

Reports from the Department of Philosophy

Vol. 30

THE REJECTION OF PRIORITARIANISM

Arvi Pakaslahti

University of Turku

Finland

Copyright© Arvi Pakaslahti & Turun yliopisto

SERIES EDITORS

Olli Koistinen

Juha Räikkä

Department of Behavioural Sciences and Philosophy

University of Turku

FIN-20014 Turku

Finland

ISSN 1457-9332

ISBN (PRINT) 978-951-29-5954-9

ISBN (PDF) 978-951-29-5955-6

Juvenes Print – Suomen Yliopistopaino Oy, Turku, Finland 2014

Philosophy

Department of Behavioural Sciences and Philosophy

The Doctoral Programme of Social and Behavioural Sciences

Faculty of Social Sciences

University of Turku, Finland

Supervised by

Professor Eerik Lagerspetz

Philosophy

Department of Behavioural Sciences and Philosophy

University of Turku, Finland

Reviewed by

Professor Matti Häyry

Department of Management Studies

Aalto University, Finland

Professor Kaisa Herne

School of Management

University of Tampere, Finland

Opponent

Professor Kaisa Herne

The originality of this dissertation has been checked in accordance with the University of Turku quality assurance system using the Turnitin OriginalityCheck service.

ACKNOWLEDGEMENTS

I would like to thank Marko Ahteensuu, Hannu Autto, Nick Beckstead, Geoffrey Brennan, Robert E. Goodin, Nils Holtug, Ville Korpela, Olli Lappalainen, Andrew Mason, Stefan Napel, Hannu Nurmi, Wlodek Rabinowicz, Hannu Salonen, Helena Siipi and Frej Klem Thomsen for particularly valuable comments, discussions and/or advice. I would also like to thank Kaisa Herne and Matti Häyry for their reports and many very useful comments. Gustaf Arrhenius provided me with some useful comments about the writing style and the clarity of exposition. Big thanks go to Juha Räikkä for encouragement, help and many valuable comments. Special thanks go to Eerik Lagerspetz and Andrew Williams. Eerik has provided me with many extremely interesting and valuable comments over the years. The discussions which I had with Andrew when I was visiting Pompeu Fabra University were extremely important in terms of the development of my thinking regarding utilitarianism, prioritarianism and intra-personal and inter-personal trade-offs. I have received interesting and useful comments also from many other people (especially in the weekly seminar of the Public Choice Research Centre) and I thank them for their comments. For practical help I would like to thank Hemmo Laiho and Hanna-Mari Salonen. I would also like to thank the personnel (especially Marko Ahteensuu, Hannu Autto, Mika Hämäläinen and Susanne Uusitalo) of the Public Choice Research Centre and the Philosophy Unit of the Department of Behavioural Sciences and Philosophy for very friendly atmosphere.

Huge thanks go to the Public Choice Research Centre, the Faculty of Social Sciences and the Philosophy Unit for extremely important financial support. I am also very grateful to Laura and Eira for extremely generous financial support. This work has been supported by the Academy of Finland.

ABSTRACT

Since the late 1990's, a group of moral doctrines called prioritarianism has received a lot of interest from many moral philosophers. Many contemporary moral philosophers are attracted to prioritarianism to such an extent that they can be called prioritarians. In this book, however, I reject prioritarianism, including not only "pure" prioritarianism but also hybrid prioritarian views which mix one or more non-prioritarian elements with prioritarianism.

This book largely revolves around certain problems and complications of prioritarianism and its particular forms. Those problems and complications are connected to risk, impartiality, the arbitrariness of prioritarian weightings and possible future individuals. On the one hand, I challenge prioritarianism through targeted objections to various specific forms of prioritarianism. All those targeted objections are connected to risk or possible future individuals. It seems to me that together they give good grounds for believing that prioritarianism is not the way to go. On the other hand, I challenge prioritarianism by pointing out and discussing certain general problems of prioritarianism. Those general problems are connected to impartiality and the arbitrariness of prioritarian weightings. They may give additional grounds for believing that all prioritarian views should be rejected.

Prioritarianism can be seen as a type of weighted utilitarianism and thus as an extension of utilitarianism. Utilitarianism is morally ultimately concerned, and morally ultimately concerned only, with some kind of maximization of utility or expected utility. Prioritarianism, on the other hand, is morally ultimately concerned, and morally ultimately concerned only, with some kind of maximization of priority-weighted utility, expected priority-weighted utility or priority-weighted expected utility. Thus prioritarianism, unlike utilitarianism, is a distribution-sensitive moral view. Besides rejecting prioritarianism, I reject also various other distribution-sensitive moral views in this book. However, I do not reject distribution-sensitivity in morality, as I end up endorsing a type of distribution-sensitive hybrid utilitarianism which mixes non-utilitarian elements with utilitarianism.

Key words: prioritarianism, utilitarianism, risk, impartiality, morality, justice

TIIVISTELMÄ

1990-luvun loppupuolelta lähtien moraalidoktriinien joukko nimeltään prioritarismi on saanut osakseen paljon mielenkiintoa moraalifilosofien keskuudessa. Monet nykymoraalifilosofit ovat niin viehättyneitä prioritarismista, että heitä voi kutsua prioritaristeiksi. Tässä kirjassa kuitenkin hylkään sekä ”puhtaan” prioritarismin että sellaiset prioritarismin sekamuodot, jotka yhdistävät yhden tai useampia ei-prioritaristisia elementtejä prioritarismiin.

Tämä kirja pyörii pääasiassa eräiden prioritarismiin ja sen tiettyihin muotoihin liittyvien ongelmien ja hankaluuksien ympärillä. Nämä ongelmat ja hankaluudet kytkeytyvät riskiin, impartialisuuteen, prioritarististen painotusten mielivaltaisuu-teen ja mahdollisiin tuleviin yksilöihin. Yhtäältä haastan prioritarismin kohdistamalla vastalauseita monia tiettyjä prioritarismin muotoja kohtaan. Kaikki kyseiset vastalauseet kytkeytyvät riskiin tai mahdollisiin tuleviin yksilöihin. Nähdäkseni yhdessä ne antavat hyviä perusteita uskoa, että prioritarismi ei ole oikea tie. Toisaalta asetan prioritarismin kyseenalaiseksi tunnistamalla ja tarkastelemalla eräitä yleisiä prioritariamiin liittyviä ongelmia. Kyseiset ongelmat kytkeytyvät impartialisuuteen ja prioritarististen painotusten mielivaltaisuu-teen. Ne saattavat antaa lisäperusteita sen puolesta, että kaikki prioritaristiset näkemykset pitäisi hylätä.

Prioritarismi voidaan ymmärtää yhdeksi painotetun utilitarismin muodoksi ja näin ollen utilitarismin jatkeeksi. Utilitarismi on moraalisesti pohjimmiltaan kiinnostunut, ja moraalisesti pohjimmiltaan kiinnostunut ainoastaan, jonkinlaisesta utiliteetin tai odotetun utiliteetin maksimoinnista. Prioritarismi taas on moraalisesti pohjim- miltaan kiinnostunut, ja moraalisesti pohjimmiltaan kiinnostunut ainoastaan, jonkinlaisesta prioriteettipainotetun utiliteetin, odotetun prioriteettipainotetun utiliteetin tai prioriteettipainotetun odotetun utiliteetin maksimoinnista. Näin ollen prioritarismi, toisin kuin utilitarismi, on jakaumasensitiivinen moraalinäkemys. Sen lisäksi että hylkään prioritarismin, hylkään tässä kirjassa myös useita muita ja- kaumasensitiivisiä moraalinäkemyksiä. En kuitenkaan hylkää jakaumasensitiivi- syyttä moraalissa, sillä päädyn kannattamaan tietyyntyyppistä jakaumasensitiivistä utilitarismin sekamuotoa, joka yhdistää ei-utilitaristisia elementtejä utilitarismiin.

Avainsanat: prioritarismi, utilitarismi, riski, impartialisuus, moraalit, oikeudenmukaisuus

CONTENTS

1 INTRODUCTION	1
2 THE REJECTION OF EX POST PRIORITARIAN VIEWS	26
2.1 Introduction	26
2.2 Ex Ante Pareto Principles and Intra-personal Trade-offs	27
2.3 Inter-personal Trade-offs and the Scope of the Diminishing Marginal Moral Significance of an Individual's Utility	44
2.4 Conclusion	54
3 THE REJECTION OF EX ANTE PRIORITARIAN VIEWS	56
3.1 Introduction	56
3.2 Ex Ante Prioritarianism, Competing Ex Post Claims and Possible Future Individuals	62
3.3 Possible Future Individuals and the Scope of Ex Ante Prioritarianism	70
3.4 The Dominance Principle	79
3.5 Conclusion	83
4 PRIORITARIANISM, IMPARTIALITY AND IMPERSONALITY	85
4.1 Introduction	85
4.2 Prioritarianism Violates Impartiality: The First Argument	86
4.3 Prioritarianism Violates Impartiality: The Second Argument	96
4.4 Normative Considerations regarding Impartiality	97
4.5 Conclusion	98
5 GIVING FAIR CHANCES TO INDIVIDUALS	99
6 THE ARBITRARINESS OF PRIORITARIAN WEIGHTINGS	112
7 CONCLUSION	123

APPENDIX A: RABINOWICZ ON THE SCOPE OF THE DIMINISHING MARGINAL MORAL SIGNIFICANCE OF AN INDIVIDUAL'S UTILITY	125
APPENDIX B: THE DOMINANCE PRINCIPLE AND CONSEQUENTIALISM	127
BIBLIOGRAPHY	128

1 INTRODUCTION

Since the late 1990's, largely due to Derek Parfit's groundbreaking Lindley Lecture "Equality or Priority?" (published in 1995) and its shorter version "Equality and Priority" (published in 1997), a group of moral doctrines called *prioritarianism* or *the Priority View* has received a lot of interest from many moral philosophers.¹ Many contemporary moral philosophers are attracted to prioritarianism to such an extent that they can be called prioritarians. In this book, however, I reject prioritarianism, including not only "pure" prioritarianism but also hybrid prioritarian views which mix one or more non-prioritarian elements with prioritarianism.² This book is, to the best of my knowledge, the most comprehensive critique of prioritarianism published to date.³

So what is prioritarianism? In order to understand the idea, structure and motivation of prioritarianism well, it is necessary to look first at *utilitarianism*, as prioritarianism can be seen as a type of *weighted* utilitarianism⁴ and thus as an extension of utilitarianism. Why it makes sense to see prioritarianism as a type of weighted

¹ Before these writings of Parfit, prioritarianism and prioritarian ideas had been scarcely discussed by moral philosophers. On prioritarianism and prioritarian ideas before these writings of Parfit, see e.g. Broome (1991, 198-200, 216-217, 221-222); McKerlie (1984); Temkin (1993, 8, 245-248); Weirich (1983).

² For example Matthew D. Adler (2012), Richard J. Arneson (1999; 2000), Campbell Brown (2005), Krister Bykvist (2010), Nils Holtug (2007c; 2007d) and Shlomi Segall (2011) are very attracted to prioritarianism. However, Arneson (1999; 2000) argues for what he calls "responsibility-catering prioritarianism", Brown (2005) argues for what he calls "threshold prioritarianism", Segall (2011) argues for what he calls "luck prioritarianism" and Adler (2012, 37, 319-320, 579-584) makes it clear that he believes in some kind of responsibility-sensitive prioritarianism rather than pure prioritarianism. Thus each of them believes in some form of hybrid prioritarianism rather than pure prioritarianism. Also Martin O'Neill (2008; 2012) and Bertil Tungodden (2003) argue for a hybrid prioritarian view. Also Thomas Porter (2012) seems to be very attracted to prioritarianism. But also Porter (2012, 361) makes it clear that he does not believe in pure prioritarianism.

³ The best-known critique of prioritarianism must be Michael Otsuka & Alex Voorhoeve's (2009) article "Why It Matters That Some Are Worse Off Than Others: An Argument against the Priority View".

⁴ See Arneson (2001, 182); Weirich (1983, 424).

utilitarianism should become apparent shortly when I start discussing prioritarianism in more detail.⁵

Utilitarianism (or more specifically, maximizing utilitarianism) is morally ultimately concerned, and morally ultimately concerned only, with some kind of maximization of utility⁶ or expected utility.⁷ This means that according to utilitarianism, only the *overall amount* of individuals' utility or expected utility – rather than for example the overall amount of individuals' utility or expected utility *and* its *distribution* between different individuals – ultimately matters from a moral point of view.⁸ (Different kinds of distinctions may be made between utility, well-being and welfare. My impression is, however, that most moral philosophers either do not make any distinctions or do not make particularly significant distinctions between these terms. I equate well-being with welfare. By one's utility I refer to how well one's life is going.⁹ I believe that how well one's life is going is determined only by her/his/its well-being.)¹⁰

⁵ Utilitarianism is a much older moral view than prioritarianism, at least as a clearly articulated moral view in academic moral philosophy. Among the most famous utilitarian writings are Bentham (1789/1823); Hare (1981); Mill (1863/1871); Sidgwick (1874/1907); Singer (1979/2011); Smart (1973).

⁶ Here I understand the maximization of utility broadly. The maximization of utility does not refer here only to what *in fact* maximizes utility. Although a utilitarian may believe that the moral rightness and wrongness (and the moral goodness and badness) of acts are determined by their actual consequences in terms of how much utility they in fact produce or would in fact produce, another utilitarian might claim that an act is morally right if and only if its agent should (from an epistemic point of view) believe that no alternative act is *more likely* to maximize utility. This kind of probabilistic utilitarian (the former utilitarian is what may be called a non-probabilistic or actual consequence utilitarian) is also morally concerned with the maximization of utility (rather than the maximization of expected utility) but in another sense than the former utilitarian.

⁷ The maximization of utility or expected utility can be either direct (i.e. act-utilitarian) or indirect (e.g. rule-utilitarian).

⁸ The overall amount of individuals' utility or expected utility does not have to be the *total* amount of individuals' utility or expected utility. It could also be the *average* amount of individuals' utility or expected utility or some kind of mixture of the total amount and the average amount. In the context of probabilistic moral views these distinctions are relevant, roughly speaking, only regarding such choices that the chooser (i.e. the agent making the choice) can justifiably (from an epistemic point of view) believe to affect in some particular way and with some particular probability the number of individuals that will come into existence. In this book I do not discuss such choices.

⁹ See Otsuka & Voorhoeve (2009, 173). See also page 57 in Chapter 3.1 of this book.

I assume that the most plausible form of utilitarianism is such a form of probabilistic utilitarianism according to which a moral agent's act (unless perhaps her act is a purely self-regarding act) is morally right if and only if it maximizes total expected utility,¹¹ and according to which an act's total expected utility depends on the information available to the agent of the act.¹² (What I mean by this is that I believe it to be the most plausible form of pure utilitarianism. In Chapter 5 I will endorse a type of hybrid utilitarianism which mixes non-utilitarian elements with utilitarianism.) I call this form of utilitarianism *expected utility utilitarianism*.¹³ (More specifically, it could be called total maximizing expected utility act-utilitarianism.) I will soon illustrate expected utility utilitarianism with examples.

The assumption that expected utility utilitarianism is the most plausible form of utilitarianism is, of course, controversial. However, any other assumption about the

¹⁰ It should be noticed that my *definition* of utility (i.e. that utility is how well one's life is going) does not imply the view that one's utility level and well-being level always coincide (although I believe they do since I believe that how well one's life is going is determined only by her/his/its well-being) and that it is even compatible with the view that non-sentient beings can have utility (although I do not believe they can), as it could be argued that how well one's life is going is not determined merely by her/his/its well-being. I should perhaps also note that I consider hedonic and preference-orientated views about well-being the best candidates for a plausible view about well-being. It seems to me that neither the plausibility of utilitarianism nor the plausibility of prioritarianism has anything to do with which one of the above views about well-being is the most plausible one. However, it might be argued that some kind of objective list theory about well-being is the most plausible view about well-being and that such a theory is incompatible both with utilitarianism and prioritarianism. I believe that objective list theories about well-being are highly implausible, but that is beyond the scope of this book.

¹¹ It may be so that a choice situation cannot be a *moral* choice situation if it is a purely self-regarding choice situation. If a choice situation is not a moral choice situation, the chooser (i.e. the agent making the choice in that choice situation) can act neither morally wrong nor morally right in that situation.

¹² See also McCarthy (1998, esp. 66-67).

¹³ There may, however, be choice situations in which *no* act by a moral agent maximizes total expected utility due to the lack of information available to her, and which nevertheless should (from the point of view of expected utility utilitarianism) be seen as *moral* choice situations. If there are such moral choice situations, then in order to avoid the conclusion that a moral agent cannot avoid acting morally wrong in those kinds of choice situations, expected utility utilitarianism should be understood so that according to it, a moral agent's act is morally wrong if and only if it does not maximize total expected utility and some other alternative act does.

most plausible form of utilitarianism would be even more controversial. For example, it would be more controversial to assume that the most plausible form of utilitarianism is such actual consequence (and thereby non-probabilistic) utilitarianism according to which a moral agent's act is morally right if and only if it in fact maximizes total utility. It would also be more controversial to assume that expected utility theory is a less plausible approach towards risk in the context of utilitarianism than some other probabilistic approach towards risk. Both of these assumptions would be more controversial because as Tim Mulgan notes, most utilitarians are probabilistic utilitarians and more specifically such probabilistic utilitarians who prefer expected utility theory to other probabilistic approaches.¹⁴

In this book my argumentation against prioritarianism, except in Chapter 4 and Chapter 6, relies on the view that expected utility utilitarianism is the most plausible form of utilitarianism. Moreover, my view that expected utility utilitarianism is the most plausible form of utilitarianism is, for a large part, based on the view that expected utility theory is the most plausible approach towards risk in the context of practical rationality¹⁵ and on the probabilistic view of morality according to which actual consequences do not determine the moral rightness and wrongness (and the moral goodness and badness) of moral agents' acts.¹⁶ If these views are implausi-

¹⁴ Mulgan (2007, 150-151).

