding people, so more good action is required in order to outweigh the harm that is caused by the evildoer. However, a consequentialist would argue that there is no intrinsic difference between harming and not-aiding, so we could provide aid in cases where 100 people are saved and 99 people are harmed, since more good is done overall. Whereas, I maintain that the bar for justifying harming is higher than the bar for justifying not-aiding. If the rescuers provide aid when the benefits do not sufficiently outweigh the harms caused by the evildoer, then the rescuers act impermissibly, but they are not blameworthy for the deaths that are caused by the evildoer.

otherwise, were it not for the rescuer's intervention. The one hundred individuals will be made better off than they would have been before the intervention. The DDE can provide a principled explanation of why saving the one hundred and foreseeing that the twenty will die is permissible, where intentionally killing the twenty in order to save the one hundred would not be. The rescuer does not intend the deaths on the small island; they are a foreseen but unintended consequence of his intention to save the one hundred. Of course, one must also take into account the asymmetry between killing and letting die. It would be impermissible, for example, to save the hundred if one thereby foreseeably spread a virus that killed ninety others. But I submit that, in *Infected Rescuer*, the number of lives saved sufficiently outweighs the number of foreseen but unintended killings. For these reasons the rescuer acts permissibly.

Is the rescuer blameworthy for the deaths of the twenty? Although the rescuer's actions have *caused* the deaths of the twenty, I submit that he is not blameworthy for the deaths. I shall adopt Scanlon's account of blameworthiness to provide an explanation as to why the rescuer in *Infected Rescuer* is not blameworthy for the deaths of the twenty.<sup>112</sup> According to Scanlon, a person is blameworthy for an action when "his action indicates something about the agent's attitudes toward others that impairs his relations with them" (2008: 145). An individual is blameworthy when their conduct reveals something about them which indicates an impairment of

<sup>&</sup>lt;sup>112</sup> If one were to use an account of blameworthiness that held that one is blameworthy if and only if one acted impermissibly, then one will get the same result in *Infected Rescuer* as well as the further cases under discussion in this paper. To demonstrate: (1) for someone to be blameworthy they must have acted wrongly. (2) The rescuer in *Infected Rescuer* did not act wrongly because their action was permissible. (C) Therefore, the rescuer is not blameworthy. The view I am defending here is consistent with this account of blameworthiness.

their relations with others. This impairment makes it appropriate for others to blame them, to hold attitudes toward them that are different from those that constitute the default moral relationship (Scanlon, 2008: 141). To illustrate, suppose that Ann drives recklessly, but through sheer good luck, she injures no one. Those around her who were endangered by her lack of concern for their safety can revise their attitudes toward Ann, given her reckless attitude towards their well-being.

The rescuer does not act in such a way that his attitudes towards those whom his action affects are impaired by failing to show a concern for the welfare of others. His attitude towards the welfare of others is a positive one: concerned with saving people from disease. One may argue that because twenty people foreseeably died as a result of the rescuer's action, his relationship to the victims is impaired. It was his actions that lead to the foreseeable deaths. He may have to offer compensation and an expression of regret to those affected. If an individual's action is the cause of harm to another, then this individual may have a duty to provide compensation to the victim (Schroeder, 1997: 349). For example, if I accidently knock over someone's vase, it is appropriate for me to apologise to the owner and offer to pay compensation. Being the cause of a death may also warrant an apology, even though one is not blameworthy for the death (Scanlon, 2008: 150). Such an apology would not constitute a "genuine apology", but will rather reflect the fact that the agent's action has turned out badly.<sup>113</sup> However, there was nothing about the rescuer's attitudes towards others which showed a failure to have concern for the welfare of others, and

<sup>&</sup>lt;sup>113</sup> According to Luc Bovens, "for a genuine apology, it is not sufficient that the offender admit that her action turned out badly – she must also recognise her culpability" (2008: 221). To the extent that the rescuer in *Infected Rescuer* is not culpable for the deaths he is not liable for a genuine apology.

consequently it is not appropriate to hold the rescuer as blameworthy for the deaths of the twenty.

Consider now the following case, which is structurally equivalent to *Aid Case*:

*Villain-enabling Rescue*: In order to rescue one hundred people deserted on an island with rising water that will soon engulf them, rescuers can build a bridge with the sole intention of allowing them to leave. If the bridge were built, all one hundred people would escape, but as a foreseen consequence, the bridge would give a villainous aggressor access to a smaller island nearby that is connected to it by a small further walkway. The small island is protected from the rising water (but it is too small to support the hundred on the larger island, which is why they need to escape via the bridge). The villainous aggressor would use the bridge to kill the twenty people that are on the smaller island.

In this case I think that it would be permissible for the rescuers to build the bridge. The intention of the rescuers is merely to save the one hundred. As a foreseen and unintended side-effect of their aiding, a villainous aggressor harms twenty people. What is important about this case is that there is now the agency of two separate individuals. Because of this, it is not immediately clear that the DDE can explain why it is permissible to provide aid in this case. The reason is that the set-up of *Villain-enabling Rescue* is structurally different from most standard applications of the DDE, where the good effect

and the foreseen negative effect arise from the actions of a single agent.<sup>114</sup> One way of applying the DDE to *Villain-enabling Rescue* is to examine what answer it gives for each agent separately. For the rescuers, they intend to save one hundred, and they foresee that twenty will be killed by someone else. Rescuing would be permissible according to the DDE. But for the villainous aggressor, his actions are clearly impermissible – intending the deaths of twenty individuals for the goal of his own sadistic satisfaction. There are no foreseen but unintended consequences of the villainous aggressor's action, since he is only able to perform his action if the rescuers perform theirs. According to the DDE, it is impermissible to intend harm to a person. The aggressor's actions would also be impermissible because no good would be done to outweigh the harm.

The following question then arises. Is it permissible for rescuers to provide aid if it follows that as a foreseen but unintended consequence another agent *intentionally* carries out a harm to others? One answer would be that it would be impermissible to provide aid, because although the rescuer does not intend harm, their actions provide the circumstances through which it is possible for another to intentionally cause harm.

Another answer is to argue that it would in fact be permissible to intervene in *Villain-enabling Rescue* because the DDE should only apply to the actions of the rescuers. I think that this is the correct application since the action of

<sup>&</sup>lt;sup>114</sup> For example, in Quinn's (1989) *Terror Bomber* case, a pilot in a just war intentionally kills twenty civilians in order to bring forward the end of the war, thereby saving many more lives. In *Strategic Bomber* a pilot in a just war aims to destroy a munitions factory in order to bring forward the end of the war, but as a foreseen and unintended consequence, the bombs also lead to the deaths of twenty civilians. Even though the same number of deaths arise in both cases for the same good purpose, the *Terror Bomber* acts impermissibly, whereas the *Strategic Bomber* acts permissibly. In *Strategic Bomber*, the intended target and the foreseen but unintended harms both arise through the actions of the pilot alone.

the villainous aggressor would not occur at all if it were not for the actions of the rescuers. I think that this second approach is superior and that the harm being inflicted through the villainous agency of another or through one's introducing an unintended threat does not alter the permissibility of rescuers providing aid.

To further illustrate, consider the following case from Jeff McMahan (2002: 239-40). Suppose that one is not a police officer or other agent of the state, but merely a morally-motivated stranger and can either prevent a terrorist attack that will kill fifty innocent people, or prevent the accidental explosion of a gas main that will kill fifty people and cause minor injuries to a few others. Suppose that any further ill effects of these disasters would be equivalent. McMahan claims that "it would not be unreasonable to suppose that one ought to prevent the explosion of the gas main" (Ibid.) [because it prevents more harm]. This is because the prevention of wrongdoing doesn't matter much independently of the prevention of its consequences. One might argue in response that there is something uniquely bad about harmful consequences that are caused by wrongdoing, and as such it is preferable to prevent the terrorist attack. I assert, however, that from the perspective of an agent who performs an action that has a foreseen consequence *not directly* caused by their agency it does not make a difference whether the foreseen harm was caused by wrongdoing or by nature. There is a distinction to be drawn between bringing evil into the world that wasn't there previously, and bringing about a set of consequences from evil intentions. It is better not to bring evil into the world. But in the scenarios under discussion in *Aid Case,* no new evildoers are brought into existence. The harms are caused by evildoers, but this is not important for assessing the badness of the

consequences themselves. Compare, for example, a person's death as a result of a natural landslide versus the same death caused by a landslide that was exactly similar but was in fact caused by a villainous aggressor. In each case the consequences are equally bad.<sup>115</sup>

One question that may arise is whether it is more readily justifiable to intervene in *Villain-enabling Rescue* than *Infected Rescuer*, since in the former case the rescuer does not owe an apology and compensation for the harms that foreseeably arose from the agency of the villainous aggressor, whereas the rescuer ought to apologise in the latter case because of his special causal role in the deaths of innocents. These can be grounds for making an intervention in *Villain-enabling Rescue* more justifiable since the rescuers lack the burden of paying compensation or making an apology. Instead it is the villainous aggressor who bears the burden of moral sanctions through his culpability for deaths of the innocents. From the point of view of the rescuers there is greater reason to intervene in *Villain-enabling Rescue* than *Infected Rescuer*. I do not think that this is a troubling implication of the view. If two actions have the same consequences, but one action has greater moral sanctions on the individual than the other, then it is reasonable for the individual to prefer the latter action.

One may question whether an intervention is more readily justifiable when it does not lead to the burden of compensation or apology than it will be in an alternative when it does, if both interventions are permissible and blameless. Although saving the 100 in both cases, where doing so is permissible and blameless, is equally justified all-things considered, there

<sup>&</sup>lt;sup>115</sup> Quinn (1989: 347) makes a similar point regarding the symmetry of the badness of consequences arising from direct harmful agency and indirect harmful agency.

may exist an agent-relative reason to prefer acting in one scenario rather than the other. Having the option of two equally permissible and blameless interventions that each saves the same number of individuals, an individual may reasonably choose one over the other if there are fewer moral sanctions.

One may argue that in a case like *Villain-enabling Rescue*, the villain has a *primary* duty to compensate, but that this duty passes on to the rescuers if he were to die. I submit that in such a case the rescuers would have a duty to fulfil the compensatory duties, given their unique causal role in allowing the villain to carry out his deeds, along with other agents who had a contributory role in enabling the villain to cause harm. It might then be argued that an intervention will become less justifiable if it were to be known that the villain will die from a heart attack after his villainous deeds. In response, both interventions would be equally justified all-things considered because each would be equally permissible and blameless, but the rescuers would have more reason to intervene in the case where the villain aggressor does not have a heart attack.

In order to now consider how the role of the separateness of the agency of the rescuers and the villainous aggressor plays a role in identifying blameworthiness for the deaths of innocents, I shall consider an example from Kamm. The example performs two functions for my argument. Firstly, it highlights how the separateness of agency can make a difference to blameworthiness. Secondly, Kamm's example provides a case where one is not blameworthy for a harm that one has personally committed since one acted at the request of a villainous aggressor. I argue that if it is plausible to think that one is not blameworthy in this case, then this provides support for

the weaker claim I defend that one is not blameworthy for purely good acts that one performs which foreseeably lead another person to cause harm.

Kamm argues that in some cases there are positive reasons to collaborate with evil (2007: 309). Kamm examines two types of cases; one where it is possible to harm others at the request of a villain, and another where there is the opportunity to harm others not at the request of an evil villain.<sup>116</sup> Kamm argues that we are "often more willing to recommend an agent harming others in the first case than in the second" (Ibid). To illustrate this point, Kamm presents the following two cases:<sup>117</sup>

> *Offer*: The Captain is about to execute twenty randomly selected villagers. The Captain offers Jim to shoot one of the villagers. If Jim accepts the offer, then the Captain will let the other nineteen villagers go. If he refuses, all twenty will be shot by the Captain.