¹⁵ This implies that I believe it to be rational in self-interested terms for an agent to act in a way that maximizes her own expected utility. It also implies that I believe it to be irrational in self-interested terms for an agent to act in a way that does not maximize her own expected utility if there is some way of acting available to her which maximizes her own expected utility.

In certain types of "extreme" choice situations (see Temkin 2001) it might be psychologically so hard for an individual to comply with expected utility theory (i.e. to maximize her own expected utility) that because of that she would fail to comply with it in those choice situations. But even if it is psychologically extremely hard for an individual to do some particular act in some particular kind of choice situation, it might be a rational act in self-interested terms and even the only rational act in self-interested terms available to her in that choice situation.

¹⁶ It could be argued that such moral views according to which actual consequences determine the moral rightness and wrongness of moral agents' acts are not decision-guiding, as actual consequences of an act may be found out only after the act has been done. Such actual consequence views also imply the view that not acting morally wrong is to a very large extent a matter of luck. Actual consequence views also face other serious problems. For the problems of maximizing actual consequence act-utilitarianism, see Miller (2003).

ble, my argumentation against prioritarianism (except in Chapter 4 and Chapter 6) may lose at least some of its force.¹⁷

It seems to me clear that the moral psychological motivation of expected utility utilitarianism is *benevolence*.¹⁸ In fact it seems to me that expected utility utilitarianism is an extremely benevolent moral doctrine. I find that a very attractive feature of expected utility utilitarianism. Moreover, expected utility utilitarianism is clearly an *impartial* moral doctrine, as it claims that *every* individual's expected utility counts morally and counts morally *equally*. I also find the impartiality of expected utility utilitarianism an attractive property of expected utility utilitarianism. However, one may object to expected utility utilitarianism for example by arguing that expected utility utilitarianism (and utilitarianism in general) is an implausible moral view, because it does not give *intrinsic* moral significance to the *distribution* of utility or expected utility between different individuals.¹⁹ In other words, it could be argued that expected utility utilitarianism and utilitarianism in general should be rejected, because utilitarianism is *distributively insensitive*. These types of objections to utilitarianism are among the most popular objections to utilitarianism in the literature criticizing utilitarianism.²⁰

Expected utility utilitarianism (and utilitarianism in general) does not give intrinsic moral significance to the distribution of utility or expected utility between different individuals, as it is not morally ultimately concerned with the utility level or the expected utility level of *each* individual (neither in an absolute nor comparative sense). Instead expected utility utilitarianism is morally ultimately concerned *only* with different individuals' *total* expected utility. It is clear, of course, that according to expected utility utilitarianism, all those individuals whose total expected

¹⁷ See Parfit (2012, 423-425).

¹⁸ See e.g. Crisp (2003, 756); Smart (1973, 7, 56).

¹⁹ Apparently the distribution of utility or expected utility between different individuals does not even have *instrumental* moral significance from the point of view of expected utility utilitarianism. Something can be said to have instrumental moral significance from the point of view of expected utility utilitarianism if it is useful from the point of view of maximizing total expected utility. For example, some particular character trait or disposition may be such that possessing it makes one more likely in many situations to act in a way that maximizes total expected utility than she would be without possessing that character trait or disposition. Thus such a character trait or disposition may be instrumentally morally significant from the point of view of expected utility utilitarianism.

²⁰ See e.g. Scheffler (1988).

utility it is morally ultimately concerned with are intrinsically morally significant individuals in the sense that all of those individuals' expected utility *counts* in utilitarian calculations. But expected utility utilitarianism does not give special moral importance to the fact that each of those individuals is a *distinct* or *separate* individual. Related to this, John Rawls famously argued that

The striking feature of the utilitarian view of justice is that it does not matter, except indirectly, how this sum of satisfactions is distributed among individuals any more than it matters, except indirectly, how one man distributes his satisfactions over time...For just as it is rational for one man to maximize the fulfillment of his system of desires, it is right for a society to maximize the net balance of satisfaction taken over all of its members...[utilitarianism extends] to society the principle of choice for one man, and then, to make this extension work, conflating all persons into one through the imaginative acts of the impartial sympathetic spectator. Utilitarianism does not take seriously the distinction between persons.²¹

On the basis of these considerations, it could be argued that expected utility utilitarianism goes astray in terms of *distributive justice* and is thus an implausible moral view.²² One group of moral doctrines, which is morally ultimately concerned with the utility level or the expected utility level of each individual and which thus cannot be accused for not taking seriously the distinction between (or separateness of) individuals, is prioritarianism. In other words, prioritarianism gives intrinsic moral significance to the distribution of utility or expected utility between different individuals. Prioritarianism (or more specifically, maximizing prioritarianism) in general may be described by saying that it is morally ultimately concerned, and morally ultimately concerned only, with some kind of maximization of *priority-weighted* utility (which it can be concerned with in a non-probabilistic or probabilistic sense, see Footnote 6 on page 2), expected priority-weighted utility or priority-weighted expected utility. If the most plausible form of utilitarianism is

²¹ Rawls (1971/1999, 23-24).

²² It may also be argued that expected utility utilitarianism goes astray in terms of distributive justice because it does not give intrinsic moral significance to the distribution of *chances* between different individuals regarding the highest possible utility level that each of them could (on the basis of the information that is available to a chooser) get as a result of a chooser's choice. I will discuss chance-orientated moral views in Chapter 5.

expected utility utilitarianism, it is natural to assume that the most plausible form of prioritarianism is the corresponding form of prioritarianism (rather than some other type of probabilistic prioritarianism which does not utilize expected utility theory or some kind of actual consequence prioritarianism). However, it could be claimed that there is not just one but two corresponding forms of prioritarianism, namely *expected weighted utility prioritarianism* and *weighted expected utility prioritarianism*.

For the sake of brevity and for the sake of using more common terminology, from now on I will call expected weighted utility prioritarianism *ex post prioritarianism* and weighted expected utility prioritarianism *ex ante prioritarianism*. According to ex post prioritarianism, a moral agent's act (unless perhaps her act is a purely self-regarding act) is morally right if and only if it maximizes total expected correctly priority-weighted utility, whereas according to ex ante prioritarianism, a moral agent's act (unless perhaps her act is a purely self-regarding act) is morally right if and only if it maximizes total correctly priority-weighted expected utility.

The basic idea of ex post prioritarianism is that an individual's utility in a foreseeable outcome of an alternative²³ has *diminishing marginal moral significance*, whereas the basic idea of ex ante prioritarianism is that an individual's expected utility in an alternative has diminishing marginal moral significance.²⁴ In other words, according to ex post prioritarianism, utility has less marginal moral significance on higher utility levels of an individual (in a foreseeable outcome of an alternative) than on lower utility levels of an individual (in a foreseeable outcome of an alternative), whereas according to ex ante prioritarianism, expected utility has less marginal moral significance on higher expected utility levels of an individual than on lower expected utility levels of an individual. Thus ex ante prioritarianism says that the *prospects* of each individual (in terms of expected utility) have diminishing marginal moral significance.²⁵ These views imply a moral significance

²³ In this book the term "alternative", when used as a noun, refers to an act.

²⁴ For diminishing marginal moral significance, see e.g. Parfit (1995, 105-106); Temkin (2003, 64). Diminishing marginal moral significance is, of course, a completely different thing than the law of *diminishing marginal utility* according to which the more money one has, the less utility she would get from N amount of extra money.

²⁵ It should be noticed that calling expected weighted utility prioritarianism simply ex post prioritarianism can be a bit confusing and perhaps even slightly misleading, because there are also other

function that leads to a strictly concave moral significance line regarding an individual's utility (in a foreseeable outcome of an alternative) or expected utility. In other words, prioritarianism leads to such a moral significance line regarding an individual's utility or expected utility which slopes upwards but bends downwards (see Figure 1 which includes three such lines).²⁶ On the other hand, utilitarianism necessarily leads to a linear moral significance line (see Figure 2), as utilitarianism does not involve any kind of weighting. (Prioritarianism can be understood, however, to include also such views according to which the diminishing marginal moral significance of an individual's utility or expected utility and the strict concavity of a moral significance line do not apply to *negative* utility or expected utility levels of an individual.²⁷ What is a negative utility or expected utility level? According to Campbell Brown, “[i]t is common to interpret the ‘zero line’ of utility as the

forms of prioritarianism which are “ex post-orientated” in the sense that they claim that prioritarian weightings should be directed to each individual's utility in different outcomes rather than each individual's prospects. However, expected weighted utility prioritarianism may be seen as the standard form of ex post prioritarianism and may thus be called simply ex post prioritarianism. Calling expected weighted utility prioritarianism ex post prioritarianism can be a bit confusing also in the sense that “ex post” may be interpreted to refer to actual consequences which take place *after* an agent's choice rather than such consequences (e.g. foreseeable consequences) which are or can be “in the head of an agent” *before* she chooses some alternative.

Otsuka & Voorhoeve (2009, 195) claim that according to the form of prioritarianism which I call ex post prioritarianism, “outcomes rather than prospects are carriers of moral value”. This seems to me a bit misleading. According to ex post prioritarianism, the moral value of an alternative is not determined by its total correctly priority-weighted utility but its total *expected* correctly priority-weighted utility. Because of this fact, it seems to me that according to ex post prioritarianism, (foreseeable) consequences and the probabilities of those consequences *together*, and only together, carry moral value. On the other hand, it is clear that according to actual consequence prioritarianism (which is actually another type of ex post prioritarianism), (actual) consequences are the carriers of moral value. According to ex post prioritarianism (meaning expected weighted-utility prioritarianism), (foreseeable) consequences clearly have moral *significance*, but it seems to me that in the case of ex post prioritarianism (foreseeable) consequences do not have such moral significance which is also moral value. This is why I use in this book the terms “moral significance line” and “diminishing marginal moral significance” rather than the terms “moral value line” and “diminishing marginal moral value”.

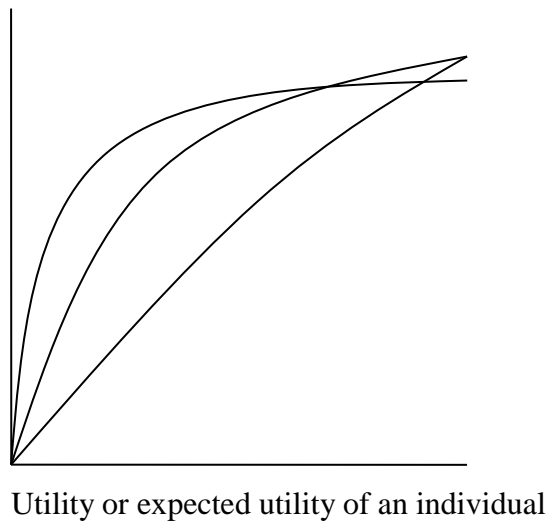
²⁶ See e.g. Brown (2005, 208-209); Hirose (2005, 374; 2011, 91, 97); Holtug (2007c, 4; 2007d, 133-134; 2010, 205-206); Peterson & Hansson (2005, 300-301).

²⁷ See Brown (2007).

boundary between those lives which are ‘worth living’ and those which are not.’²⁸ Paul Weirich writes that “[o]ur zero point for utility levels is indifference to the future, which is roughly indifference to dying now.”²⁹ However, in this book I do not explore negative utility and expected utility levels in more detail. Thus I ignore all moral philosophical complications regarding negative utility and expected utility levels that prioritarianism faces.³⁰ Related to this, in all figures in this book in which I present moral significance lines, the beginning of the x-axis represents the zero point of an individual’s utility or expected utility.)

FIGURE 1. Prioritarian moral significance lines.

Moral significance




²⁸ Brown (2007, 337).

²⁹ Weirich (1983, 428).

³⁰ On some of those complications, see Brown (2007).

FIGURE 2. A utilitarian moral significance line.



As I already said, it seems to me clear that the moral psychological motivation of expected utility utilitarianism is benevolence. However, obviously ex post prioritarianism and ex ante prioritarianism are also motivated by benevolence.³¹ But it seems to me that unlike expected utility utilitarianism, ex post prioritarianism and ex ante prioritarianism are also motivated by *compassion*,³² as they claim that an individual's utility in a foreseeable outcome of an alternative or an individual's expected utility has more marginal moral significance on lower utility or expected utility levels of an individual than on higher utility or expected utility levels of an individual. Of course, the view that the distribution of utility or expected utility between different individuals has intrinsic moral significance and the view that utilitarianism is an implausible moral view because it does not give special moral importance to the fact that each individual is a distinct or separate individual are important motivations of prioritarianism. As David McCarthy puts it, "prioritarianism is typically motivated by the claim that utilitarianism is distributively insen-

³¹ See Driver (2012, 82).

³² See Crisp (2003, 756-757). Certainly an expected utility utilitarian may be a person who in many situations *feels* a lot of compassion towards some or many individuals, but this is a different thing. An expected utility utilitarian does not give any special *role* to compassion in morality. One might argue that expected utility utilitarianism is an implausible moral doctrine because of that. However, as I see it, it may also be the case that compassion should not be given any special role in morality.

sitive and ignores the separateness of persons".³³ However, these are motivations of *all* moral views which give intrinsic moral significance to the distribution of utility or expected utility between different individuals. I think a further motivation of prioritarianism is compassion.

It is also important to see that prioritarianism is a *non-comparative* rather than *comparative* moral view in the sense that it is not concerned with how high (neither cardinally nor ordinally) utility or expected utility levels different individuals have *compared* to each other. Prioritarianism (unless it is mixed with some comparative moral view) is concerned with utility or expected utility levels of individuals only in an *absolute* sense. Thus prioritarianism is concerned with the distribution of utility or expected utility between different individuals in an absolute rather than comparative sense.³⁴ If one does not keep that in mind, she may easily confuse prioritarianism with one or two other distribution-sensitive moral views, namely *the rank-weighted priority view* and/or such weighted utilitarianism which is a mixture of utilitarianism and *egalitarianism*.

According to prioritarianism, an individual's utility or expected utility has diminishing marginal moral significance in a *cardinal* sense. According to the rank-weighted priority view, on the other hand, the moral significance of an individual's utility or expected utility depends non-lexically on how high it is in an *ordinal* sense. This kind of ordinality of the rank-weighted priority view makes it a comparative moral view. But Julia Driver confuses prioritarianism with the rank-weighted priority view. Driver writes that

implementing [prioritarianism] raises a number of difficult issues. For example, *how much* extra weight does one give to the least well-off members of the group? Does one give them double weight? Triple? How does one non-arbitrarily decide on a cutoff between least well-off and the rest of the group? The bottom 5, 10, or 20 percent?...one could adopt a sliding scale along which one weights different groups differently depending on where they fall on the scale – the lower 5 percent being weighted more heavily than the next 5 percent, and so on.³⁵

³³ McCarthy (2006, 345).

³⁴ See e.g. Parfit (1995, 103-104, 123-124); Temkin (2003, 64-65, 68-70).

³⁵ Driver (2012, 82).

In this passage Driver discusses prioritarianism as if prioritarianism gave priority to worse-off individuals over better-off individuals on the basis of ordinal considerations. Driver uses expressions like “the least well-off members of the group”, “the bottom 5, 10, or 20 percent”, “the lower 5 percent” and “the next 5 percent”. These expressions are clearly connected to how well-off different individuals are in an ordinal sense rather than how well-off different individuals are in a cardinal sense.³⁶ I will come back to the rank-weighted priority view in Chapter 4.2 in which I reject it.³⁷

According to utility-based egalitarianism (including also such weighted utilitarianism which is a mixture of utilitarianism and egalitarianism), it is intrinsically morally significant how equally utility or expected utility is distributed between different individuals. (From now on, whenever I use the term “egalitarianism”, I refer only to utility-based egalitarianism rather than egalitarianism more widely. One important type of non-utility-based egalitarianism is resource egalitarianism.) Egalitarianism, unlike the rank-weighted priority view, does not include an ordinal element. But despite that, egalitarianism is clearly a comparative moral view in the sense that it is concerned with how high utility or expected utility levels different individuals have compared to each other.³⁸ I will come back to egalitarianism later in this chapter and in Chapter 2.2 I reject it.

I can now illustrate with examples how, according to expected utility utilitarianism, ex post prioritarianism and ex ante prioritarianism, the total moral value of an alternative is calculated.³⁹ Consider

³⁶ Also Brad Hooker (2000, 56-57; 2003/2008) confuses prioritarianism with the rank-weighted priority view.

³⁷ On the rank-weighted priority view, see Adler (2012, 352-354).

³⁸ Weirich (1983) is an example of a philosopher who fails to distinguish prioritarianism from egalitarianism.

³⁹ At least in the context of aggregative moral doctrines it makes sense to talk about “the total moral value of an alternative”. From the points of view of such moral doctrines, the total moral value of each alternative act in some choice situation is what determines their moral betterness.

Multi-individual Case under Risk. You can choose between two alternatives. You are not one of the individuals mentioned in the information below that you have.⁴⁰

Alternative 1: A 100 per cent probability that the utility levels of 10,000 specific individuals will be 4.

Alternative 2: A 100 per cent probability that the utility levels of 5,000 specific individuals (of the same 10,000 individuals) will be 7. A 50 per cent probability that the utility levels of 5,000 specific individuals (of the same 10,000 individuals) will be 2 and a 50 per cent probability that their utility levels will be 1.

According to expected utility utilitarianism, Alternative 2 is a morally better alternative than Alternative 1 (and thus the morally right alternative, as there are only these two alternatives available to you), as its total expected utility is higher. This is so because $10,000 \times (1 \times 4) = 40,000$ and $5,000 \times (1 \times 7) + 5,000 \times (0.5 \times 2 + 0.5 \times 1) = 42,500$.⁴¹

According to ex post prioritarianism, on the other hand, Alternative 1 is a morally better alternative than Alternative 2 if the weightings are done with a square root. This is so because $10,000 \times (1 \times \sqrt{4}) = 20,000$ and $5,000 \times (1 \times \sqrt{7}) + 5,000 \times (0.5 \times \sqrt{2} + 0.5 \times \sqrt{1}) \approx 19,264$. (It may be asked why a square root rather than some other strictly concave function? This relates to a problem, perhaps a big one, regarding prioritarianism which I will discuss in Chapter 6. In this and many other examples in this book I use a square root,⁴² because I need to use *some* strictly concave function in those examples in order to illustrate and discuss prioritarianism and different problems which different forms of prioritarianism face, and because a square root is a mathematically very straightforward function.)⁴³

⁴⁰ I also assume that in this and my other examples in this book you cannot communicate with any of the individuals mentioned in the information that you have.