> *Scan*: The same scenario as *Offer*, but now Jim is hiding in the bushes and is in possession of an infallible brainscanning device that tells him that the Captain will kill all twenty unless Jim kills one. If Jim does kill the one, then the Captain will release the nineteen others, and if Jim does not kill the one, then the Captain will kill all twenty.<sup>118</sup>

<sup>&</sup>lt;sup>116</sup> It is possible for Jim to refuse the offer, but this would lead to all the villagers being killed by the Captain.

<sup>&</sup>lt;sup>117</sup> These cases are adapted from Bernard Williams (1973).

<sup>&</sup>lt;sup>118</sup> A less abstract example can also be constructed. For example, a judge in a town has the option of scapegoating one person in order to stop an angry mob from killing a greater number. The scapegoat would have been killed by the mob anyway. The mob calls for the one to be scapegoated, and if he is not, then they will kill the greater number. This is

Kamm introduces the notion of differential responsibility to explain why it would be more permissible for Jim to shoot in Offer than in Scan. If Jim shoots in Offer then the Captain will be solely blameworthy for the death of the villager.<sup>119</sup> If Jim shoots in *Scan* then, she claims, he will be partly blameworthy for the death of the villager.<sup>120</sup> What makes the Captain solely blameworthy for the negative consequences of Jim's action in Offer is that the Captain created the evil scheme whereby the contingency of the better outcome (one villager dying rather than all twenty) is reliant on Jim's killing. If the Captain did not make the offer, or create the unjust state of affairs, then no villagers would have died. The Captain's act in Offer is "the initiation of a lethal plot" (Kamm, 2007: 312), and "we lay the negative consequences of the act at the Captain's doorstep and not just negatively, that is, because he allowed Jim to kill, but positively, as something he brought about because he made the offer" (Ibid: 311). Although the structure of these cases differs somewhat from the Aid Case, by virtue of the fact that rescuers do not directly kill in order to save more from harm and because the person Jim kills would have been killed by the Captain anyway, there is a valuable insight for the distribution of blameworthiness in Aid Case.

Returning to *Villain-enabling Rescue*, it is important in these cases that the aims of the rescuers and the villainous aggressor are *separate*. There is no explicit collaboration, and the rescuers do not *intend* harm (unlike Jim in

sufficiently similar to *Offer*. An approximation of *Scan* can be made if one now assumes that the judge *knows* that this is what the mob are planning to do, so he can choose to scapegoat the one because he knows that they plan to kill a greater number if he doesn't. <sup>119</sup> Kamm argues that this would be true "even though Jim kills someone and the Captain winds up killing no one" (2007: 311).

<sup>&</sup>lt;sup>120</sup> Kamm specifically has in the mind the idea of *moral responsibility* and *moral blameworthiness*. Although there is a sense in which Jim would be *causally* responsible for the death of the one if he chooses to shoot with respect to the fact that the negative consequences are attributable to his agency (Rosen, 2010: 684).

*Offer*). The evil plan of the villainous aggressor exists separately from the agency of the rescuers plan to save the one hundred people. The villainous aggressor uses the efforts of the rescuers as a means to carry out his evil project. Why should the rescuers be blameworthy for the actions that result from the evil plan from the agency of another? One might think that the rescuers bear some moral responsibility for the bad effects of their action of building a bridge because the villainous aggressor would not have killed the twenty people had they not built it. In response, I think that the fact that an intention to aid others is *used* by another to carry out a distinct and evil plan to harm others does not make those who aid others morally responsible for the harms resulting from this evil plan.

This can be developed further by seeing how the separateness of the aims of the rescuers and the villainous aggressor can have implications for liability to blame. According to Elinor Mason (forthcoming) one can only be morally responsible for those actions that one has a "right relationship" with. For example, if one performs an action whilst being coerced their action is "not [their] own in the right way for him to be responsible for it" (Ibid.). Mason illustrates this with *Offer*. The captain dominates Jim's will, and as such Jim is not responsible for his action. He doesn't identify it as *his* action. Something similar can be said in *Aid Case*. The rescuers cannot identify with the action of the villainous aggressor as *their* action. This might seem obvious, since the actions of the villainous aggressor are *his* actions. But what is crucial is that, furthermore, the rescuers do not share the aims and desires of the aggressor. If the rescuers *did* share the aims and desires of the aggressor, and wished for him to carry out the harms as an indirect

consequence of their rescuer, then this would be grounds for thinking that they identify with the actions of the aggressor in the right way.

One might argue that if one foresees particular consequences from one's action then one relates in the right way, to some degree, to the consequences. This might be the case when there are direct foreseeable consequences that flow from one's actions. But when there are foreseeable *indirect* consequences, it is difficult to see how mere foreseeability can entail that one is partly responsible for the act if one both did not share the aims of the act nor carry out the act. The foreseen harm that occurs in *Villain-enabling Rescue* is not only causally distant from the action of the rescuers, but if no other agent had intervened there would be no harmful side effect. This is in contrast to *Infected Rescuer* where the rescuer's action of saving was also the direct cause of the foreseen harm to the islanders. This lack of a direct causal connection between the rescuers aiding and the harm caused by the villainous aggressor is what absolves the rescuers from blame in Villain*enabling Rescue*. After all, even in *Infected Rescuer*, the rescuer's causal connection to the foreseen harm of the twenty was enough to make the rescuer liable for an apology or for compensation, but not enough to make him blameworthy for the deaths.