⁴¹ For the sake of clarity, in various calculations in this book the probability and utility values are in brackets even when that is not necessary.

⁴² See also Hirose (2005, 377; 2011, 98, 105) who uses a square root in the context of prioritarianism.

⁴³ Of course, no particular strictly concave function would make much sense as such in the context of prioritarianism before it is determined which qualitative well-being or utility levels different numerical utility levels (should) stand for and vice versa. For example, the number one might stand

Also according to ex ante prioritarianism (if the weightings are done with a square root), Alternative 1 is a morally better alternative than Alternative 2. This is so because $10,000 \times \sqrt{(1 \times 4)} = 20,000$ and $5,000 \times \sqrt{(1 \times 7)} + 5,000 \times \sqrt{(0.5 \times 2 + 0.5 \times 1)} \approx 19,352$. But as can be seen, the difference between the total moral value of Alternative 1 and the total moral value of Alternative 2 is bigger according to ex post prioritarianism than what it is according to ex ante prioritarianism. So although both forms of prioritarianism claim that Alternative 2 is a morally worse alternative than Alternative 1 and thus give the same prescription regarding which alternative you should morally choose, the moral worseness of Alternative 2 is slightly bigger according to ex post prioritarianism than what it is according to ex ante prioritarianism. (In chapters 2, 3 and 5 I will look into such choice situations in which ex post prioritarianism and ex ante prioritarianism give not merely different results regarding the total moral value of some alternative but also different moral prescriptions.) This is not the case in the next example in which there is certainty regarding both alternatives. It should be noticed that under certainty both ex post prioritarianism and ex ante prioritarianism necessarily produce identical results, because under certainty individuals' utility and expected utility levels coincide.

Multi-individual Case under Certainty. You can choose between two alternatives. You are not one of the individuals mentioned in the information below that you have.

Alternative 1: A 100 per cent probability that the utility levels of 10,000 specific individuals will be 4.

Alternative 2: A 100 per cent probability that the utility levels of 5,000 specific individuals (of the same 10,000 individuals) will be 7 and a 100 per cent probability that the utility levels of 5,000 specific individuals (of the same 10,000 individuals) will be 1.5.

Both according to ex post prioritarianism and ex ante prioritarianism (if the weightings are done with a square root), Alternative 1 is a morally better alterna-

for "extremely low well-being", number two for "very low well-being", number four for "moderately low well-being", number seven for "high well-being", and so on.

tive than Alternative 2. There is also no disagreement between ex post prioritarianism and ex ante prioritarianism on how much it is morally better. This is so because $10,000 \times (1 \times \sqrt{4}) = 20,000$ and $5,000 \times (1 \times \sqrt{7}) + 5,000 \times (1 \times \sqrt{1.5}) \approx 19,352$, and because $10,000 \times \sqrt{(1 \times 4)} = 20,000$ and $5,000 \times \sqrt{(1 \times 7)} + 5,000 \times \sqrt{(1 \times 1.5)} \approx 19,352$. On the other hand, according to expected utility utilitarianism, Alternative 2 is a morally better alternative than Alternative 1, because $10,000 \times (1 \times 4) = 40,000$ and $5,000 \times (1 \times 7) + 5,000 \times (1 \times 1.5) = 42,500$.

In the literature those philosophers who take up a positive attitude towards prioritarianism or some hybrid prioritarian view naturally reject utilitarianism because of its distribution-insensitivity. Besides this, those philosophers tend to argue for prioritarianism in a way that consists of one or both of the following two parts. First, they often point out some problems that some other distribution-sensitive moral views face but which prioritarianism does not face. On the basis of those problems, they argue that prioritarianism is superior to those other distribution-sensitive moral views. Secondly, they often defend prioritarianism or some particular form of prioritarianism against some objections to it.

This kind of argumentation for prioritarianism can be characterized as *negative*, as strictly speaking it tries to give good grounds for believing that prioritarianism is a *more* plausible moral view than various other moral views rather than tries to give good grounds for believing that prioritarianism is a plausible moral view in its own right. There is nothing wrong as such with this kind of argumentation for prioritarianism. However, the case for prioritarianism might be made stronger by also invoking *positive* argumentation for prioritarianism. But there is a lack of positive arguments for prioritarianism (i.e. arguments according to which there are good grounds for believing that prioritarianism is a plausible moral view in its own right) in the prioritarian literature,⁴⁴ at least if one's statement that she finds priori-

⁴⁴ On how those who are attracted to prioritarianism tend to argue for it, see also McCarthy (2008, 14); Tungodden (2003, 24). A nice example of how prioritarians tend to argue for prioritarianism comes from Holtug (2007d, 155) who writes in the end of his paper that

tarianism or some particular form of prioritarianism *intuitively* appealing or attractive (such statements are common in the prioritarian literature) does not count as a positive *argument* for it.

What might a good positive argument for prioritarianism be like? I think one could develop a positive argument for prioritarianism, and perhaps a good one, by pointing out that equally big increments or decrements (i.e. equally big increments or decrements in an absolute sense) in an individual's utility or expected utility increase or decrease an individual's utility or expected utility *proportionally* more on lower utility or expected utility levels of an individual than on higher utility or expected utility levels of an individual.⁴⁵ On the basis of this, and perhaps also by appealing to compassion and arguing for a special role of compassion in morality, one might develop an argument whose conclusion is that an individual's utility or expected utility has diminishing marginal moral significance. However, although this kind of positive argument for prioritarianism might make the case for prioritarianism stronger than invoking merely negative argumentation for prioritarianism, I do not think it would make the case for prioritarianism strong enough. If prioritarianism is looked into comprehensively, such problems can be found which cast serious doubt both on pure prioritarianism and hybrid prioritarian views. So although on the face of it, I in fact find prioritarianism intuitively quite appealing, I believe that all prioritarian views should be rejected. I believe that those problems – and various problems of various other distribution-sensitive moral views and certain problems of pure expected utility utilitarianism – give good grounds for believing that a certain type of distribution-sensitive utilitarian view (which is thus a hybrid rather than pure utilitarian view) is, on the one hand, a more plausible moral view than any form of prioritarianism and any other distribution-sensitive

I have defended prioritarianism against various objections. And in the process of doing so, I have suggested that prioritarianism is superior to egalitarianism. However, I have not provided anything like a positive case for prioritarianism. So, for all I have said, prioritarianism may be more plausible than egalitarianism and yet not plausible enough to be justifiable.

⁴⁵ For example, an increase of an individual's expected utility from 4 to 6 is in an absolute sense an equally big increase as an increase of some other individual's expected utility from 10 to 12. But proportionally the former increase is bigger, as the former increase is 50 per cent and the latter increase is 20 per cent.

moral view which I am aware of, and, on the other hand, a more plausible moral view than pure expected utility utilitarianism. I will discuss and endorse such a hybrid utilitarian view in Chapter 5.

Of those problems that some non-prioritarian distribution-sensitive moral views face but which prioritarianism does not face, the one that has been discussed most in the prioritarian literature is the so called *Levelling Down Objection*.⁴⁶ Parfit writes that

I am concerned with people's being *equally well off*. To count as Egalitarians, in my sense, this is the kind of equality in which we must believe... We may believe that inequality is *bad*. On such a view, when we should aim for equality, that is because we shall thereby make the outcome better. We can then be called *Teleological* – or, for short *Telic* – Egalitarians... If inequality is bad, its disappearance must be in one way a change for the better, *however this change occurs*. Suppose that those who are better off suffer some misfortune, so that they become as badly off as everyone else. Since these events would remove the inequality, they must be in one way welcome, on the Telic View, even though they would be worse for some people, and better for no one. This implication seems to many to be quite absurd. I call this *the Levelling Down Objection*.⁴⁷

I find the Levelling Down Objection very powerful and I believe that telic egalitarianism should be rejected because of it. On the other hand, prioritarianism (except that kind of hybrid prioritarianism which mixes prioritarianism with telic egalitarianism) clearly avoids the Levelling Down Objection, because according to prioritarianism, the goodness of an outcome is not determined at all by how equally utility is distributed in it between different individuals. Since prioritarianism avoids the Levelling Down Objection and telic egalitarianism does not, I believe that prioritarianism is a more plausible moral view (or more precisely, a less implausible moral view, as I shall reject prioritarianism because of other reasons) than telic egalitarianism. Also some prioritarians make it clear that they consider prioritarianism a more plausible moral view than telic egalitarianism because of the Levelling Down Objection.⁴⁸

⁴⁶ See also McCarthy (2008, 14).

⁴⁷ Parfit (1995, 84, 98).

⁴⁸ E.g. Arneson (1999); Holtug (2007d).

It should be noticed, however, that *deontic* egalitarianism is not vulnerable to the Levelling Down Objection.⁴⁹ Parfit describes deontic egalitarianism in the following way:

Our view may instead be *Deontological* or, for short, *Deontic*. We may believe we should aim for equality, not to make the outcome better, but for some other moral reason. We may believe, for example, that people have rights to equal shares...On [the Deontic View], it is not in itself good if people are equally well off, or bad if they are not...On the Telic View, we believe that inequality is in itself bad, or unfair. On the Deontic View, our concern about equality is only a concern about what we should do.⁵⁰

However, it can be asked whether deontic egalitarianism makes much sense in its own right. If the distinction between individuals should not be given special importance in an egalitarian fashion in the evaluations of the goodness of outcomes, why should it actually be given special moral importance in an egalitarian fashion in the moral evaluations of acts? It seems to me that deontic egalitarianism is an *ad hoc* solution for an egalitarian for avoiding the Levelling Down Objection. However, I will not explore deontic egalitarianism here in more detail. I will also not discuss any other such form of egalitarianism which definitely or possibly avoids the Levelling Down Objection⁵¹ and neither will I explore the plausibility of the Levelling Down Objection itself. Instead, in Chapter 2.2 I will give a general argument against egalitarianism which can be directed also against those forms of egalitarianism that avoid the Levelling Down Objection.

Also *leximin* and *sufficientarianism* – which are also distribution-sensitive moral views – have been discussed extensively in the prioritarian literature. The idea of *leximin* is, roughly speaking, that a moral agent should morally act in a way that maximizes the utility or the expected utility of the worst-off (the worst-off in terms of utility or expected utility), or if more than one alternative maximize it, act in a

⁴⁹ Or at least deontic egalitarianism is not vulnerable to the Levelling Down Objection in the form in which Parfit formulated the Levelling Down Objection. According to O'Neill (2008, 142-143), a form of deontic egalitarianism is vulnerable to the Levelling Down Objection if we reformulate the Levelling Down Objection.

⁵⁰ Parfit (1995, 84, 88, 94).

⁵¹ On one such form of egalitarianism, see Holtug (2007a).

way that maximizes the utility or the expected utility of the second worst-off, or if more than one alternative maximize it, act in a way that maximizes the utility or the expected utility of the third worst-off, etc. Leximin is seen by some moral philosophers as a type of prioritarianism. Richard J. Arneson refers to leximin as “an extreme version of prioritarianism”,⁵² Iwao Hirose as “one special version of Prioritarianism”⁵³ and “the extreme form of Prioritarianism”,⁵⁴ Andrew Mason as “extreme prioritarianism”,⁵⁵ and Tim Mulgan as “lexical prioritarianism”.⁵⁶ I think, however, that leximin should not be considered the lexical or extreme form of prioritarianism, because leximin, unlike prioritarianism (as I have presented it), is a comparative moral view in the sense that it is concerned with how high utility or expected utility levels different individuals have compared to each other.⁵⁷ This is a big structural difference between leximin and prioritarianism.

I find leximin a completely implausible moral view. To see how implausible leximin is, consider

Multi-individual Case under Certainty II. You can choose between two alternatives. You are not one of the individuals mentioned in the information below that you have.

Alternative 1: A 100 per cent probability that the utility levels of 9,999 specific individuals will be extremely low and a 100 per cent probability that the utility level of one specific individual will be extremely low and even slightly lower than each of the 9,999 individuals’ utility.

Alternative 2: A 100 per cent probability that the utility levels of the same 9,999 individuals will be extremely high and a 100 per cent probability that the utility level of the same single individual will be extremely low and even slightly lower than it would be as a result of you choosing Alternative 1.

⁵² Arneson (2000, 341).

⁵³ Hirose (2005, 383).

⁵⁴ Hirose (2011, 102-103).

⁵⁵ Mason (2006, 114).

⁵⁶ Mulgan (2007, 111).

⁵⁷ See Rabinowicz (2001, 148, 153, 162-163; 2002, 8, 13-14).

Leximin, unlike prioritarianism, says that Alternative 1 is a morally better alternative than Alternative 2. But I do not think it makes any sense to consider Alternative 1 a morally better alternative than Alternative 2. Alternative 2 is a *considerably* better alternative (both from ex ante and ex post perspectives) for 9,999 individuals, whereas Alternative 1 is only a *slightly* better alternative (both from ex ante and ex post perspectives) for only *one* individual. Moreover, it should be noticed that as a result of you choosing Alternative 1 the utility of each of the 9,999 individuals would be *close to* as low as the utility of the single individual would be as a result of you choosing Alternative 2. On the basis of these considerations, it is hard for me to see how you could act morally wrong by choosing Alternative 2.⁵⁸

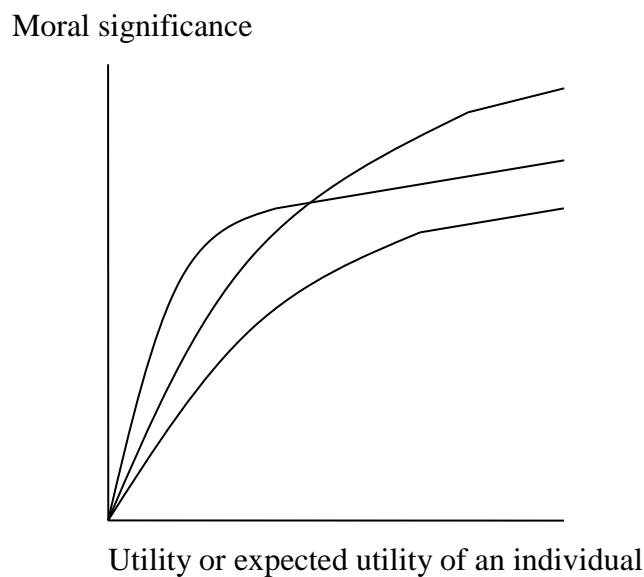
Sufficientarianism is structurally more similar to prioritarianism than leximin is. According to sufficientarianism, the correct moral significance line is strictly concave up to a certain utility or expected utility level of an individual and linear after that (see Figure 3 which includes three sufficientarian moral significance lines).⁵⁹ Thus sufficientarianism is structurally very similar to prioritarianism. In this book I concentrate on prioritarianism rather than sufficientarianism because of two reasons. First, although prioritarianism and sufficientarianism are structurally very similar to each other, prioritarianism is structurally simpler than sufficientarianism. Thus concentrating on sufficientarianism would be more complex. Secondly, I am inclined to believe that prioritarianism is a more plausible moral view than sufficientarianism. As I see it, through the idea of the diminishing marginal moral significance of an individual's utility or expected utility, prioritarianism and sufficientarianism give a special role to compassion in morality. But *if* compassion should be given a special role in morality, I do not think there could be any *specific* morally unique utility or expected utility level of an individual on which the special

⁵⁸ Many moral philosophers have rejected leximin along similar lines. See e.g. Broome (2004, 23-29); Brown (2007, 351-352); Crisp (2003, 752; 2006, 153-154); Holtug (2007b, 8; 2010, 224-225). I will come back to leximin in Chapter 4.2 and Chapter 5.

⁵⁹ *Lexical* sufficientarianism adds one extra feature which is that the point in which the moral significance line changes from strictly concave to linear is also a lexical threshold. This means that according to lexical sufficientarianism, any increment (or decrement) in an individual's utility or expected utility below the threshold, regardless of how small, is morally more significant than any increment (or decrement) in an individual's utility or expected utility above the threshold.

role of compassion should completely vanish.⁶⁰ However, the question whether it is prioritarianism or sufficientarianism which is more plausible is not relevant regarding my conclusions about various specific forms of prioritarianism and prioritarianism in general, because all my arguments against specific forms of prioritarianism can be directed against the corresponding forms of sufficientarianism and because all my arguments about the general problems of prioritarianism apply also to sufficientarianism.

FIGURE 3. Sufficientarian moral significance lines.



Now let's have a brief look at certain problems and complications of prioritarianism and its particular forms around which the rest of the book largely revolves. Those problems and complications are connected to *risk*, *impartiality*, *the arbitrariness of prioritarian weightings* and *possible future individuals*.