To summarise, I have argued that in *Infected Rescuer* the rescuer is not blameworthy for the deaths of the twenty because he acted permissibly. He owes an apology and/or compensation to the victim in virtue of being the cause of the harm. Due to the differences between *Villain-enabling Rescue* and *Infected Rescuer* – that the rescuers were not the cause of the deaths in *Villainenabling Rescue*, and that it was another agent who was responsible for the negative effects – the rescuers are no more blameworthy in *Villain-enabling* 

*Rescue* than they are in *Infected Rescuer*. Therefore, rescuers are not blameworthy in *Villain-enabling Rescue*. To connect this conclusion to the dialectic of the paper, it is important to note that *Villain-enabling Rescue* is structurally equivalent to *Aid Case*. Given that this is the case, I have demonstrated that (other things being equal) rescuers in *Aid Case* are not blameworthy for the harms that arise through the agency of another. The first part of the *Moral Purity Account* utilises the applied version of the DDE to state when it would be permissible for rescuers to intervene in *Aid Case*. The second part of the account is comprised of the account of blameworthiness previously discussed. The account of permissibility utilises the DDE. This is also implicitly used by the "spattered hands" account to determine when it is permissible to intervene in *Aid Case* (Rubenstein, 2015: 108).<sup>121</sup>

To further illustrate the view of permissibility outlined in the *Moral Purity Account*, consider the following case:

*Tsunami Terror*: There are one hundred individuals on the larger island. A tsunami will shortly crash into the island and kill all one hundred. They could leave and save themselves, but due to a reasonable distrust of outsiders, the one hundred do not believe the warnings about the tsunami and refuse to evacuate the island. As a means of getting the hundred to leave, the rescuers could kill the twenty on the smaller island (who would be safe from the

<sup>&</sup>lt;sup>121</sup> Further evidence for this can be found when Rubenstein says that "the spattered hands framework suggests, first, that in addressing this question, INGOs should focus primarily on the likely consequences of their actions. Rather than simply "doing no harm," they should, as the DDE suggests, seek to implement a principle of "proportionality." (2015: 105).

tsunami) as a means of terrorising the one hundred into evacuating the larger island.

The account claims that it would be impermissible for the rescuers to save the hundred by killing the twenty in this case. It is impermissible to intentionally kill the twenty as a means of saving the one hundred. A further explanation for the impermissibility of killing the twenty is that if one does so, one is thereby employing *opportunistic agency*. This is the *use* of others to further our own ends, and is something that the DDE "strongly discriminates against" (Quinn, 1989: 344). Further, on the *Moral Purity Account* the rescuers would be blameworthy for the deaths of the twenty. They intentionally cause harm to individuals, and the harmful effects flow from their agency. This is in contrast to *Villain-enabling Rescue*, where the intentional harm flowed from the agency of the villainous aggressor. This makes a difference to the blameworthiness that can be ascribed to rescuers in each case.

Consider now the following variation:

*Tsunami Terror-enabling Case*: This case has a similar structure to *Tsunami Terror*. The rescuers can either refrain from saving the one hundred, or they can intend to build a bridge so as to facilitate the evacuation of the island. The rescuers are unable to persuade the islanders to leave. Knowing this, the bridge is built by the rescuers *because* they foresee that a villainous aggressor will use the bridge

to kill twenty on the small island, which then spreads terror in the large island, and persuades them to leave.<sup>122</sup>

Here the bridge is built *knowing* that the villainous aggressor will use it, and only because he will use it. Building the bridge *because* it will lead to the twenty being killed by the villainous aggressor does not imply that one *intends* the death of the twenty (Kamm, 2007: 95). To illustrate, the bridge isn't built *in order* for the villainous aggressor to kill the twenty. The bridge is built *because* it is known that the villainous aggressor will kill the twenty, thereby facilitating the evacuation of the one hundred. The rescuers would not build the bridge unless it was expected that the villainous aggressor would kill the twenty and facilitate the escape of the one hundred. For it is consistent with building it only *because* they foresee that the aggressor will come that the rescuers wouldn't, if the rescuer would not otherwise come without encouragement, encourage or induce the villainous aggressor over to the smaller island, knowing that he will kill twenty.

I believe that because the deaths of the twenty are not intended by the rescuers, it is permissible for them to save the one hundred. This is true, even though the bridge is built only because the villainous aggressor kills the twenty. On the *Moral Purity Account*, the rescuers would not be blameworthy for the deaths. The deaths are intended and caused by the agency of another, not by the rescuers.<sup>123</sup>

<sup>&</sup>lt;sup>122</sup> This example shares a similar structure to Kamm's "Party Case" (2007: 95).

<sup>&</sup>lt;sup>123</sup> The *Moral Purity Account* is therefore also consistent with what is called the Doctrine of Triple Effect (DTE): "The DTE states: 'A greater good that we cause and whose expected existence is a condition of our action, but which we do not necessarily intend, may justify a lesser evil that we must not intend but may have as a condition of action" (Kamm, 2007: 118).

It is worth considering how different types of behaviour may be collaboration with evil and therefore impermissible. In *Somalia Famine* rescuers had to pay militants in order to access those who required aid. When one pays-off evil-doers, one is giving a direct financial contribution to those who are likely to use the resources to continue carrying out harm. To see where the *Moral Purity Account* stands on such cases, consider the following modification to *Aid Case*:

> *Ransom Payment Rescue*: There are one hundred individuals held hostage on an island who will soon die if they are not given life-saving medicines. The island is guarded by a villainous aggressor, who will grant access to the island for a ransom payment. If the aggressor receives the ransom payment he will then be enabled to kill twenty separate individuals. Rescuers can decide whether to pay the ransom or to refrain.

If the rescuers refrain from paying the ransom the one hundred will die, and the twenty others will live. If the rescuers pay the ransom, then they will save the one hundred, but will also be providing the villainous aggressor with the means to kill the twenty. Would it be permissible to pay the ransom? One answer is that it would be impermissible because the ransom payment takes the form of bribery, and bribery is a *mala in se* – a wrong in itself. Another answer is to suggest that the payment is permissible because the rescuers do not intend the deaths of the twenty since their intention is to pay-off the aggressor in order to access the one hundred. If the payment is interpreted as a bribe, then I argue it would be impermissible for the rescuers to save the one hundred. This is because the act *itself* would be

wrong. But if instead the payment is viewed as a mere *enabling* of the saving of the one hundred then it would be permissible to pay the ransom. This is because *Ransom Payment Rescue* will have the same structure as *Villain-Enabling Rescue*, in virtue of the payment both functioning as a means to save the one hundred, and as an enabler of the villainous aggressor's plan to kill the twenty.