In recent years many moral philosophers have discussed prioritarianism in the context of risk. Although for a long time risk was largely ignored in the literature

⁶⁰ On sufficientarianism and the relative plausibility of prioritarianism and sufficientarianism, see e.g. Arneson (1999); Brown (2005); Crisp (2003; 2006, 152-162); Holtug (2007d, 149-155; 2010, 226-236); Mason (2006, 122-126). Arneson (1999) and Holtug (2007d, 149-155; 2010, 226-236) reject sufficientarianism in favour of prioritarianism, whereas Crisp (2003) rejects prioritarianism in favour of lexical sufficientarianism and Crisp (2006, 152-162) in favour of moderate, non-lexical, sufficientarianism. Brown's (2005) view, on the other hand, is a mixture of prioritarianism and lexical sufficientarianism.

that concentrated on prioritarianism, it could be said that risk has become the main theme in the prioritarian literature.⁶¹ In this book Chapter 2, Chapter 3 and Chapter 5 revolve around risk. The plausibility of prioritarianism and the question what form of prioritarianism is the most plausible form of prioritarianism are very much connected to risk. Because of certain problems that ex post prioritarianism and ex ante prioritarianism can face in choice situations that involve risk but which they cannot face in any choice situations in which there is certainty, both ex post prioritarianism and ex ante prioritarianism may seem to work much better in those kinds of moral choice situations in which there is certainty than in many such moral choice situations that involve risk. Related to this, the following two important questions can be formulated. First, even if ex post prioritarianism and ex ante prioritarianism (or some kinds of desert-sensitive versions of them in case desert should be given intrinsic moral significance)⁶² work well in all such moral choice situations in which there is certainty, does either of them work well enough in all moral choice situations involving risk? My answer will be negative. Secondly, if ex post prioritarianism does not work well enough in all moral choice situations involving risk, could ex post prioritariness, in order to overcome this problem, plausibly limit the scope of ex post prioritarian reasoning (or to put it differently, plausibly limit the scope of the diminishing marginal moral significance of an individual's utility) to certain kinds of moral choice situations? Also to this question my answer will be negative.

Impartiality and the arbitrariness of prioritarian weightings have been discussed much less in the prioritarian literature than risk. My main question regarding impartiality is whether prioritarianism violates impartiality or not. This question has been discussed by Krister Bykvist⁶³ and Brad Hooker.⁶⁴ I will discuss it in Chapter 4 and I will do it more comprehensively than Bykvist and Hooker have done. I will argue that prioritarianism violates impartiality. The question whether

⁶¹ Several writings on prioritarianism revolve at least partly around risk. These writings include Adler (2012); Bognar (2012); Crisp (2011); McCarthy (2006; 2008; 2012); O'Neill (2012); Otsuka (2012); Otsuka & Voorhoeve (2009; 2011); Parfit (2012); Porter (2012); Rabinowicz (2001; 2002); Voorhoeve & Fleurbaey (2012); Williams (2012).

⁶² I will discuss desert quite a lot in Chapter 2.2.

⁶³ Bykvist (2010, 71-72).

⁶⁴ Hooker (2000, 27, 59-62).

prioritarianism violates impartiality is important, as moral philosophers widely believe that plausible moral reasoning implies impartiality.

Questions related to the arbitrariness of prioritarian weightings have been explored even less in the prioritarian literature than the question whether prioritarianism violates impartiality. I will discuss questions related to the arbitrariness of prioritarian weightings in Chapter 6. Since prioritarianism involves weighting of each individual's utility or expected utility, one may wonder how it could be determined *non-arbitrarily* how much prioritarian weighting (where this means more than no prioritarian weighting at all) on each utility or expected utility level of an individual is the correct amount of prioritarian weighting. Moreover, if it cannot be determined non-arbitrarily, what *kind* of problem is it for prioritarianism? Is it such a problem because of which we should perhaps reject prioritarianism? I will argue that it may be.

In Chapter 3.3 I will give an argument against ex ante prioritarianism that is connected to possible future individuals and which I have not come across in the prioritarian literature. I will reject ex ante prioritarianism, because I believe that it cannot include possible future individuals in its scope. Because of this problem, I will also discuss such a moral view which is partly ex ante prioritarian and partly ex post prioritarian and such a moral view which is partly ex ante prioritarian and partly utilitarian, as these views can include possible future individuals in their scope. But I will reject both of these moral views because of other kinds of problems they face related to possible future individuals.

It should be noticed that in this book I challenge prioritarianism both through *targeted objections* to specific forms of prioritarianism (i.e. through objections connected to risk and possible future individuals) and through *general problems* of prioritarianism (i.e. through problems connected to impartiality and the arbitrariness of prioritarian weightings). I will give targeted objections to various specific forms of prioritarianism in Chapter 2, Chapter 3 and Chapter 5.⁶⁵ I believe that together these arguments give good grounds for believing that prioritarianism is not the way to go. In Chapter 4 and Chapter 6 I will point out and discuss general

⁶⁵ On the other hand, I will also reject two targeted objections to ex ante prioritarianism. I will do that in Chapter 3.1 and Chapter 3.2. The targeted objection that I reject in Chapter 3.2 is connected to possible future individuals.

problems of prioritarianism. These problems may give additional grounds for believing that all prioritarian views should be rejected.

Before going to the next main chapter, I would like to point out three things. First, most of my examples in this book are very empty examples in the sense that they lack many practical details of the choice situation. In other words, they lack a “story”. That makes those examples shorter than they would otherwise be and may make it easier for many readers to concentrate on those moral questions and problems which I want to illustrate with those examples.

Secondly, and partly related to the emptiness of most of my examples, I should perhaps note that in order to illustrate and discuss various moral questions and problems clearly, many of my examples are numerical examples in the sense that I use numbers in them to describe individuals’ well-being levels. I am aware that there are quite a few people who do not like such examples. I ask those people to bear with me. I see the numbers only as tools for illustrating various moral questions and making clear and interesting arguments. A different question is to what extent moral agents are justified (from an epistemic point of view), or whether they are ever justified, to assign numerical utility levels to different individuals in real life. But this question is beyond the scope of this book. Also the question how well expected utility utilitarianism, *ex post* prioritarian views and *ex ante* prioritarian views could cope without any (inter-personal) cardinal utility scale is beyond the scope of this book.

Thirdly, it should be noticed that prioritarian weightings can also be *backward-looking*. In other words, they do not have to be merely *forward-looking*. Consider

Past and Future. You know that you can benefit either eighty-year-old John or eighty-year-old Mike and that as a result of that either John’s or Mike’s utility will be somewhat higher than it would otherwise be. You know that John’s utility is currently very high. You also know that if you benefit John, his utility will be extremely high, whereas it would otherwise remain very high. You also know that Mike’s utility is currently extremely low. You also know that if you benefit Mike, his utility will be very low, whereas it would otherwise remain extremely low. You also know that John’s utility has been extremely low for the most of his life and that Mike’s utility has been extremely high for the most of his life.

According to such a prioritarian who holds that prioritarian weightings should not be backward-looking at all, it would be morally wrong to benefit John. But clearly one could hold that prioritarian weightings should be both forward-looking and backward-looking. Such a prioritarian would say that it would be morally wrong to benefit Mike. The question whether prioritarian weightings should (in case a plausible moral doctrine includes prioritarian weightings which I will deny) be also backward-looking is, however, outside the scope of this book.⁶⁶ All the other examples in this book are formulated so that this question is not relevant regarding them, because each of the other examples is such that there is no information available to a chooser (i.e. the agent making the choice) about those individuals' past utility levels whose utility is at stake non-trivially. (I will explain in Chapter 2.2 what I mean by an individual's utility being at stake non-trivially.)

⁶⁶ Similar distinctions between forward-looking and backward-looking views have been made and discussed by Holtug (1999, 25-26) in the context of prioritarianism, Benbaji (2005, 329-330) and Crisp (2003, 762-763; 2006, 161-162) in the context of sufficientarianism, and McKerlie (1989) in the context of egalitarianism. Benbaji (2005, 329) writes that

a person who has led a great life for eighty years can be in need of romance or remission of the cancer from which he suffers. Hence, meeting his needs is not of special importance, according to a whole-life version of the doctrine of sufficiency...[Another] principle mandates assigning priority to one's needs in any case where meeting this need is the only way to prevent one's *present* condition from becoming bad...the doctrine of sufficiency – if it is built on the distinction between bad and good lives – assigns priority to people whose lives are bad.

2 THE REJECTION OF EX POST PRIORITARIAN VIEWS

2.1 Introduction

In this main chapter I will show that there are both *intra-personal trade-off choice situations*, or more specifically, other-regarding intra-personal trade-off choice-situations (i.e. other-regarding, rather than purely self-regarding, choice situations which involve trade-offs within lives but not between or across lives)¹, and *inter-personal trade-off choice situations* (i.e. other-regarding choice situations which involve trade-offs between or across lives)² in which ex post prioritarianism violates a moral principle that I call *the restricted desert-sensitive weak ex ante Pareto principle*.³ I will also argue that the restricted desert-sensitive weak ex ante Pareto principle is a plausible moral principle. I will also discuss whether an ex post prioritarian could plausibly argue that other-regarding intra-personal trade-off choice situations do not belong to the scope of the diminishing marginal moral significance of an individual's utility (i.e. that an individual's utility does not have diminishing marginal moral significance in other-regarding intra-personal trade-off choice situations). I will also discuss whether an ex post prioritarian could tenably go even further and plausibly argue that such inter-personal trade-off choice situations in which an ex post prioritarian view can violate the restricted desert-sensitive weak ex ante Pareto principle also do not belong to the scope of the diminishing marginal moral significance of an individual's utility (i.e. that an individual's utility does not have diminishing marginal moral significance also in such inter-personal

¹ As I already noted in Footnote 11 on page 3, it may be the case that no purely self-regarding choice situation is a *moral* choice situation. Because of this I will concentrate on such intra-personal trade-off choice situations which are other-regarding choice situations.

² Naturally no inter-personal trade-off choice situation can be a purely self-regarding choice situation.

³ My demonstration that there are both other-regarding intra-personal trade-off choice situations and inter-personal trade-off choice situations in which ex post prioritarianism violates the restricted desert-sensitive weak ex ante Pareto principle is ultimately based on Otsuka & Voorhoeve (2009). See also Adler (2012, 503-504, 520-521); Rabinowicz (2001; 2002).

trade-off choice situations in which an ex post prioritarian view can violate the restricted desert-sensitive weak ex ante Pareto principle). I will reject both limitations of the scope of the diminishing marginal moral significance of an individual's utility by an ex post prioritarian. I will conclude that both pure ex post prioritarianism and all such forms of ex post prioritarianism which are partly non-prioritarian should be rejected.

2.2 Ex Ante Pareto Principles and Intra-personal Trade-offs

Perhaps the best way to approach the question whether all forms of ex post prioritarianism should be rejected is through discussing various ex ante Pareto principles. First consider

The weak ex ante Pareto principle: If in some choice situation (unless perhaps it is a purely self-regarding choice situation) some particular alternative maximizes everyone's expected utility (where "everyone" refers to each of those individuals whose utility is at stake non-trivially in that choice situation) and every other alternative is such that it does not maximize anyone's expected utility (where "anyone" refers to any of those individuals whose utility is at stake non-trivially in that choice situation), a moral agent acts morally wrong if she does not choose the alternative that maximizes everyone's expected utility.

By individuals whose utility is at stake in some choice situation I refer, roughly speaking, to individuals whose utility, on the basis of the information that is available to a chooser, will or may be affected by the chooser's choice in some particular foreseeable way. An individual's utility may also be at stake in a way that I call trivial if she is the chooser. I consider the chooser's utility to be at stake trivially rather than non-trivially in some choice situation if what maximizes her expected utility in that choice situation depends merely on which alternative she considers to be the morally best alternative in that choice situation. We can assume that most moral agents are morally motivated agents in the sense that they generally get or would get some sort of satisfaction from acting in a way that they consider a

morally better way of acting than any alternative way of acting. So if in some choice situation some particular alternative does not maximize a chooser's expected utility merely because she considers it a morally wrong alternative despite the facts that it maximizes everyone else's expected utility and that none of the other alternatives maximizes the expected utility of any of those other individuals, the weak ex ante Pareto principle, at least as I have formulated it, claims that the chooser acts morally wrong if she does not choose the alternative that maximizes everyone else's expected utility. I should perhaps also note that whenever in this book I talk about the total expected utility, the total expected priority-weighted utility or the total priority-weighted expected utility of some alternative, it does not include a chooser's own expected utility, expected priority-weighted utility or priority-weighted expected utility, as I assume that in all my examples that contain information about some individual's or individuals' utility levels, a chooser's utility either is not at stake at all⁴ or is at stake merely trivially, except for one example in Footnote 21 on page 39.

Consider also

The strong ex ante Pareto principle: If in some choice situation some particular alternative maximizes everyone's expected utility and if at least regarding one individual that alternative is the only alternative that maximizes her expected utility, a moral agent acts morally wrong if she does not choose the alternative that maximizes everyone's expected utility.⁵

⁴ A chooser's utility may not be at stake even trivially in some choice situation if in that choice situation she considers all the alternatives morally equally good alternatives.

⁵ Allan Gibbard (1984, 263) formulates the strong ex ante Pareto principle in the following way:

The [strong ex ante Pareto] principle says...that where prospective interests are unopposed, they should prevail. More precisely, let P and Q be alternative policies. Then, the principle says, if both (1) the prospects of at least one person are better under P than under Q , and (2) no one's prospects are better under Q than under P , then ethically speaking, P is better than Q .

Adler (2012, 503) formulates ex ante Pareto-superiority so that "[i]f each person's ex ante well-being with action b is at least as large as her ex ante well-being with action a , and at least one person's ex ante well-being with b is strictly larger, b is a morally better action than a ."

It may look like that the weak ex ante Pareto principle makes a very weak claim. It certainly makes a weaker claim than the strong ex ante Pareto principle, which may also look like making a very weak claim. I believe, however, that the weak ex ante Pareto principle is an implausible moral principle. And if the weak ex ante Pareto principle is an implausible moral principle, then naturally also the strong ex ante Pareto principle is an implausible moral principle. Consider

Possible Future Individuals and Existing Individuals. You can choose between two alternatives. You are not one of the individuals mentioned in the information below that you have.

Alternative 1: A 100 per cent probability that 10,000 individuals will come into existence and that their utility levels will be 1. A 100 per cent probability that the utility levels of 10 specific (and already existing) individuals will be 10.

Alternative 2: A 100 per cent probability that 10,000 different individuals than in Alternative 1 will come into existence and that their utility levels will be 10. A 100 per cent probability that the utility levels of the same 10 already existing individuals will be 1.

Alternative 2 maximizes total expected utility, because $10,000 \times (1 \times 1) + 10 \times (1 \times 10) = 10,100$ and $10,000 \times (1 \times 10) + 10 \times (1 \times 1) = 100,010$. But it seems to me that Alternative 1 maximizes the expected utility of each of those individuals whose utility is at stake non-trivially and that Alternative 2 does not maximize the expected utility of any of those individuals, because it seems to me that possible future individuals cannot have properties, including also the property of expected utility.⁶ But I am convinced that Alternative 2 is a morally much better alternative

⁶ Even if in Possible Future Individuals and Existing Individuals the 10,000 individuals that will come into existence if you choose Alternative 1 are the same individuals as the individuals that will come into existence if you choose Alternative 2, it would not imply that their identities are fixed at the time you are choosing between the alternatives. So even if the individuals that will come into existence are the same regardless of which alternative you choose, it would not necessarily have to be the case that it is already determined which individuals will come into existence. Thus even if the individuals that will come into existence are the same regardless of which alternative you choose, at the time you are choosing between the alternatives those possible future individuals that will in fact

due to its much higher total expected utility. I believe that Possible Future Individuals and Existing Individuals shows that the weak ex ante Pareto principle is an implausible moral principle (and thus that the strong ex ante Pareto principle is an implausible moral principle) if it is true that possible future individuals cannot have properties. (In Chapter 3.3 I will discuss briefly the question whether possible future individuals can have properties. One might think that the claim that in Possible Future Individuals and Existing Individuals Alternative 2 maximizes total expected utility is incompatible with the claim that Alternative 2 does not maximize the expected utility of any of the individuals whose utility is at stake non-trivially. But I think I manage to show in Chapter 3.3 that they are not incompatible with each other.) Thus consider next what may be called

The restricted weak ex ante Pareto principle: If in some choice situation some particular alternative maximizes everyone's expected utility, if every other alternative in that choice situation is such that it does not maximize anyone's expected utility, and if the total expected utility of each alternative in that choice situation consists only of existing individuals' expected utility, a moral agent acts morally wrong if she does not choose the alternative that maximizes everyone's expected utility.

I am perplexed with whether the restricted weak ex ante Pareto principle is a plausible moral principle. I am perplexed with that because some ways of giving intrinsic moral significance to desert are incompatible with the restricted weak ex ante Pareto principle. Consider

Multi-individual Case under Risk Combined with Desert. You can choose between two alternatives. You are not one of the individuals mentioned in the information below that you have.

Alternative 1: A 100 per cent probability that the utility levels of 10,000 specific individuals will be 4.

come into existence could be merely possible future individuals rather than necessary future individuals that will necessarily come into existence.

Alternative 2: A 50 per cent probability that the utility levels of the same 10,000 individuals will be 8 and a 50 per cent probability that their utility levels will be 1.

Somehow you also have good grounds for believing that the moral deservingness level of each of the 10,000 individuals, based on how morally meritoriously each of them has acted in the past, is 4.

It could be argued that although in Multi-individual Case under Risk Combined with Desert by choosing Alternative 2 you would maximize the expected utility of each of the 10,000 individuals and by choosing Alternative 1 you would not maximize the expected utility of any of them, Alternative 1 is a more just and fairer alternative than Alternative 2 and thus a morally better alternative than Alternative 2, because you have good grounds for believing that by choosing Alternative 1 (and only by choosing Alternative 1) all the 10,000 individuals would, with a 100 per cent probability, end up having that utility level which they morally deserve to have. This is a *meritoriousness*-orientated way of giving intrinsic moral significance to desert (but not the only possible meritoriousness-orientated way of giving intrinsic moral significance to desert). Later in this sub-chapter I will point out another, non-meritoriousness-orientated, way of giving intrinsic moral significance to desert, which is a very different way of giving intrinsic moral significance to desert.

My intuitions say that desert has intrinsic moral significance. It is, however, very unclear to me in what way(s) and how much desert has intrinsic moral significance.⁷ Since some ways of giving intrinsic moral significance to desert are incompatible with the restricted weak ex ante Pareto principle and since it is so unclear to me in what way(s) desert should be given intrinsic moral significance, I am much more comfortable with endorsing what may be called

The restricted desert-sensitive weak ex ante Pareto principle: If in some choice situation some particular alternative maximizes everyone's expected utility, if every other alternative in that choice situation is such that it does not maximize

⁷ On different ways of giving intrinsic moral significance to desert, see e.g. Arneson (1999); Feldman (1995); Temkin (2011).

anyone's expected utility, and if the total expected utility of each alternative in that choice situation consists only of existing individuals' expected utility, a moral agent acts morally wrong if she does not choose the alternative that maximizes everyone's expected utility, provided that there is no information available to her about anyone's deservingness or things related to evaluating anyone's deservingness (where "anyone" refers to any of those individuals whose utility is at stake non-trivially in that choice situation).⁸

However, there are various ways of challenging the restricted desert-sensitive weak ex ante Pareto principle. Some of them appeal to the view that there are types of acts that are intrinsically morally bad or even intrinsically morally wrong. Consider

Voluntary Euthanasia. You are in a situation in which you can perform euthanasia on a human being whom you know to be terminally ill and to be suffering immensely. She has told you that her life is not worth living any more and that

⁸ Multi-individual Case under Risk Combined with Desert is outside the scope of the restricted desert-sensitive weak ex ante Pareto principle (i.e. the restricted desert-sensitive weak ex ante Pareto principle is silent about the moral betterness of Alternative 1 and Alternative 2 in Multi-individual Case under Risk Combined with Desert), because in Multi-individual Case under Risk Combined with Desert there *is* information available to you about the deservingness levels of the individuals whose utility is at stake non-trivially.

I should perhaps note that I am also comfortable with endorsing the restricted desert-sensitive *strong* ex ante Pareto principle.

The restricted desert-sensitive strong ex ante Pareto principle: If in some choice situation some particular alternative maximizes everyone's expected utility, if at least regarding one individual that alternative is the only alternative that maximizes her expected utility, and if the total expected utility of each alternative in that choice situation consists only of existing individuals' expected utility, a moral agent acts morally wrong if she does not choose the alternative that maximizes everyone's expected utility, provided that there is no information available to her about anyone's deservingness or things related to evaluating anyone's deservingness.

However, in order to plausibly reject pure ex post prioritarianism and other forms of ex post prioritarianism in this main chapter, there is no need for me to endorse anything stronger than the restricted desert-sensitive weak ex ante Pareto principle.

she wants you to kill her. She is the only individual whose utility is at stake non-trivially. You would maximize her expected utility by performing euthanasia and only by performing euthanasia.⁹

One may hold that by performing voluntary euthanasia you cannot act unjustly or unfairly in a situation in which by performing voluntary euthanasia on some specific individual you would maximize everyone's expected utility, in which by not performing voluntary euthanasia on that individual you would not maximize anyone's expected utility, in which the total expected utility of each alternative consists only of existing individuals' expected utility, and in which there is no information available to you about the deservingness, or things related to evaluating the deservingness, of any of the individuals whose utility is at stake non-trivially. But one might also hold that although you would not act unjustly or unfairly by performing voluntary euthanasia in such a situation, performing it would be morally wrong, because performing euthanasia on a human being is *disrespecting the sanctity of the human life* and is thus intrinsically morally bad and perhaps even intrinsically morally wrong.

The supporter of this kind of deontological moral view¹⁰ could clearly be charged with "heartlessness" towards the terminally ill individual who is suffering immensely.¹¹ Related to this, it seems to me that this kind of deontological moral view implies the view that mankind exists to serve morality rather than the other way around, at least in some such situations in which receiving euthanasia would clearly be best for the patient and in which the patient has made it clear that she wants to receive euthanasia, as this kind of deontological moral view implies the view that moral considerations related to justice, fairness and well-being can be

⁹ I assume that if in Voluntary Euthanasia you are not legally allowed to perform euthanasia, you know that you can perform euthanasia in such a way or in such circumstances that there will not be legal repercussions on you.

¹⁰ By a deontological moral view I refer, roughly speaking, to such a moral view, according to which there are act-types (e.g. committing a suicide, performing euthanasia, murdering, cheating, lying and/or stealing) that are intrinsically morally bad or even intrinsically morally wrong, and which puts special moral significance, importance or emphasis on an agent herself not doing such acts.

¹¹ See Smart (1973, 5-6).

morally outweighed by some other kinds of moral considerations. But it seems to me that morality exists to serve mankind (and sentient beings more generally) rather than the other way around. As William Frankena claimed, “[m]orality is made for man, not man for morality”.¹² For the sake of simplicity, I call this kind of view *the Morality for Mankind View*.¹³ It is a metanormative view about the nature and purpose of morality.¹⁴

¹² Frankena (1963/1973, 116). How about *involuntary* euthanasia? Consider

Involuntary Euthanasia. You are in a situation in which you can perform euthanasia on a human being whom you know to be terminally ill and to be suffering immensely. You also know, unlike her, that her suffering will get soon even worse unless you kill her. She has told you that she does not want you to kill her, because she mistakenly believes that she is not terminally ill and that she will get better. But you have good grounds for believing that if she knew that she is terminally ill and that her suffering will get soon even worse, she would tell you that her life is not worth living any more and that she wants you to kill her. She is the only individual whose utility is at stake non-trivially. You would maximize her expected utility by performing euthanasia and only by performing euthanasia.

The case for performing euthanasia is stronger in Voluntary Euthanasia than in Involuntary Euthanasia. Nevertheless, it seems to me that if “morality is made for man”, then also in Involuntary Euthanasia you would act morally better by performing euthanasia than by not performing euthanasia. (This, of course, does not imply that there are situations in which involuntary euthanasia should be legal from a moral point of view, as the legalisation of any kind of involuntary euthanasia would probably have many bad consequences.) But I strongly believe that if in Involuntary Euthanasia it would be morally wrong to perform euthanasia, it would be morally wrong because it is *against the will* of the patient. In other words, I do not think that performing euthanasia in Involuntary Euthanasia would be intrinsically morally bad or wrong. On another example which is important regarding the plausibility of the restricted desert-sensitive weak ex ante Pareto principle and which raises similar questions, see Footnote 35 on pages 50-51 in the next sub-chapter. On involuntary euthanasia, see Singer (1979/2011, 158, 176-178).

¹³ This term is misleading in the sense that I do not believe in an anthropocentric morality. I believe that all sentient non-human animals are intrinsically morally significant beings.

¹⁴ Gibbard (1984, 261) interprets Frankena’s slogan “as suggesting that questions of morality can most fundamentally be addressed by considering human benefits and human harms—those benefits and harms to which our accepting various alternative moral principles would tend to lead.” Gibbard (1984, 263) also clearly believes that the view that morality is made for man and not man for morality means that morality should serve humanity. (Gibbard uses the term “humanity” rather than the term “mankind”.) See also Footnote 19 on page 37.

I think the Morality for Mankind View has at least five important implications. The discussion of one of them I postpone until Chapter 3.4. First, as I already implied, I think the Morality for Mankind View implies that moral considerations related to justice, fairness and well-being cannot be morally outweighed by some other kinds of moral considerations. Actually it seems to me that the Morality for Mankind View implies the view that morality is fundamentally only about justice, fairness and well-being.¹⁵

Secondly, it seems to me that the Morality for Mankind View implies that all non-decision-guiding moral doctrines should be rejected. In order to fully serve mankind, morality must serve guidance to moral agents in choice situations. Such moral views, according to which actual consequences determine the moral rightness and wrongness of moral agents' acts (and which thus are incompatible with the restricted desert-sensitive weak ex ante Pareto principle), may not be decision-guiding, as actual consequences of an act – unlike foreseeable consequences of an act based on the information that an agent has – may be found out only after the act has been done. As I see it, all non-decision-guiding moral doctrines are incompatible with the Morality for Mankind View.

Thirdly, it seems to me that the Morality for Mankind View implies that *if* we should give intrinsic moral significance to the distribution of utility or expected utility between different individuals and thus should be morally ultimately concerned with the utility level or the expected utility level of each individual, we should be concerned with it *for the sake of* (at least some of) *those individuals themselves*, and only for the sake of (at least some of) those individuals themselves, rather than for the sake of something else. But if one is an egalitarian and thus is concerned with how equally utility or expected utility is distributed between different individuals, she is concerned with the distribution of utility or expected utility between those individuals partly (if she is morally ultimately concerned with the utility level or the expected utility level of each individual because of egalitarian and some non-egalitarian reasons) or only (if she is morally ultimately concerned with the utility level or the expected utility level of each individual only

¹⁵ It is not clear to me whether there is some philosophically important difference between justice and fairness. I assume, however, that there is at least not such a difference between them which would be significant regarding the plausibility of the main conclusions of this main chapter. On possible differences between justice and fairness, see Temkin (2011, 66-69).

because of egalitarian reasons) *for the sake of an abstract relation* rather than only for the sake of those individuals themselves. This does not seem to me to be compatible with the Morality for Mankind View.¹⁶ Of course, arguably no one or hardly anyone is a pure egalitarian in the sense that she is morally concerned *only* with how equally utility or expected utility is distributed between different individuals.¹⁷ But if that how equally utility or expected utility is distributed between

¹⁶ It should be noticed that this argument against egalitarianism is more pervasive than the Levelling Down Objection, because it, unlike the Levelling Down Objection, can be used against any form of egalitarianism that gives intrinsic moral significance to how equally utility or expected utility is distributed between different individuals. On the other hand, it should also be noticed that prioritarianism is clearly concerned with the distribution of utility or expected utility between different individuals for the sake of some individuals themselves rather than for the sake of something else, as it is concerned with the distribution of utility or expected utility between different individuals for the sake of worse-off individuals (worse-off individuals in terms of utility or expected utility).

¹⁷ See Peterson & Hansson (2005, 302). Consider

Multi-individual Case under Certainty III. You can choose between two alternatives. You are not one of the individuals mentioned in the information below that you have.

Alternative 1: A 100 per cent probability that the utility levels of 10,000 specific individuals will be 4.

Alternative 2: A 100 per cent probability that the utility levels of the same 10,000 individuals will be 5.

Both alternatives have equally equal distributions of utility (in terms of the foreseeable outcomes of those alternatives) and expected utility between the 10,000 individuals. But even if it makes sense to claim that it is intrinsically morally significant how equally utility or expected utility is distributed between different individuals, would it make sense to claim, as a pure egalitarian might do, that Alternative 2 is not a morally better alternative than Alternative 1 but is instead morally as good an alternative as Alternative 1? I think not. See also Holtug (1999, 18-19).

I wrote “as a pure egalitarian might do” rather than “as a pure egalitarian does”, because what a pure egalitarian would say about the moral betterness of Alternative 1 and Alternative 2 would ultimately depend on how high utility or expected utility levels *all* individuals, whose utility or expected utility levels you have information about, have compared to each other rather than how high utility or expected utility levels the 10,000 individuals have compared to each other. For example, if in Multi-individual Case under Certainty III you lived in a world in which very few individuals’ utility is lower than 5 and almost everyone’s utility is 5 or higher and in which you are also aware of these facts, it would be clear that a pure egalitarian would claim that Alternative 2 is a morally better alternative than Alternative 1. Egalitarians in general hold that it is morally relevant

different individuals does not have any intrinsic moral significance, then all hybrid egalitarian moral doctrines, including also all *prigitarian*¹⁸ moral doctrines (i.e. those moral doctrines which are a mixture of prioritarianism and egalitarianism), are implausible moral doctrines.

Fourthly, and partly related to the first three implications of the Morality for Mankind View, it seems to me that the Morality for Mankind View implies that the restricted desert-sensitive weak ex ante Pareto principle should be accepted.¹⁹

Related to this, consider first

Multi-individual Case under Risk II. You can choose between two alternatives. You are not one of the individuals mentioned in the information below that you have.

Alternative 1: A 100 per cent probability that the utility levels of 10,000 specific individuals will be 4.

Alternative 2: A 50 per cent probability that the utility levels of the same 10,000 individuals will be 8 and a 50 per cent probability that their utility levels will be 1.

According to expected utility utilitarianism, Alternative 2 is a morally better alternative than Alternative 1, because $10,000 \times (1 \times 4) = 40,000$ and $10,000 \times (0.5 \times 8 + 0.5 \times 1) = 45,000$. Also according to ex ante prioritarianism (regardless of the priority function), Alternative 2 is a morally better alternative than Alternative 1. If the weightings are done with a square root, the total moral value of Alternative 1 comes out as 20,000 and the total moral value of Alternative 2 approximately as 21,213. This is so because $10,000 \times \sqrt{1 \times 4} = 20,000$ and $10,000 \times \sqrt{0.5 \times 8 + 0.5 \times 1} \approx 21,213$. But according to ex post prioritarianism (if the weightings are done with a square root), Alternative 1 is a morally better alternative than Alternative 2.

how high utility or expected utility levels all individuals, whose utility or expected utility levels you have information about, and not merely how high utility or expected utility levels those individuals whose utility is at stake, have compared to each other. See also Parfit (2012, 407-408, 417).

¹⁸ This term comes from Paula Casal (2007, 309).

¹⁹ See also Gibbard (1984, 263, 278-282). Gibbard (1984, 263) writes that “[the strong ex ante Pareto principle] should be acceptable to anyone who accepts that morality should serve humanity—that morality is made for man and not man for morality.”

This is so because $10,000 \times (1 \times \sqrt{4}) = 20,000$ and $10,000 \times (0.5 \times \sqrt{8} + 0.5 \times \sqrt{1}) \approx 19,142$.

Since in Multi-individual Case under Risk II ex post prioritarianism (combined with the square root weightings) says that you act morally wrong if you choose Alternative 2, since there are no possible future individuals involved, since there is no information available to you about the deservingness, or things related to evaluating the deservingness, of any of the individuals whose utility is at stake non-trivially, and since Alternative 2 maximizes the expected utility of each of those individuals whose utility is at stake non-trivially and Alternative 1 does not maximize the expected utility of any of those individuals, we can see that in Multi-individual Case under Risk II ex post prioritarianism can violate the restricted desert-sensitive weak ex ante Pareto principle. But if expected utility theory is – as I believe it is – the most plausible approach towards risk in the context of practical rationality, it would be irrational in self-interested terms (on the basis of the information that you have about the alternatives) for each of the 10,000 individuals to choose Alternative 1 if she could make the choice instead of you. Because of this, it seems to me that by choosing Alternative 2 you would not act unjustly or unfairly *towards* any of the 10,000 individuals. I think these considerations and the fact that there is no information available to you about the deservingness, or things related to evaluating the deservingness, of any of the 10,000 individuals give – together with the view that morality exists to serve mankind – very good grounds for believing that by choosing Alternative 2 you would not act morally wrong and that by choosing Alternative 1 you would act morally wrong.

It also seems to me reasonable to hold that by choosing the alternative (i.e. Alternative 2) which gives better prospects (in terms of expected utility) to each of the 10,000 individuals than any other alternative (i.e. Alternative 1) does, you would *treat* each of those individuals better than you would treat any of them by choosing any other alternative (i.e. Alternative 1),²⁰ provided that expected utility theory is the most plausible approach towards risk in the context of practical rationality in which case it would be irrational in self-interested terms (on the basis of the information that you have about the alternatives) for each of the 10,000 individuals to

²⁰ See also Temkin (2011, 75) who connects, in the context of egalitarianism, concerns for individuals' prospects with concerns for how individuals are treated.

choose Alternative 1 if she could make the choice instead of you.²¹ If by choosing Alternative 2 you would treat each of those 10,000 individuals better than you would treat any of them by choosing Alternative 1, then by choosing Alternative 2 you would also treat each of those individuals justly and fairly. It is hard for me to see how by treating everyone else justly and fairly you could *act* unjustly or unfairly (as a whole, all things considered) in a case like Multi-individual Case under Risk II in which there is no information available to you about the deservingness, or things related to evaluating the deservingness, of any of the individuals whose utility is at stake non-trivially.²² It is equally hard for me to see how by treating everyone else justly and fairly you could act morally wrong in a case like Multi-individual Case under Risk II if the Morality for Mankind View is a plausible view.²³

Moreover, it seems to me that the fact that Multi-individual Case under Risk II does not involve inter-personal trade-offs gives even more reason to believe that you would not act morally wrong by choosing Alternative 2 and that you would act morally wrong by choosing Alternative 1. In Multi-individual Case under Risk II the nice or attractive features of Alternative 2 regarding each of the 10,000 indi-

²¹ Consider also

Self-regarding One-individual Case under Risk. You can choose between two alternatives.

Alternative 1: A 100 per cent probability that your utility level will be 4.

Alternative 2: A 50 per cent probability that your utility level will be 8 and a 50 per cent probability that your utility level will be 1.

If expected utility theory is the most plausible approach towards risk in the context of practical rationality, you would act rationally in self-interested terms by choosing Alternative 2 and irrationally in self-interested terms by choosing Alternative 1. In that case I think it is reasonable to hold that you would treat yourself better by choosing Alternative 2 than you would treat yourself by choosing Alternative 1.

²² If there was such information available to you and there would be good grounds for believing that none of the individuals would deserve to be treated in the best possible way, the situation would be morally more complex.

²³ Cf. O'Neill (2012) and Porter (2012) with my arguments in this and the previous paragraph. O'Neill and Porter do not consider arguments like these and fail to see the untenability of ex post prioritarianism in the context of other-regarding intra-personal trade-off choice situations. Cf. also Parfit (2012, 423) with my arguments.

viduals (i.e. that it maximizes the expected utility of each of the 10,000 individuals and that it is the only alternative that maximizes the expected utility of at least one of those 10,000 individuals) do *not* come at the expense of making trade-offs *between* (or across) the lives of those 10,000 individuals. The choice situation in Multi-individual Case under Risk II is an intra-personal trade-off choice situation rather than an inter-personal trade-off choice situation. Thus Multi-individual Case under Risk II involves merely intra-personal trade-offs, trade-offs *within* lives.²⁴ If Multi-individual Case under Risk II was transformed into a case which involves inter-personal trade-offs but which is nevertheless identical with Multi-individual Case under Risk II in terms of how high expected utility each of the 10,000 individuals has in Alternative 1 and Alternative 2, there would be better grounds for believing (although I do not think good enough grounds for believing) that Alternative 1 is a more just alternative from the point of view of distributive justice than Alternative 2 and that choosing Alternative 2 would consequently be morally wrong. (I will come back to this in the next sub-chapter in which I discuss such a case.) But since Multi-individual Case under Risk II does not involve inter-personal trade-offs, it seems to me that there is an additional good reason for believing

²⁴ The fact that the choice situation in Multi-individual Case under Risk II is not an inter-personal trade-off choice situation (although Multi-individual Case under Risk II is a multi-individual case) can be seen more easily if it is reduced to a one-individual case. (If the choice situation in Multi-individual Case under Risk II was an inter-personal trade-off choice situation, it naturally could not be reduced to a one-individual case. This does not mean, however, that all multi-individual cases in which the choice situation is an intra-personal trade-off choice situation could be reduced to one-individual cases. But I will not go into that here, as it is not relevant regarding my arguments in this main chapter.)