#### 5.5 Rubenstein and Responsibility

The *Moral Purity Account* differs from Rubenstein's "spattered hands" account because it claims that rescuers are not blameworthy for the negative effects of aiding others if their intentions are *solely* to save others, and that any negative effects come from the evil projects of another. The "spattered hands" account claims that rescuers would have *some* blameworthiness because they foresaw that their actions would indirectly lead to some causing harm to others. A virtue of the "spattered hands" account is that it takes the felt experience of the rescuers into consideration. To recall, the account states that the hands of rescuers become:

"Dirty through more indirect means; they are spattered by dirt created primarily by others. Nonetheless dirt still clings to them: INGO's good intentions do not shield them from moral responsibility for the predictable effects of their actions, even if those effects are indirect" (Rubenstein, 2015: 105).

I have claimed that rescuers are not blameworthy for the side effects of their actions caused by a villainous aggressor. Further, because the harms occurring as a side effect arise from the agency of another, the rescuers do

not compromise their integrity as individuals. According to Bernard Williams (1973), for an individual to act with integrity is for the individual's actions and decisions to flow from his projects and attitudes which he is most closely identified (116-7). The integrity of an individual is compromised when obliged to act in a particular way, as determined by the projects and plans of another, or of an impartial viewpoint. For example, in the *Offer* case, above, Jim's integrity is compromised if he shoots villager without previously wanting to. A utilitarian would suggest that the correct thing to do would be to shoot the villager, for this increases the sum total of utility. However, doing so will compromise the integrity of Jim. Nothing of this sort occurs in *Aid Case*.<sup>124</sup>

My preferred view is that rescuers may suffer from "agent regret" in *Aid Case*. Agent regret is a variety of regret whereby an agent themselves is causally responsible for an unintended harm that flows from their agency. For example, a conscientious lorry driver who runs over a small child who skips in front of the lorry before the driver has the opportunity to stop or swerve feels a distinct type of regret, because it was *he* who made the death happen, even though it was not intentional. This is not true of a spectator to the incident, who may legitimately feel regret about what has happened, but not agent regret.<sup>125</sup>

Agent regret captures the felt experience of individuals whose good actions lead to another exploiting those actions in order to cause harm. Rescuers

<sup>&</sup>lt;sup>124</sup> It is not clear that the integrity of the rescuers is compromised when they carry out their benevolent projects and attitudes. If it were the case that the rescuers had to intentionally cause harm to some in order to save others, this may be a cause for compromised integrity. <sup>125</sup> I shall leave aside the question of whether individuals *ought* to feel agent regret in such cases. Instead, I present this account as a possibility that is consistent with the *Moral Purity Account*.

may regret that *their* actions have been part of a causal chain that led to harm being caused.<sup>126</sup> But this does not mean that moral dirt clings to them in the sense that they have a share of the blameworthiness for the harms that are caused, but instead that the rescuers wish that it was not them that made the conditions appropriate for a villainous aggressor to carry out their evil projects. Further, feeling agent regret does not entail that the agent is in anyway blameworthy for the harm that may inadvertently flow from their agency. This is consistent with the *Moral Purity* account. This is a more plausible account of the felt experience of the rescuers in *Aid Case*.

The view I am developing can allow for the alteration of the moral relationship between the victim and the "perpetrator" even when blame is inappropriate. To illustrate, recall the lorry driver case. It may be appropriate for the child's parents to not want to say hello to the lorry driver in the morning, let him look after their other children, or rely on him to fulfil a particular duty (Scanlon, 2008: 148). These forms of stigma or negative attitudes towards those who are causally responsible for a state of affairs may be appropriate, even though these individuals are not blameworthy for the harms that they caused. In *Aid Case* it may be appropriate for those who bear the harms of the villainous aggressor to modify their moral relationship with the rescuers whose actions *indirectly* caused harms through the agency of another. For example, individuals may become less trustworthy of rescuers if it is known that their aiding led to the circumstances in which a villainous aggressor killed innocents.

<sup>&</sup>lt;sup>126</sup> Bernard Williams argues that "the sentiment of agent regret is by no means restricted to *voluntary agency*. It can extend far beyond what one intentionally did to almost anything for which one was causally responsible in virtue of something one intentionally did" (1993: 43). To this extent, it is possible for an agent to feel regret about a harm caused by a villainous aggressor, if it was the case that this agent's act was *causally responsible* for the circumstances that allowed for the villainous aggressor to cause harm.

To this extent, the *Moral Purity Account* provides a plausible account of the three criteria listed in Section 5.3.3 It adequately captures the fact that the rescuers only intend for the one hundred to be saved, and foresee that the agency of another will kill the twenty. It is also sensitive to the separateness of the agency of the rescuers and the villainous aggressor; that the rescuers only intend to carry out a plan that saves individuals, whereas the villainous aggressor intends only to kill innocent people. The *Moral Purity* account is also sensitive to the "felt experience" of the rescuers in *Aid Case*.

#### 5.6 Conclusion

In this chapter, I have argued that the *Moral Purity Account* provides a more plausible conceptualisation of *Aid Case* than the "spattered hands" approach. The "spattered hands" approach argues that sometimes rescuers share a degree of blameworthiness for the harms caused by another. I have argued that this is not the case. The *Moral Purity Account* also provides an independently plausible explanation for both the separateness of the agency of the rescuers and the villainous aggressor, along with the "felt experience" of rescuers acting in *Aid Case*.

# Bibliography

Al-Najjar, Nabil I., and Weinstein, Jonathan, 2009, "The Ambiguity Aversion Literature: A Critical Assessment", *Economics and Philosophy*, Vol. 25, pp. 249-84.

Arneson, Richard, 1997. 'Postscript to "Equality and Equal Opportunity for Welfare"', in Equality: Selected Readings, ed. Louis Pojman and Robert Westmoreland (Oxford: Oxford University Press), pp. 238-41.