Other-regarding One-individual Case under Risk. You can choose between two alternatives. You are not the individual mentioned in the information below that you have.

Alternative 1: A 100 per cent probability that the utility level of one specific individual will be 4.

Alternative 2: A 50 per cent probability that the utility level of the same individual will be 8 and a 50 per cent probability that her utility level will be 1.

As the name of this case suggests, it is an other-regarding one-individual case rather than a self-regarding one-individual case. If you were the individual mentioned in the information about the alternatives and the case would thus be a self-regarding one-individual case, you might not be even making a moral choice when you are choosing between the alternatives.

that Alternative 2 is a more just alternative from the point of view of distributive justice than Alternative 1.

There are, of course, other weightings combined with which ex post prioritarianism would not violate the restricted desert-sensitive weak ex ante Pareto principle in Multi-individual Case under Risk II. However, even such ex post prioritarianism whose moral significance function leads to such a strictly concave moral significance line that is very close to a linear utilitarian moral significance line (see Figure 1 on page 9 which contains such a strictly concave moral significance line) would violate the restricted desert-sensitive weak ex ante Pareto principle in some possible other-regarding intra-personal trade-off choice situations. Since ex post prioritarianism (regardless of the priority function) cannot avoid violating the restricted desert-sensitive weak ex ante Pareto principle in every possible other-regarding intra-personal trade-off choice situation, there are good grounds for believing that ex post prioritarianism should be rejected.

I strongly believe that expected utility utilitarianism (unlike ex post prioritarianism) cannot give an implausible moral prescription in any such intra-personal trade-off choice situation in which there is no information available to a chooser about the deservingness, or things related to evaluating the deservingness, of any of the individuals whose utility is at stake non-trivially. I also strongly believe that ex ante prioritarianism (unlike ex post prioritarianism) cannot give an implausible moral prescription in any such intra-personal trade-off choice situation in which there are no possible future individuals involved²⁵ and in which there is no information available to a chooser about the deservingness, or things related to evaluating the deservingness, of any of the individuals whose utility is at stake non-trivially.²⁶ However, I am inclined to believe that expected utility utilitarianism and ex

²⁵ As I already noted on page 23, I believe that ex ante prioritarianism cannot include possible future individuals in its scope. I will discuss this in Chapter 3.3. It seems to me clear that ex ante prioritarianism should be rejected if it is true that it cannot include possible future individuals in its scope.

²⁶ Crisp (2011, 106) gives an example in which the alternatives available to an agent are incomparable from the point of view of the restricted desert-sensitive weak ex ante Pareto principle (i.e. in which the restricted desert-sensitive weak ex ante Pareto principle is silent about the moral betterness of the alternatives):

ante prioritarianism cannot avoid giving an implausible moral prescription in every possible such choice situation in which there is information available to a chooser about the deservingness, or things related to evaluating the deservingness, of some or all of the individuals whose utility is at stake non-trivially, because I am inclined to believe that desert should be given intrinsic moral significance in some way(s). I am quite sympathetic to the view that desert should be given intrinsic moral significance in a meritoriousness-orientated way. I am also quite sympathetic, or perhaps very sympathetic, to the view that desert should be given intrinsic moral significance in a *responsibility*-orientated way. Related to the latter view, consider the following case by Richard J. Arneson:

Imagine that a parent has a child with a 50 per cent chance of a disability. The parent must decide whether to move to the city (where the disability could be catered for more effectively) or to the suburbs (where an able-bodied child would have a better life than in the city).

Crisp (2011, 107) argues that

If the parent moves to the suburbs and the child becomes disabled, one can imagine the child (presumably when he has grown up a little) complaining as follows: ‘I accept that the expected utility of each move was the same for me. But you should have paid greater attention to the fact that I had two possible futures: one in which I was disabled, and one in which I was not. You should then have asked yourself which possible future person should be given priority to the other – and the answer would have been clear: the one who is worse off.’

Both according to expected utility utilitarianism and ex ante prioritarianism, the alternatives are morally equally good. Ex post prioritarianism, on the other hand, claims that it would be morally wrong for the parent not to move with her child to the city. Crisp writes as if the choice situation was an inter-personal trade-off choice situation involving possible future individuals. But the choice situation is clearly an intra-personal trade-off choice situation and does not involve possible future individuals. Crisp (2011, 107-108) also makes it clear in the next two paragraphs that he considers the choice situation an intra-personal trade-off choice situation. Since Crisp’s argument is so misleading, it is implausible. Moreover, in Chapter 3.1 I will discuss another case in which the choice situation is an intra-personal trade-off choice situation and in which two alternatives are incomparable from the point of view of the restricted desert-sensitive weak ex ante Pareto principle due to the fact that both of them maximize the expected utility of each of the individuals whose utility is at stake non-trivially. (In that particular case each individual comes to the same thing as the only individual.) I believe that expected utility utilitarianism and ex ante prioritarianism, unlike ex post prioritarianism, do not give an implausible moral prescription in that example.

Suppose that a national park service rescue team can choose just one of three lifesaving missions. Each involves significant risk of severe harm to rescue workers but promises a significant net saving of lives. Suppose these risks and benefits are the same for each of the three rival missions. The park rescue team must choose either to assist (a) a party of stranded schoolchildren caught in an unanticipated blizzard while on a school outing, (b) a party of experienced climbers who carefully chose to pursue a difficult route under hazardous conditions which then suddenly turned desperate, or (c) a party of tourists who ignored warning signs and the stern advice of park rangers to venture on a foolhardy hike across a treacherous steep slope, rendered more treacherous by their mid-hike alcohol consumption. One might suppose that the rescue team's policy should be set in part by consideration of its incentive effects on the behavior of future park visitors, but suppose the park is about to be shut down and there are no such incentive effects to consider.²⁷

The problem is that neither expected utility utilitarianism, ex ante prioritarianism (regardless of the priority function) nor ex post prioritarianism (regardless of the priority function) can judge (b) and (c) as morally worse alternatives than (a). Arneson continues that

I take it to be a datum in this case that the fully voluntary choice of the climbers to shoulder the risk they take and the grossly reckless conduct of the hikers reduce their moral claims to be aided by comparison with the claim of the stranded schoolchildren. This is the basic idea of the responsibility-catering element in responsibility-catering prioritarianism. A proposed theory of justice that excludes it excludes a factor that is intrinsically morally important.²⁸

If desert has intrinsic moral significance (e.g. in some meritoriousness-orientated and/or responsibility-orientated way), then some sort of desert-sensitive expected utility utilitarianism, desert-sensitive ex ante prioritarianism or desert-sensitive ex post prioritarianism could be developed.²⁹ But the problem of ex post prioritarianism and desert-sensitive ex post prioritarianism is that they can give implausible moral prescriptions even in such intra-personal trade-off choice situations in which the total expected utility of each alternative consists only of existing individuals'

²⁷ Arneson (2000, 348).

²⁸ Arneson (2000, 348).

²⁹ See e.g. Arneson (1999; 2000); Feldman (1995); Segall (2011).

expected utility and in which there is no information available to a chooser about the deservingness, or things related to evaluating the deservingness, of any of the individuals whose utility is at stake non-trivially.

2.3 Inter-personal Trade-offs and the Scope of the Diminishing Marginal Moral Significance of an Individual's Utility

An ex post prioritarian could make an interesting defensive move. She might argue that the fact that (desert-sensitive) ex post prioritarianism can violate the restricted desert-sensitive weak ex ante Pareto principle in Multi-individual Case under Risk II does not imply that all forms of ex post prioritarianism should be rejected, as ex post prioritarianism can be mixed with expected utility utilitarianism. She could invoke Andrew Williams who writes that

[The Restrictive Priority View] assumes the individuals have claims on each other's beneficence but that the content of those claims differs significantly depending on whether we need to resolve normative conflicts within rather than across lives. In cases involving only intrapersonal conflicts this Restrictive Priority View assumes individuals have a claim that conflicts are resolved in a way that maximizes their expected utility...In contrast, in cases of interpersonal conflict individuals have claims to be benefited that become stronger as their absolute position worsens. On the Restrictive View, then, the conviction that benefiting people matters more as they become worse off does not apply to all our decisions but instead is triggered by a specific context in which we face interpersonal conflicts...the Restrictive View affirms a discontinuity between the factors that should govern the distribution of benefits within lives and across lives[.]³⁰

The Restrictive Priority View is a mixture of ex post prioritarianism and expected utility utilitarianism. It limits the scope of the diminishing marginal moral significance of an individual's utility to inter-personal trade-off choice situations. Williams writes about the rationale of the Restrictive Priority View in the following way:

³⁰ Williams (2012, 323-324).

The rationale [for the Restrictive View] starts with the assumption that when an individual or set of individuals makes trade-offs within their own lives considered in isolation from the lives of others each has undefeated reasons to maximize her own expected utility; there is no distinct additional requirement to attach greater weight to benefits that might fall at a lower absolute level of advantage. The rationale then insists that the normative situation does not change when the same decision must be made by a second decision-maker acting on behalf of the previous individuals...According to the rationale, the normative situation then changes when we move from intrapersonal to interpersonal conflict resolution. It does so because when a purely intrapersonal conflict is resolved by withholding some benefit to an individual such a decision does not necessarily come at any particular individual's expense. The decision-maker can always appeal to the possibility that withholding some benefit produces a greater benefit for that very same individual that he himself had reason to prefer, and perhaps actually does prefer. In contrast, when a decision-maker resolves an interpersonal conflict by withholding some benefit to an individual such a justification appealing to the unity within a life is not always available. Instead the decision comes at the expense of some individual for the sake of some separate individual, and some other type of justification is necessary to show that sufficient weight was attached to the losing individual. The rationale insists that these specific conditions result in the weight of our reasons to benefit individuals diminishing as the recipients enjoy greater benefits on some absolute scale. Perhaps further investigation will show such insistence to be unsound, but the pedigree of the rationale does at least suggest that the Restricted View is not a merely an *ad hoc* manoeuvre.³¹

Along similar lines, one could argue that

Multi-individual Case under Risk II does not belong to the *scope* of the diminishing marginal moral significance of an individual's utility, because the choice situation in it is an intra-personal trade-off choice situation rather than an inter-personal trade-off choice situation. Expected utility utilitarianism is an implausible moral doctrine, because it ignores the morally important fact that each individual is a distinct (or separate) individual. Ex post prioritarianism, on the other hand, does not ignore this morally important fact. However, the fact that each individual is a distinct individual is a morally important fact *only* in inter-personal trade-off choice situations. Thus it makes sense to hold that those moral choice situations which are inter-personal trade-off choice situations belong to the scope of the diminishing marginal moral significance of an individual's utility and those

³¹ Williams (2012, 326-327).

moral choice situations which are intra-personal trade-off choice situations belong to the scope of the non-diminishing, linear, marginal moral significance of an individual's utility. In other words, an ex post prioritarian principle should be invoked in inter-personal trade-off choice situations and a utilitarian principle should be invoked in other-regarding intra-personal trade-off choice situations. Consequently the correct moral doctrine is what may be called *utilitarian-adjusted ex post prioritarianism*.

However, even if it is plausible to argue in this way that Multi-individual Case under Risk II does not belong to the scope of the diminishing marginal moral significance of an individual's utility and that utilitarian-adjusted ex post prioritarianism (or desert-sensitive utilitarian-adjusted ex post prioritarianism) is a more plausible moral view than pure ex post prioritarianism (or desert-sensitive ex post prioritarianism),³² I think the former would nevertheless be an implausible moral view. I believe so because despite the limitation of the scope of the diminishing marginal moral significance of an individual's utility, utilitarian-adjusted ex post prioritarianism is incompatible with the restricted desert-sensitive weak ex ante Pareto principle. Consider the following case in which not only ex post prioritarianism but also utilitarian-adjusted ex post prioritarianism (as well as desert-sensitive versions of them) – but naturally neither expected utility utilitarianism nor ex ante prioritarianism – can violate the restricted desert-sensitive weak ex ante Pareto principle:

³² I should perhaps note that if one holds in the way I described that an ex post prioritarian principle should be invoked in inter-personal trade-off choice situations and a utilitarian principle should be invoked in other-regarding intra-personal trade-off choice situations, then it clearly seems to imply that she believes that a utilitarian principle should also be invoked in what may be called other-regarding *non-trade-off choice situations* which involve neither inter-personal nor intra-personal trade-offs. In the previous sub-chapter in Footnote 17 on pages 36-37 I presented a case which I called Multi-individual Case under Certainty III and in which the choice situation is an other-regarding non-trade-off choice situation. In that particular case Alternative 2 would maximize the utility of each of the 10,000 individuals with a 100 per cent probability and there is a 100 per cent probability that Alternative 1 would not maximize the utility of any of them. Since it is *not* the case that *both* alternatives *might* or *would* maximize the utility of at least one of those 10,000 individuals, the choice situation in Multi-individual Case under Certainty III includes neither inter-personal nor intra-personal trade-offs.

Multi-individual Case under Risk and Unspecificity. You can choose between two alternatives. You are not one of the individuals mentioned in the information below that you have.

Alternative 1: A 100 per cent probability that the utility levels of 10,000 specific individuals will be 4.

Alternative 2: A 100 per cent probability that the utility levels of 5,000 unspecific individuals (of the same 10,000 individuals) will be 8 and a 100 per cent probability that the utility levels of 5,000 unspecific individuals (of the same 10,000 individuals) will be 1.

According to expected utility utilitarianism, Alternative 2 is a morally better alternative than Alternative 1. This is so because $10,000 \times (1 \times 4) = 40,000$ and $10,000 \times (0.5 \times 8 + 0.5 \times 1) = 45,000$. Also according to ex ante prioritarianism (regardless of the priority function), Alternative 2 is a morally better alternative than Alternative 1. If the weightings are done with a square root, the total moral value of Alternative 1 comes out as 20,000 and the total moral value of Alternative 2 approximately as 21,213. This is so because $10,000 \times \sqrt{(1 \times 4)} = 20,000$ and $10,000 \times \sqrt{(0.5 \times 8 + 0.5 \times 1)} \approx 21,213$. But according to ex post prioritarianism (combined with the square root weightings) and utilitarian-adjusted ex post prioritarianism (combined with the square root weightings), Alternative 1 is a morally better alternative than Alternative 2. This is so because $10,000 \times (1 \times \sqrt{4}) = 20,000$ and $10,000 \times (0.5 \times \sqrt{8} + 0.5 \times \sqrt{1}) \approx 19,142$.³³

³³ It may be worth pointing out here that in Multi-individual Case under Risk and Unspecificity the total expected utility of Alternative 2 can be calculated also with the calculation $5,000 \times (1 \times 8) + 5,000 \times (1 \times 1) = 45,000$ and the total expected priority-weighted utility of Alternative 2 can be calculated also with the calculation $5,000 \times (1 \times \sqrt{8}) + 5,000 \times (1 \times \sqrt{1}) \approx 19,142$. However, the total priority-weighted expected utility of Alternative 2 clearly could not be calculated with the calculation $5,000 \times \sqrt{(1 \times 8)} + 5,000 \times \sqrt{(1 \times 1)} \approx 19,142$. The fact that the total expected utility of Alternative 2 can be calculated also with the calculation $5,000 \times (1 \times 8) + 5,000 \times (1 \times 1) = 45,000$ and the fact that the total priority-weighted expected utility of Alternative 2 cannot be calculated with the calculation $5,000 \times \sqrt{(1 \times 8)} + 5,000 \times \sqrt{(1 \times 1)} \approx 19,142$ reflect the fact that expected utility utilitarianism, unlike ex ante prioritarianism, is not morally ultimately concerned with each individual's expected utility level, as it is clear that by using the calculation $5,000 \times (1 \times 8) + 5,000 \times (1 \times 1) = 45,000$, one calculates the total expected utility of Alternative 2 without first calculating each

The choice situation in Multi-individual Case under Risk and Unspecificity is an inter-personal trade-off choice situation, albeit not at the level of prospects but merely at the level of (foreseeable) outcomes. Since there are inter-personal trade-off choice situations in which utilitarian-adjusted ex post prioritarianism can violate the restricted desert-sensitive weak ex ante Pareto principle (as Multi-individual Case under Risk and Unspecificity shows), it seems to me that utilitarian-adjusted ex post prioritarianism is an implausible moral view.³⁴ However, a supporter of utilitarian-adjusted ex post prioritarianism might argue that

Since the choice situation in Multi-individual Case under Risk and Unspecificity is an inter-personal trade-off choice situation rather than an intra-personal trade-off choice situation, there are much better grounds for believing that you would act morally wrong by choosing Alternative 2 in Multi-individual Case under Risk and Unspecificity than there are for believing that you would act morally wrong by choosing Alternative 2 in Multi-individual Case under Risk II. In Multi-individual Case under Risk and Unspecificity the nice features of Alternative 2 regarding each of the 10,000 individuals (i.e. that it maximizes the expected utility of each of the 10,000 individuals and that it is the only alternative that maximizes the expected utility of at least one of those 10,000 individuals) come at the expense of making trade-offs *between* the lives of those 10,000 individuals. They do *not* come *merely* at the expense of making trade-offs *within* the lives of those 10,000 individuals unlike in Multi-individual Case under Risk II. Since the fact that each individual is a distinct (or separate) individual has special moral importance, it is the inter-personal trade-offs – rather than the fact that Alternative 2 maximizes the expected utility of each of the 10,000 individuals and the fact that Alternative 2 is the only alter-

individual's expected utility (which is 4.5) in Alternative 2. Ex post prioritarians, on the other hand, can calculate the total expected priority-weighted utility of Alternative 2 in two different ways, because they are morally ultimately concerned with each individual's utility level (in foreseeable outcomes of different alternatives) rather than each individual's expected utility level.

³⁴ There are, of course, other weightings combined with which utilitarian-adjusted ex post prioritarianism would not violate the restricted desert-sensitive weak ex ante Pareto principle in Multi-individual Case under Risk and Unspecificity. However, even those versions of utilitarian-adjusted ex post prioritarianism whose moral significance functions regarding inter-personal trade-off choice situations lead to such strictly concave moral significance lines which are very close to a linear utilitarian moral significance line (see Figure 1 on page 9 which contains such a strictly concave line) would violate the restricted desert-sensitive weak ex ante Pareto principle in some possible inter-personal trade-off choice situations. See also Adler (2012, 503-504, 522).

native that maximizes the expected utility of at least one of those 10,000 individuals – that we should be morally concerned with in Multi-individual Case under Risk and Unspecificity. Thus the restricted desert-sensitive weak ex ante Pareto principle should be discarded.

On the basis of this argument, it seems to me that there are better grounds for believing that you would act morally wrong by choosing Alternative 2 in Multi-individual Case under Risk and Unspecificity than there are for believing that you would act morally wrong by choosing Alternative 2 in Multi-individual Case under Risk II. But I also believe that there are *not good enough* grounds for believing that you would act morally wrong by choosing Alternative 2 in Multi-individual Case under Risk and Unspecificity.

The problem with the argument of the supporter of utilitarian-adjusted ex post prioritarianism is that instead of openly confronting the problem that in Multi-individual Case under Risk and Unspecificity utilitarian-adjusted ex post prioritarianism can prescribe you not to choose an alternative that maximizes the expected utility of each of the 10,000 individuals and which is the only alternative that maximizes the expected utility of at least one of those individuals, it *assumes* that there is a morally important difference – rather than tells why there is a morally important difference – between other-regarding intra-personal trade-off choice situations and *all kinds* of inter-personal trade-off choice situations. In the previous sub-chapter on pages 38-39 I gave two such arguments (i.e. the argument that by choosing Alternative 2 you would not act unjustly or unfairly towards any of the 10,000 individuals and the argument that by choosing Alternative 2 you would treat each of the 10,000 individuals justly and fairly) for the view that you would not act morally wrong by choosing Alternative 2 in Multi-individual Case under Risk II which can be invoked for holding that also in Multi-individual Case under Risk and Unspecificity you would not act morally wrong by choosing Alternative 2. I believe that in the context of Multi-individual Case under Risk and Unspecificity those arguments give good grounds for believing that there is not a morally important difference between other-regarding intra-personal trade-off choice situations and all kinds of inter-personal trade-off choice situations. Thus I believe that utilitarian-adjusted ex post prioritarianism gives special moral importance to the fact that each individual is a distinct individual *in a wrong way*. It seems to me that

if special moral importance should be given to the fact that each individual is a distinct individual, it should be given in such a way which implies the view that there is a morally important difference between other-regarding intra-personal trade-off choice situations and those kinds of inter-personal trade-off choice situations which are inter-personal trade-off choice situations also at the level of prospects rather than in a way which implies the view that there is a morally important difference between other-regarding intra-personal trade-off choice situations and all kinds of inter-personal trade-off choice situations.

Alternatively, rather than questioning the plausibility of the restricted desert-sensitive weak ex ante Pareto principle,³⁵ a supporter of utilitarian-adjusted ex post

³⁵ But what if each of the 10,000 individuals was risk-averse and you were also aware of that? Consider

Multi-individual Case under Risk and Unspecificity Combined with Risk-aversiveness. You can choose between two alternatives. You are not one of the individuals mentioned in the information below that you have.

Alternative 1: A 100 per cent probability that the utility levels of 10,000 specific highly risk-averse individuals will be 4.

Alternative 2: A 100 per cent probability that the utility levels of 5,000 unspecific individuals (of the same 10,000 individuals) will be 8 and a 100 per cent probability that the utility levels of 5,000 unspecific individuals (of the same 10,000 individuals) will be 1.

You also know that each of the 10,000 individuals is so risk-averse that she would, on the basis of the information that you have about the alternatives, consider it irrational in self-interested terms to choose Alternative 2 if she could make the choice instead of you and would consequently choose Alternative 1.

Provided that expected utility theory is the most plausible approach towards risk in the context of practical rationality in which case it would be irrational in self-interested terms (on the basis of the information that you have about the alternatives) for each of the 10,000 individuals in Multi-individual Case under Risk and Unspecificity Combined with Risk-aversiveness to choose Alternative 1 if she could make the choice instead of you, it seems to me that by choosing Alternative 2 you would treat each of the 10,000 individuals better than you would treat any of them by choosing Alternative 1. However, perhaps you would treat each of those individuals *more respectfully* by choosing Alternative 1 than you would treat any of them by choosing Alternative 2, as by choosing Alternative 1 you would show respect towards their views about risk or their risk attitudes. Thus it may make sense to make a distinction between treating someone better and treating someone more respectfully.

prioritarianism might argue for a further limitation of the scope of the diminishing marginal moral significance of an individual's utility in the following way:

Actually it is not so that all inter-personal trade-off choice situations belong to the scope of the diminishing marginal moral significance of an individual's utility. Only those inter-personal trade-off choice situations belong to the scope of the diminishing marginal moral significance of an individual's utility which are also *inter-personal conflict choice situations*. All other types of moral choice situations belong to the scope of the non-diminishing, linear, marginal moral significance of an individual's utility. Since in Multi-individual Case under Risk and Unspecificity it would be rational in self-interested terms (on the basis of the information that you have about the alternatives) for each of the 10,000 individuals to choose Alternative 2 if she could make the choice instead of you, the choice situation in Multi-individual Case under Risk and Unspecificity is not an inter-personal trade-off choice situation at the level of prospects. The choice situation in it is an inter-personal trade-off choice situation merely at the level of (foreseeable) outcomes. Thus the choice situation in Multi-individual Case under Risk and Unspecificity is not an inter-personal conflict choice situation. However, cases like Multi-individual Case under Risk, Multi-individual Case under Certainty, Multi-individual Case under Certainty II, Past and Future, and Possible Future Individuals and Existing Individuals are cases in which there is an inter-personal conflict choice situation, because in those cases it would be rational in self-interested terms (on the basis of the information that you have about the alternatives) for some of the individuals to choose Alternative 1 (and irrational to choose Alternative 2) if she could make the choice instead of you and rational in self-interested terms for some of the individuals to choose Alternative 2 (and irrational to choose Alternative 1) if she could make the choice instead of you.³⁶

Even if expected utility theory is the most plausible approach towards risk in the context of practical rationality, it would be more controversial to claim that you would not act morally wrong by choosing Alternative 2 in Multi-individual Case under Risk and Unspecificity Combined with Risk-aversiveness than that you would not act morally wrong by choosing Alternative 2 in Multi-individual Case under Risk and Unspecificity. Nevertheless, if morality exists to serve mankind (and sentient beings more generally), I believe it makes sense to hold that in Multi-individual Case under Risk and Unspecificity Combined with Risk-aversiveness you should (from a moral point of view) *take care of* or *look after* the 10,000 individuals in the most reasonable way and thus choose Alternative 2 rather than respect their implausible views about risk or their irrational risk attitudes and choose Alternative 1.

³⁶ On pages 44-45 I cited Williams who talks about "interpersonal conflicts". If I understand Williams correctly, he refers with that to all inter-personal trade-off choice situations.

Limiting the scope of the diminishing marginal moral significance of an individual's utility in this way even more (i.e. limiting the scope of the diminishing marginal moral significance of an individual's utility to those inter-personal trade-off choice situations which are inter-personal trade-off choice situations also at the level of prospects) seems to me, however, arbitrary and inconsistent. Why would only inter-personal conflict choice situations belong to the scope of the diminishing marginal moral significance of an individual's utility? Since ex post prioritarian views claim that prioritarian weightings should be directed to each individual's utility in foreseeable outcomes of different alternatives rather than each individual's expected utility, they are not morally ultimately concerned with each individual's prospects. This means that *from an ex post prioritarian perspective* it should be considered morally irrelevant whether some inter-personal trade-off choice situation is an inter-personal trade-off choice situation both at the level of (foreseeable) outcomes and at the level of prospects or merely at the level of (foreseeable) outcomes. Thus I do not think that the second limitation of the scope of the diminishing marginal moral significance of an individual's utility is tenable. But since the second limitation is untenable, ex post prioritarian views should be rejected because of the untenable implications they can have in such inter-personal trade-off choice situations which are inter-personal trade-off choice situations merely at the level of (foreseeable) outcomes.

I have not yet said anything against the first limitation of the scope of the diminishing marginal moral significance of an individual's utility (i.e. the limitation of the scope of the diminishing marginal moral significance of an individual's utility to inter-personal trade-off choice situations). Now I will. Although the first limitation of the scope of the diminishing marginal moral significance of an individual's utility makes much more sense than the second limitation, it nevertheless seems to me implausible. Consider

Multi-individual Case under Risk III. You can choose between two alternatives. You are not one of the individuals mentioned in the information below that you have.

Alternative 1: A 100 per cent probability that the utility levels of 10,000 specific individuals will be 4. A 90 per cent probability that the utility level of one additional specific individual will be 9 and a 10 per cent probability that her utility level will be 10.

Alternative 2: A 50 per cent probability that the utility levels of the same 10,000 individuals will be 8 and a 50 per cent probability that their utility levels will be 1. A 100 per cent probability that the utility level of the same additional individual will be 9.

Multi-individual Case under Risk II is a case in which the choice situation is not an inter-personal trade-off choice situation at the level of prospects. The choice situation in it is not an inter-personal trade-off choice situation even at the level of (foreseeable) outcomes. The choice situation in it is merely an intra-personal trade-off choice situation. In Multi-individual Case under Risk III the choice situation is, however, an inter-personal trade-off choice situation both at the level of (foreseeable) outcomes and at the level of prospects. However, is it plausible to claim that we should apply a utilitarian principle to Multi-individual Case under Risk II but that we should apply an ex post prioritarian principle to Multi-individual Case under Risk III just because of *one* additional individual in the latter case whose expected (priority-weighted) utility in it is only *a tiny bit* higher in Alternative 1 than in Alternative 2? Are these two cases *so* fundamentally different from a moral point of view that it makes sense to apply *one* criterion of moral rightness and goodness to one of these cases and *another* criterion of moral rightness and goodness to the other case? I am inclined to believe that the correct answer is “no”. Moreover, it seems to me to be against “the spirit” of prioritarianism – especially if prioritarianism is motivated by compassion – that a prioritarian principle comes into play *because of* a single individual whose utility will, with a 100 per cent probability regardless of which alternative you choose, be *higher* than each of the 10,000 individuals’ utility. On the basis of these considerations, it seems to me that an ex post prioritarian cannot plausibly leave other-regarding intra-personal trade-off choice situations outside the scope of the diminishing marginal moral significance of an individual’s utility. But this implies that ex post prioritarians are stuck with applying an ex post prioritarian principle to other-regarding intra-personal trade-off choice situations and thus cannot avoid violating the restricted desert-

sensitive weak ex ante Pareto principle in every possible other-regarding intra-personal trade-off choice situation.

2.4 Conclusion

I believe that (1) (desert-sensitive) ex post prioritarianism (regardless of the priority function) cannot avoid having untenable implications in every possible other-regarding intra-personal trade-off choice situation, as (desert-sensitive) ex post prioritarianism (regardless of the priority function) may prescribe you not to act in a way that maximizes everyone's expected utility and which is the only way to maximize at least someone's expected utility in such an other-regarding intra-personal trade-off choice situation in which the total expected utility of each alternative consists only of existing individuals' expected utility and in which there is no information available to you about the deservingness, or things related to evaluating the deservingness, of any of the individuals whose utility is at stake non-trivially. I also believe that (2) ex post prioritarians cannot plausibly leave other-regarding intra-personal trade-off choice situations outside the scope of the diminishing marginal moral significance of an individual's utility. It also seems to me that (3) (desert-sensitive) ex post prioritarianism (regardless of the priority function) cannot avoid having untenable implications in every possible such inter-personal trade-off choice situation which is an inter-personal trade-off choice situation merely at the level of (foreseeable) outcomes, as (desert-sensitive) ex post prioritarianism (regardless of the priority function) may prescribe you not to act in a way that maximizes everyone's expected utility and which is the only way to maximize at least someone's expected utility in such an inter-personal trade-off choice situation which is an inter-personal trade-off choice situation merely at the level of (foreseeable) outcomes, in which the total expected utility of each alternative consists only of existing individuals' expected utility, and in which there is no information available to you about the deservingness, or things related to evaluating the deservingness, of any of the individuals whose utility is at stake non-trivially. I also believe that (4) ex post prioritarians cannot plausibly leave any inter-

personal trade-off choice situations outside the scope of the diminishing marginal moral significance of an individual's utility.

In order to justifiably reject all forms of ex post prioritarianism, it is sufficient that either (1) and (2) or (3) and (4) are true. I am inclined to believe that (2) is true, quite convinced that (3) is true, very convinced that (1) is true, and extremely convinced that (4) is true.

3 THE REJECTION OF EX ANTE PRIORITARIAN VIEWS

3.1 Introduction

Michael Otsuka & Alex Voorhoeve argue against ex post prioritarianism on the basis of a case in which “a young adult...is now in perfect health but who receives the distressing news that she will soon develop [either a slight impairment or a very severe impairment] and has a 50 percent chance of developing each”.¹ They describe the slight impairment as “a condition that renders it difficult for one to walk more than 2 km” and the very severe impairment as “a condition that leaves one bedridden, save for the fact that one will be able to sit in a chair and be moved around in a wheelchair for part of the day if assisted by others”.² Otsuka & Voorhoeve continue describing the case by writing that

Suppose there is a treatment that is available for each of these conditions, but, in order for it to be effective, it must be taken before it is known which impairment she will suffer. Moreover, she cannot take both treatments. The treatment for the slight impairment would completely eliminate this mild disability: it would restore her to perfect health. It would, however, be completely ineffective against the very severe impairment. The treatment for the very severe impairment would, by contrast, be completely ineffective against the slight impairment but move her up from the highly debilitating condition of the very severe impairment to [a severe impairment which is a] somewhat less debilitating condition[.]³

Otsuka & Voorhoeve describe the severe impairment as “a condition in which one is no longer bedridden; rather, one is able to sit up on one’s own for the entire day but requires the assistance of others to move about.”⁴ Otsuka & Voorhoeve continue describing the case by writing that

¹ Otsuka & Voorhoeve (2009, 171).

² Otsuka & Voorhoeve (2009, 171).

³ Otsuka & Voorhoeve (2009, 172).

⁴ Otsuka & Voorhoeve (2009, 172).

Suppose that you are a morally motivated stranger who could come to the assistance of this young adult by providing her with one or the other treatment. Suppose, further, that you learn that she is indifferent between the two treatments and that you judge her preferences to be an accurate measure of her utility, where “utility” is understood to refer to how well her life is really going (or would go).⁵

Otsuka & Voorhoeve make the following normative conclusion:

Given these suppositions, it would be reasonable for you to share her indifference between these two treatments. Had she preferred the treatment for the slight impairment rather than the very severe impairment, then it would have been reasonable for you to provide her with the treatment for the slight impairment. Had she preferred the treatment for the very severe impairment, then it would have been reasonable for you to provide her with that treatment. In other words, it would be reasonable for you to provide a treatment that maximizes the expected increase in the utility of the recipient. This conclusion is justified at least insofar as this individual is the only person whom you can benefit and you are considering her fate in isolation from any consideration of how well off or badly off anybody else is (yourself included)...when we consider a *single* person in isolation who will develop either the slight impairment or the very severe impairment and who has an equal chance of developing each, it is reasonable to provide her with a treatment that maximizes the expected increase in her utility and therefore that we can reasonably be indifferent between supplying this person with the treatment for the slight impairment and supplying her with the treatment for the very severe impairment when both treatments yield an equally large increase in her expected utility. For on the [Ex Post] Priority View the treatment for the very severe impairment will have a higher expected moral value simply by virtue of the fact that the initial state from which her utility would be increased would be at a lower absolute level.⁶

⁵ Otsuka & Voorhoeve (2009, 173).

⁶ Otsuka & Voorhoeve (2009, 173-174, 178). Otsuka & Voorhoeve (2009, 178) also write that

It follows that someone who holds the [Ex Post] Priority View would conclude that he has decisive moral reason to supply this person with the treatment for the very severe impairment rather than the slight impairment even in some cases in which the person would rationally prefer the treatment for the latter. Such a Prioritarian would conclude that he ought to provide treatment for the very severe impairment even while accepting as correct the person’s own judgment that this treatment would do her less expected good than the treatment for the slight impairment. The [Ex

I also believe that in Otsuka & Voorhoeve's case it would be "reasonable" for you to share the young adult's indifference between the treatments and that it would thus be reasonable, or more specifically, "morally reasonable" (and consequently morally right, or in other words, morally permissible and acceptable), for you to provide her with either of the treatments that maximizes the expected increase in her utility. Since in Otsuka & Voorhoeve's example both treatments maximize the young adult's expected utility, they are incomparable from the point of view of the restricted desert-sensitive weak ex ante Pareto principle. In other words, the restricted desert-sensitive weak ex ante Pareto principle is silent about the moral betterness of the treatments. Nevertheless, I am very convinced that it would not be morally wrong to supply the young adult with the treatment for the slight impairment, which you might do if you believe in ex ante prioritarianism or expected utility utilitarianism and consequently consider both alternatives morally equally good. Thus I am very convinced that ex post prioritarianism (regardless of the priority function), unlike ex ante prioritarianism (combined with any priority function) and expected utility utilitarianism, gives an implausible moral prescription in Otsuka & Voorhoeve's example.