Ashford, Elizabeth, 2003, "The Demandingness of Scanlon's Contractualism", *Ethics*, Vol. 113, No. 2, pp. 273-302.

Binet, Laurence, 2016, "Rwandan Refugee Camps in Zaire and Tanzania 1994-1995", *Medicins San Frontieres*.

Binmore, Ken (2009). *Rational Decisions*. Princeton: Princeton University Press.

Binmore, Ken, & Stewart, Lisa, & Voorhoeve, Alex, 2012, "How Much Ambiguity Aversion? Finding Indifferences Between Ellsberg's Risky and Ambiguous Bets", *Journal of Risk & Uncertainty*, Vol. 45, pp. 215-38.

Bovens, Luc, 2008, "Apologies", *Proceedings of the Aristotelian Society*, Vol. CVIII, Part 3, pp. 219-239.

Bovens, Luc, 2015, 'Concerns for the Poorly Off in Ordering Risky Prospects', *Economics and Philosophy*, Vol. 31, pp. 397-429.

Bradley, Richard and Drechsler, Mareile, 2014, "Types of Uncertainty", *Erkenntnis*, Vol. 79, Issue 6, pp. 1225-1248.

Broome, John, 1984, "Uncertainty and Fairness", *The Economic Journal*, Vol 94. Issue 375, pp. 624-632.

Broome, John, 1998, "Kamm on Fairness", *Philosophy and Phenomenological Research*, Vol. 58, No. 4, pp. 955-61.

Broome, John, 1990, "Fairness", *Proceedings of the Aristotelian Society*, Vol. 91, pp. 87-101.

Broome, John, 1991, *Weighing Goods*: *Equality, Uncertainty and Time*, Blackwell: Oxford.

Broome, John, 2015, "Equality versus Priority: A Useful Distinction", *Economics and Philosophy*, Vol. 31, Issue 2, pp. 219-228.

Crisp, Roger, 2011, "In Defence of the Priority View: A Response to Otsuka and Voorhoeve", *Utilitas*, Vol. 23, No. 1, pp. 105-108.

Darwall, Stephen, 1977, "Two Kinds of Respect", *Ethics*, Vol. 88, No. 1, pp. 36-49.

Diamond, Peter. 1967, "Cardinal Welfare, Individualistic Ethics, and Interpersonal Comparison of Utility: Comment", *Journal of Political Economy*, 75: 765-6.

Dillon, Robin S., "Respect", *The Stanford Encyclopedia of Philosophy* (Fall 2015 Edition), Edward N. Zalta (ed.), URL = <http://plato.stanford.edu/archives/fall2015/entries/respect/>.

Ellsberg, Daniel, 1961, "Risk, Ambiguity, and the Savage Axioms", *The Quarterly Journal of Economics*, Vol. 75, No. 4, pp. 643-69.

Feinberg, Joel, 1987, *The Moral Limits of the Criminal Law Volume 1: Harm to Others*, Oxford University Press: Oxford.

Finkelstein, Claire, 2003, "Is Risk a Harm?", 151, University of Pennsylvania Law Review, pp. 963-1001.

Frick, Johann, 2013, "Uncertainty and Justifiability to Each Person", in *Inequalities in Health: Concepts, Measures, and Ethics*, eds. Nir Eyal, Samia A. Hurst, Ole F. Norheim, and Dan Wikler, Oxford University Press: Oxford, pp. 130-145.

Frick, Johann, 2015, "Contractualism and Social Risk", *Philosophy and Public Affairs*, Vol. 43., No. 3., pp. 175-223.

Fried, Barbara, 2012, "Can Contractualism Save Us from Aggregation?", *Journal of Ethics*, Vol. 16, pp. 39-66.

Gilboa, Itzhak, & Schmeidler, David, 1989, "Maximin Expected Utility with Non-Unique Prior", *Journal of Mathematical Economics*, 18, 141–153.

Gilboa, Itzhak, & Postlewaite, Andrew & Schmeidler, David, 2009, "Is it Always Rational to Satisfy Savage's Axioms?", *Economics and Philosophy*, Vol. 25, pp. 285-96. Greaves, Hilary, 2015, "Antiprioritarianism", *Utilitas*, Vol. 27, Issue 1, pp. 1-42.

Griffin, James, 1988, *Well-Being: Its Meaning, Measurement and Moral Importance*, Oxford University Press: Oxford.

Hansson, Sven Ove, 2003, "Ethical Criteria of Risk Acceptance", *Erkenntnis*, Vol. 59, Issue 3, pp. 291-309.

Hare, Caspar, 2012, "Obligations to Merely Statistical People", *The Journal of Philosophy*, Vol. 109, Issue 5-6, pp. 378-390.

Hayenhjelm, Madeleine, & Jonathan Wolff, 2011, "The Moral Problem of Risk Impositions: A Survey of the Literature", *European Journal of Philosophy*, Vol. 20, Issue S1, pp. 26-51.

Holm, Sune, 2016, "A Right against Risk-Imposition and the Problem of Paralysis", *Ethical Theory and Moral Practice*, Vol. 19, Issue 4, pp. 917-30.

Hooker, Brad, 2005, "Fairness", *Ethical Theory and Moral Practice*, Vol. 8, pp. 329-352.

Hurwicz, L., 1951, *Optimality criteria for decision making under ignorance* (Vol. 370). Cowles Commission Discussion Paper, Statistics.

Hyams, Keith, 2015, "Hypothetical Choice, Egalitarianism and the Separateness of Persons", *Utilitas*, Vol. 27, No. 2, pp. 217-39.

IPCC, 2014a, "Climate Change 2014: Impacts, Adaptation, and Vulnerability. Part B: Regional Aspects. Contribution of Working Group II to the Fifth Assessment Report of the Intergovernmental Panel on Climate Change" [Barros, V. R., C. B. Field, D. J. Dokken, M. D. Mastrandrea, K. J. Mach, T. E. Bilir, M. Chatterjee, K. L. Ebi, Y. O. Estrada, R. C. Genova, B. Girma, E. S. Kissel, A. N. Levy, S MacCracken, P. R. Mastrandrea, and LL. White (eds.)]. Cambridge University Press, Cambridge, United Kingdom and New York, NY. USA, 688 pp.

IPCC, 2014b, "Summary for Policymakers", In: *Climate Change 2014: Impacts, Adaptation, and Vulnerability. Part A: Global and Sectoral Aspects. Contribution of Working Group II to the Fifth Assessment Report of the Intergovernmental Panel on Climate Change, C. B. Field et al. eds.* Cambridge University Press, Cambridge, United Kingdom and New York, NY, USA, pp. 1-32.

James, Aaron, 2012, "Contractualism's (not so) Slippery Slope", *Legal Theory*, Vol. 18, Issue 3, pp. 263-292.

Kagan, Shelly, 1994, "Me and My Life", *Proceedings of the Aristotelian Society*, Vol. 94, pp. 309-324.

Kamm, Frances, 1993, *Morality*, *Mortality*, Vol. I, Oxford University Press: Oxford.

Kamm, F. M., 2007, *Intricate Ethics: Rights, Responsibilities, and Permissible Harm*, Oxford University Press: Oxford.

Kirkpatrick, James & Nick Eastwood, 2015, "Broome's Theory of Fairness and the Problem of Quantifying the Strength of Claims", *Utilitas*, 27:1, pp. 82-91.

Kornhauser, Lewis and Lawrence G. Sager, 1988, "Just Lotteries", *Rationality and Society*, Vol. 27, No. 4, pp. 483-516.

Kumar, Rahul, 2003, "Who Can Be Wronged?", *Philosophy and Public Affairs*, 31, No. 2, pp. 99-118.

Kumar, Rahul, 2015, "Risking and Wronging", *Philosophy & Public Affairs*, Vol. 43, No. 1, pp. 27-51.

Lange, Benjamin, 2016, "Restricted Prioritarianism or Competing Claims?', *Utilitas*, pp. 1-16.

Lazar, Seth, 2017, "Risky Killing: How Risks Worsen Violations of Objective Rights", *Journal of Moral Philosophy*, online first, pp. 1-26.

Lazenby, Hugh, 2014, "Broome on Fairness and Lotteries", *Utilitas*, Vol. 26, No. 4, pp. 331-345.

Lenman, James, 2008, "Contractualism and Risk Imposition", *Politics*, *Philosophy & Economics*, Vol. 7, No. 1, pp. 99-122.

Lepora, Chiara, and Goodin, Robert E, 2013, *On Complicity and Compromise*, Oxford University Press: Oxford.

Mason, Elinor, forthcoming, "Consequentialism and Moral Responsibility", in *Consequentialism: New Directions, New Problems?*, Christian Seidel ed., Oxford University Press. McCarthy, David, 1997, "Rights, Explanation, and Risks", *Ethics*, Vol. 107, No. 2, pp. 205-225.

McCarthy, David, 2008, "Utilitarianism and Prioritarianism II". *Economics and Philosophy*, Vol. 24, Issue 1, pp. 1-33.

McCarthy, David, 2015. "Distributive Equality", Mind 124 (496): 1045-1109.

McCarthy, David, forthcoming, "The Priority View", *Economics and Philosophy*.

Munoz-Darde, Veronique, 2013, "Global Justice: Imposed and Shared Risks", in *Spheres of Global Justice*: *Volume 2*, ed. Jean-Cristophe Merle, Springer Science & Business Media, pp. 539-554.

Nagel, Thomas, 1979, *Mortal Questions*, Cambridge University Press: Cambridge.

Nagel, Thomas, 1991, *Equality and Partiality*, Oxford University Press: Oxford.

Nissan-Rozen, Ittay, 2017, "How to be an Ex-Post Egalitarian and an Ex-Ante Paretian", *Analysis*, online first, pp. 1-10.

Nozick, Robert, 1974, Anarchy, State, and Utopia, Blackwell: Oxford.

Oberdiek, John, 2009, "Towards a Right against Risking", *Law and Philosophy*, Vol. 28, No. 4, pp. 367-92.

Oberdiek, John, 2012, "The Moral Significance of Risking", *Legal Theory*, 18, pp. 339-56.

Oberdiek, John, ms., *Imposing Risk: A Normative Framework*, Oxford University Press: Oxford.

O'Neill, Martin, 2012, "Priority, Preference and Value", *Utilitas*, Vol. 24, Issue 3, pp. 332-348.

Otsuka, Michael, and Voorhoeve, Alex, 2009, "Why It Matters That Some Are Worse Off Than Others: An Argument Against the Priority View", *Philosophy & Public Affairs*, Vol. 37, No. 2, pp. 171-199.

Otsuka, Michael, and Voorhoeve, Alex, 2011, "Reply to Crisp", Vol. 23, No. 1, pp. 109-114.
Otsuka, Michael, 2006, "Saving Lives, Moral Theory, and the Claims of Individuals", *Philosophy & Public Affairs*, Vol. 34, No. 2, pp. 109-135.

Otsuka, Michael, 2012, "Prioritarianism and the Separateness of Persons", *Utilitas*, Vol. 24, Issue 3, pp. 365-380.

Otsuka, Michael, 2015, "Prioritarianism and the Measure of Utility", *The Journal of Political Philosophy*, Vol. 23, No. 1, pp. 1-22.

Otsuka, Michael, 2015, "Risking Life and Limb: How to Discount Harms by their Improbability", in I. Glenn Cohen, Norman Daniels, and Nir Eyal, eds., *Identified versus Statistical Lives: An Interdisciplinary Perspective*, Oxford University Press: Oxford.

Otsuka, Michael, ms., "Determinism and the Value and Fairness of Equal Chances".

Parfit, Derek, 1984, Reasons and Persons, Oxford University Press: Oxford.

Parfit, Derek, 2002, "Equality or Priority?", in *The Ideal of Equality*, ed. Matthew Clayton and Andrew Williams, Basingstoke, pp. 81-125.

Parfit, Derek, 2011, *On What Matters: Volume One*, Oxford University Press: Oxford.

Parfit, Derek, 2012, "Another Defence of the Priority View", *Utilitas*, Vol. 24, pp. 399-440.

Perry, Stephen, 1995, "Risk, Harm, and Responsibility" in *Philosophical Foundations of Tort Law*, David G. Owen ed., Oxford University Press: Oxford.

Perry, Stephen, 2001, "Responsibility for Outcomes, Risk, and the Law of Torts", in Gerald Postema (ed.), *Philosophy and the Law of Torts*, Cambridge University Press: Cambridge.

Perry, Stephen, 2003, "Harm, History, and Counterfactuals", *San Diego Law Review*, 40, pp. 1283-1314.

Piller, Christian, forthcoming, "Treating Broome Fairly", Utilitas, pp. 1-25.

Placani, Adriana, forthcoming, "When the Risk of Harm Harms", Law and Philosophy.

Quinn, Warren, 1989, "Actions, Intentions, and Consequences: The Doctrine of Double Effect", *Philosophy & Public Affairs*, Vol. 18, No. 4, pp. 334-351.

Rachels, James, 2001, "Killing and Letting Die", in Lawrence Becker and Charlotte Becker (eds.), *Encyclopedia of Ethics, Volume 2, 2<sup>nd</sup> Edition*, Routledge: London, pp. 947-950.

Rawls, John, 1999, *A Theory of Justice*: *Revised Edition*, Harvard University Press: Cambridge, Mass..

Raz, Joseph, 1988, The Morality of Freedom, Oxford University Press: Oxford.

Reibetanz, Sophia, 1998, "Contractualism and Aggregation", *Ethics*, Vol. 108, pp. 296-311.

Rendall, Matthew, 2013, "Priority and Desert", *Ethical Theory and Moral Practice*, Vol. 16, No. 5, pp. 939-951.

Rosen, Gideon, 2003, "Culpability and Ignorance", *Proceedings of the Aristotelian Society*, Vol. 103, pp. 61-84.

Rosen, Gideon, 2010, "Kamm on Collaboration", *Philosophy and Phenomenological Research*, Vol. 80, No. 3, pp. 681-693.

Rubenstein, Jennifer, 2015, *Between Samaritans and States: The Political Ethics of Humanitarian INGOs*, Oxford University Press: Oxford.

Saunders, Ben, 2009, "A Defence of Weighted Lotteries in Life Saving Cases", *Ethical Theory and Moral Practice*, Vol. 12, No. 3, pp. 279-290.

Scanlon, T.M., 1998, *What We Owe To Each Other*, Harvard University Press: London.

Scanlon, T. M., 2008, *Moral Dimensions: Permissibility, Meaning, Blame,* Harvard University Press: Cambridge, Mass.

Scheffler, Samuel, 1985, "The Role of Consent in the Legitimation of Risky Activity", in M. Gibson (ed.), *To Breathe Freely*, Rowman and Allanheld: Totowa, NJ, pp. 75-88.

Schroeder, Christopher, 2003, "Causation, Compensation, and Moral Responsibility", in *The Philosophical Foundations of Tort Law*, ed. David Owen, Oxford University Press: Oxford, pp. 348-362. Segall, Shlomi, 2014, "In Defense of Priority (and Equality)", *Politics*, *Philosophy and Economics*, 30, pp. 1-22.

Sen, Amartya, 1993, "Internal Consistency of Choice", *Econometrica*, Vol. 61, No. 3, pp. 495-521.

Sher, George, 1980, "What Makes a Lottery Fair?", Nous, 14, pp. 203-216.

Stone, Peter, 2007, "Why Lotteries are Just", *The Journal of Political Philosophy*, Vol. 15, No. 3, pp. 276-295.

Temkin, Larry S., 1993, Inequality, Oxford University Press: Oxford.

Temkin, Larry, 2001, "Inequality: A Complex, Individualistic, and Comparative Notion", *Philosophical Issues*, Vol. 11, pp. 327-53.

Thomson, Judith Jarvis, 1990, *The Realm of Rights*, Harvard University Press: London.

Trautmann, Stefan T., and van de Kuilen, Gijs, 2015, "Ambiguity Attitudes" in *Blackwell handbook of Judgment and Decision Making*, Gideon Keren and George Wu eds.

Voorhoeve, Alex, 2014, "How Should We Aggregate Competing Claims?" *Ethics*, Vol. 125, pp. 64-87.

Voorhoeve, Alex, and Fleurbaey, Marc, 2012, "Egalitarianism and the Separateness of Persons", *Utilitas*, Vol. 24, Issue 3, pp. 381-398.

Voorhoeve, Alex, and Fleurbaey, Marc, 2013, "Decide as You Would with Full Information! An Argument Against ex ante Pareto", in Nir Eyal, Samia A. Hurst, Ole F. Norheim, and Dan Wikler, Oxford University Press: Oxford, pp. 113-128.

Voorhoeve, Alex and Fleurbaey, Marc, 2016, "Priority or Equality for Possible People?", *Ethics*, Vol. 126, pp. 929-54.

Voorhoeve, Alex, and Otsuka, Michael, forthcoming, 'Equality versus Priority', in S. Oslaretti (ed.), *Oxford Handbook of Distributive Justice*, Oxford University Press: Oxford.

Walzer, Michael, 1973, "Political Action: The Problem of Dirty Hands", *Philosophy & Public Affairs*, Vol. 2, No. 2, pp. 160-180.

Wasserman, David, 1996, "Let Them Eat Chances: Probability and Distributive Justice", *Economics and Philosophy*, 12, pp. 29-49.

Williams, Andrew, 2012, "The Priority View Bites the Dust?", *Utilitas*, Vol. 24, No. 3, pp. 315-331.

Williams, Bernard, 1993, "Moral Luck", in *Moral Luck*, ed. Daniel Statman, State University of New York Press: Albany, pp. 35-57.

Williams, Bernard, with J.J.C. Smart, 1973, *Utilitarianism: For and Against*, Cambridge University Press: Cambridge.