I believe that the young adult has a moral claim on you acting justly and fairly towards her. To put it differently, it seems to me that she has a moral claim on being treated (in a non-medical sense) justly and fairly by you. If you did not supply with any treatment to the young adult, you would ignore these ex ante moral claims of hers and would thus act morally wrong. But since it seems to me that it would not be irrational in self-interested terms (on the basis of the information that you have about the alternatives) for the young adult to choose to receive the treatment for the slight impairment if she rather than you would actually decide which treatment she will receive, I believe that by supplying her with the treatment for the slight impairment you would act justly and fairly towards her and would treat her justly and

Post] Priority View therefore unreasonably mandates provision of a treatment with a lower expected utility for the person concerned.

If it is implausible to hold that in Otsuka & Voorhoeve's case you act morally wrong if you supply the adult with the treatment for the slight impairment, it is even more implausible to hold that in these latter kinds of cases you act morally wrong if you supply the individual with the treatment for the slight impairment.

fairly. It is very hard for me to see how acting in such a way could be morally wrong in this example of Otsuka & Voorhoeve.

However, the passages “[t]his conclusion is justified at least insofar as this individual is the only person whom you can benefit and you are considering her fate in isolation from any consideration of how well off or badly off anybody else is (yourself included)” and “when we consider a *single* person in isolation” hint that according to Otsuka & Voorhoeve, there are moral choice situations in which an alternative that maximizes everyone’s expected utility is not a morally right alternative (i.e. a morally permissible and acceptable alternative) because of *comparative* considerations. Otsuka & Voorhoeve believe in a moral view which puts special moral significance on individuals’ ex post claims in a comparative way:

We might be moved by the thought that, in any scenario involving more than one person, the allocation of resources must be justifiable to each person taken separately in a manner that brings interpersonal considerations to bear that cannot apply in the case of one person considered in isolation. In [multi-person cases] one must justify any claim on resources in light of the comparative strength of the claims of others. Those who are relatively worse off have stronger claims to a given increment of improvement simply by virtue of the fact that it is, other things equal, harder to justify improving the situation of someone who is better off rather than someone who is worse off.⁷

On the basis of this kind of *Competing Claims View*,⁸ Otsuka & Voorhoeve go on to reject ex ante prioritarianism. They do that with the help of a case in which

You are confronted with a group of people, each of whom you know will either develop the very severe or the slight impairment and each of whom has an equal chance of developing either impairment. You also know that half the people will end up suffering from the very severe impairment, and half from the slight impairment. You can either supply everyone with a treatment that will surely improve a recipient’s situation if and only if she turns out to suffer the very severe impairment or supply everyone with a treatment that will surely improve a recipient’s situation if and only if she turns out to suffer the

⁷ Otsuka & Voorhoeve (2009, 183-184).

⁸ That is how Voorhoeve & Fleurbaey (2012, 397-398) call it.

slight impairment. Both treatments, if effective, provide the same increase in utility over nontreatment.⁹

Otsuka & Voorhoeve call this case “Multi-person case with risk and inversely correlated outcomes”.¹⁰ According to Otsuka & Voorhoeve, you should (from a moral point of view) supply everyone with the treatment for the very severe impairment:

From the ex ante perspective of each individual whom we can supply with one of the two treatments, both treatments will be equally valuable. The ex ante view holds that we should therefore be indifferent between providing everyone with the treatment for the slight impairment and providing everyone with the treatment for the very severe impairment. Focusing exclusively on the ex ante perspective, however, means that we fail to take account of the legitimate claims of that half of the group who will, ex post, due to bad brute luck, end up very badly off and worse off than others. Consider, from an ex post perspective, the distributions of well-being that would result from providing everyone with either the treatment for the slight impairment or the treatment for the very severe impairment. If one supplies everyone with the treatment for the slight impairment, then the resulting distribution of utility will be one in which half are in full health, and half are very badly off. If, by contrast, one supplies everyone with the treatment for the very severe impairment, then the resulting distribution will be one in which half are somewhat badly off, and half are badly off (though not as badly off as suffering the very severe impairment without treatment). Ex ante, though we do not know which individuals will end up with the very severe impairment, we do know that whoever does end up with this condition has a stronger claim on the treatment for the very severe impairment than whoever will end up with the slight impairment has on the treatment for the slight impairment. This provides us with a decisive reason to provide everyone with the treatment for the very severe impairment, since we know we will thereby be providing treatment to those who will turn out to be the people who had the strongest claim on it.¹¹

In Multi-person Case with Risk and Inversely Correlated Outcomes both treatments maximize the expected utility of each of the individuals of the group that

⁹ Otsuka & Voorhoeve (2009, 197).

¹⁰ Otsuka & Voorhoeve (2009, 197).

¹¹ Otsuka & Voorhoeve (2009, 197-198).

you are confronted with. Thus also in this case both treatments are incomparable from the point of view of the restricted desert-sensitive weak ex ante Pareto principle. Nevertheless, I am quite convinced that in Multi-person Case with Risk and Inversely Correlated Outcomes it would not be morally wrong to supply everyone with the treatment for the slight impairment, which you might do if you believe in ex ante prioritarianism or expected utility utilitarianism and consequently consider both alternatives morally equally good. Thus I am quite convinced that ex post prioritarianism (regardless of the priority function), unlike ex ante prioritarianism (combined with any priority function) and expected utility utilitarianism, gives an implausible moral prescription also in Multi-person Case with Risk and Inversely Correlated Outcomes.

I believe that Otsuka & Voorhoeve put too much moral importance or emphasis on (foreseeable) outcomes. I think each of the individuals of the group that you are confronted with has a moral claim on you acting justly and fairly towards her. To put it differently, it seems to me that each of them has a moral claim on being treated (in a non-medical sense) justly and fairly by you. I believe that also in Multi-person Case with Risk and Inversely Correlated Outcomes it is these kinds of ex ante claims rather than some kinds of ex post claims that we should be morally concerned with. If you did not supply with any treatment to any of the individuals of the group that you are confronted with, you would ignore the ex ante moral claims of each of them and would thus act morally wrong. But since it seems to me that it would not be irrational in self-interested terms (on the basis of the information that you have about the alternatives) for any of the individuals of the group that you are confronted with to choose to receive the treatment for the slight impairment if she rather than you would actually decide which treatment she will receive, I believe that by providing each of them with the treatment for the slight impairment you would act justly and fairly towards each of them and would treat each of them justly and fairly. I strongly believe that acting in such a way cannot be morally wrong in Multi-person Case with Risk and Inversely Correlated Outcomes. Thus I believe that Otsuka & Voorhoeve's argument against ex ante prioritarianism is implausible and that also in this example of Otsuka & Voorhoeve ex ante prioritarianism does not give an implausible moral prescription.

Since ex ante prioritarianism is compatible with the restricted desert-sensitive weak ex ante Pareto principle, one might think that ex ante prioritarianism is a

more plausible moral view than ex post prioritarianism and utilitarian-adjusted ex post prioritarianism. Moreover, if I am right that ex ante prioritarianism does not give an implausible prescription in Otsuka & Voorhoeve's latter case, it may look like that ex ante prioritarianism, or some sort of desert-sensitive ex ante prioritarianism in case desert should be given intrinsic moral significance, cannot also go astray in such choice situations in which there is a utilitarian tie. Consequently one may think that ex ante prioritarianism or some sort of desert-sensitive ex ante prioritarianism is a plausible moral view.

In the next sub-chapter I will reject another argument that Otsuka has given against ex ante prioritarianism.¹² In Chapter 3.3, however, I will reject ex ante prioritarianism because of the problem it, I believe, faces regarding possible future individuals (i.e. that it cannot include possible future individuals in its scope). Because of that problem, I will formulate a view which I call *ex post-adjusted ex ante prioritarianism*. However, I will reject it because I find it too problematic in the context of possible future individuals (though problematic in a different way than ex ante prioritarianism). I will also formulate a view which I call *utilitarian-adjusted ex ante prioritarianism*, but I will reject it because of a problem it faces regarding certain types of choice situations that involve both possible future individuals and existing individuals. In Chapter 3.4 I will discuss a very different kind of problem, a problem regarding *the dominance principle*, which ex ante prioritarian views face. I believe that ex ante prioritarian views should be rejected also because of that problem.

3.2 Ex Ante Prioritarianism, Competing Ex Post Claims and Possible Future Individuals

After arguing against ex ante prioritarianism with Voorhoeve, Otsuka has discussed ex ante prioritarianism in more detail elsewhere.¹³ Otsuka formulates a case in which

¹² Otsuka (2012).

¹³ Otsuka (2012).

There are two people, each of whom you know will develop either the very severe or the slight impairment and each of whom has an equal chance of developing either impairment. You also know that their risks are inversely correlated: i.e., whenever one of them would suffer the very severe impairment, then the other would suffer the slight impairment. You can either supply both with a treatment that will surely improve a recipient's situation if and only if he turns out to suffer the very severe impairment or supply both with a treatment that will surely improve a recipient's situation if and only if he turns out to suffer the slight impairment. An effective treatment for the slight impairment would provide a somewhat greater increase in utility than would an effective treatment for the very severe impairment.¹⁴

Otsuka calls this case “Two-person case with risk and inversely correlated outcomes”.¹⁵ It is clearly very similar to Otsuka & Voorhoeve's Multi-person Case with Risk and Inversely Correlated Outcomes. The main difference is that in Two-person Case with Risk and Inversely Correlated Outcomes only the provision of the treatment for the slight impairment, rather than a provision of either treatment, would maximize the expected utility of each of the individuals.

Importantly connected to Two-person Case with Risk and Inversely Correlated Outcomes, Otsuka quotes what he wrote with Voorhoeve¹⁶ and what I already quoted in the previous sub-chapter on page 59:

one must justify any claim on resources in light of the comparative strength of the claims of others. Those who are relatively worse off have stronger claims to a given increment of improvement simply by virtue of the fact that it is, other things equal, harder to justify improving the situation of someone who is better off rather than someone who is worse off.¹⁷

Otsuka continues that “the strength of an individual's claim is a function of how much he might benefit in welfare terms, and from what baseline level of well-being, relative to another who might benefit if goods are distributed to him in-

¹⁴ Otsuka (2012, 373-374).

¹⁵ Otsuka (2012, 373).

¹⁶ Otsuka & Voorhoeve (2009, 183-184).

¹⁷ Otsuka (2012, 371).

stead.”¹⁸ If in Two-person Case with Risk and Inversely Correlated Outcomes you supply the individuals with the treatment for the slight impairment, then that individual who turns out to suffer the very severe impairment would, according to Otsuka, have what he calls a *competing-claims-based complaint*.¹⁹ But according to ex ante prioritarianism (regardless of the priority function), you act morally wrong in Two-person Case with Risk and Inversely Correlated Outcomes if you supply the individuals with the treatment for the very severe impairment. Related to this, Otsuka writes that

However things turn out in [Two-person Case with Risk and Inversely Correlated Outcomes], we are certain that one person will be worse off than the other. We can then ask, rhetorically, how can one justify providing a benefit of a given size to someone who will already be better off in order to make him better off still, when one could instead provide nearly as large a benefit to someone else who will be worse off, and who would not even reach the (unimproved) level of the better off person if she (the worse off person) is benefited?²⁰

However, Otsuka realizes that an ex ante prioritarian could respond to this rhetorical question in the following way:

Here’s how we can justify this. In this case, treatment for the slight impairment was in each person’s ex ante self-interest. Our morally motivated stranger can and should take such gambles on their behalves, in order to promote their rational self-interests. The competing-claims-based complaint against administering treatment for the slight impairment is overridden by such an ex ante justification for taking such rational gambles on each person’s behalf.²¹

Thus one could argue that these kinds of justifications that Otsuka calls *ex ante prudential justifications*²² always morally override competing ex post claims and that because of that in Two-person Case with Risk and Inversely Correlated Out-

¹⁸ Otsuka (2012, 371).

¹⁹ Otsuka (2012, 374).

²⁰ Otsuka (2012, 376).

²¹ Otsuka (2012, 377).

²² Otsuka (2012, 377).

comes you should (from a moral point of view) supply the individuals with the treatment for the slight impairment. In Two-person Case with Risk and Inversely Correlated Outcomes there is an ex ante prudential justification available to you for supplying the individuals with the treatment for the slight impairment, as by supplying the individuals with the treatment for the slight impairment you are taking a rational gamble on *each* of those individuals' behalf.

However, Otsuka gives an interesting counter-argument. But before presenting his counter-argument, I should note that in that counter-argument Otsuka talks about how ex ante prioritarianism “ignores the moral significance of the separateness of persons”.²³ I cannot see how ex ante prioritarianism could *ignore* the moral significance of the separateness of persons, as ex ante prioritarianism is distributively sensitive rather than insensitive. The question is rather whether giving priority to worse-off individuals over better-off individuals in the way ex ante prioritarianism does is *a right way* of giving moral significance to the separateness of persons and whether the separateness of persons should be given moral significance (also) *in some other way(s)*. So whenever Otsuka says “ignores the moral significance of the separateness of persons”, the reader should read it something like “goes at least partly astray in terms of how it gives moral significance to the separateness of persons”. Keeping this in mind, let's have a look at what Otsuka writes:

In response, I would point out that, in order to vindicate my charge that the prioritarian ignores the moral significance of the separateness of persons because he is insensitive to the presence or absence of competing claims, I don't need to reject this claim that ex ante prudential justifications override competing ex post claims when the two come into conflict. It would be sufficient, for my purposes, to show that the presence of competing claims provides a reason of some force, even if not necessarily an overriding reason, to provide treatment for the very severe impairment in the case of inversely correlated risks...If competing claims have any positive moral weight, even if these claims are outweighed by an ex ante prudential justification whenever the two come into conflict, then it follows that there is a respect in which prioritarianism ignores the moral significance of the separateness of persons. In order to show that ex ante prioritarianism is vulnerable to the separateness-of-persons objection under discussion, I don't even need

²³ Otsuka (2012, 377, 379).

to establish that competing claims complaints carry any weight *at all* in cases in which there is an *ex ante* prudential justification for a conflicting course of action. All I ultimately need to show is that there are *some* pairs of cases – perhaps cases in which there are no such prudential justifications – in which the presence or absence of competing claims makes a moral difference that the prioritarian fails to register.²⁴

In order to show that there are such pairs of cases and that “*ex ante* prioritarianism is vulnerable to the separateness-of-persons objection”, Otsuka presents a case in which

One person will suffer the slight impairment and another person will suffer the very severe impairment. The person who will suffer the very severe impairment, is, however, either of two possible persons. This is because one’s choice of treatment makes a difference to the identity of the person with the very severe impairment. If you provide the treatment for the slight rather than the very severe impairment, a very severely impaired person will come into existence whose position could not have been improved by treatment for the very severe impairment. It could not have been improved for the following reason: if you provide treatment for the very severe rather than the slight impairment, then a *different* very severely impaired person will come into existence, and his position will be moved from very severe impairment to severe impairment by such treatment. The identity of the slightly impaired person, however, is unaffected by one’s choice of treatment. There is, moreover, no uncertainty regarding the identity of those who would exist and would benefit under each alternative.²⁵

Otsuka calls this case “Two-actual-person case in which either of two possible persons might be very severely impaired”.²⁶ For the sake of brevity, I will call it “Two-actual-person Case Involving Two Possible Persons”. Otsuka also makes it clear that also in this case an effective treatment for the slight impairment would

²⁴ Otsuka (2012, 377).

²⁵ Otsuka (2012, 377-378). Otsuka (2012, 378) adds that “[I]et us assume that the lives of those who are very severely impaired or severely impaired would be worth living and that such people would not regret their existence.”

²⁶ Otsuka (2012, 377).

provide a somewhat greater increase in an individual's utility than an effective treatment for the very severe impairment would provide.²⁷

Since Two-actual-person Case Involving Two Possible Persons is quite complicated, it may be useful to present it also in the following form:

Two-actual-person Case Involving Two Possible Persons. You can supply either with a treatment for a very severe impairment or a treatment for a slight impairment.

The provision of the treatment for the very severe impairment: A very severely impaired individual, individual I, will come into existence. I's position will be moved from very severe impairment to severe impairment thanks to the treatment for the very severe impairment. Thus I's utility will increase because of this treatment. On the other hand, since you supply with the treatment for the very severe impairment, you do not increase the utility of individual K who is slightly impaired.

The provision of the treatment for the slight impairment: A very severely impaired individual, individual J, will come into existence. Since you supply with the treatment for the slight impairment, you do not increase J's utility. On the other hand, since you supply with the treatment for the slight impairment, you increase K's utility. You increase K's utility somewhat more than you would increase I's utility if you supplied with the treatment for the very severe impairment.

Otsuka writes about Two-actual-person Case Involving Two Possible Persons in the following way:

In [Two-actual-person Case Involving Two Possible Persons] there would be no competing claims complaint against the provision of treatment for the slight impairment. This is because, if one provides treatment for the slight impairment, then the person who is thereby brought into existence, suffering the very severe impairment, could not have benefited from the alternative treatment for the very severe impairment. He would therefore have no claim to treatment for the very severe impairment...the merely possible person who could have benefited from the treatment for the very severe impairment

²⁷ Otsuka (2012, 379).

would not have a competing claims complaint against provision of treatment for the slight impairment either, since he has no claim to be brought into existence. Given that no one would have a competing claim to provision of treatment for the very severe impairment in the event that treatment for the slight impairment is provided, it is, intuitively, much less difficult to justify provision of the treatment for the slight impairment in [Two-actual-person Case Involving Two Possible Persons] than in [Two-person Case Involving Certainty].²⁸

Two-person Case Involving Certainty is a case in which

One person will develop the slight impairment and the other will develop the very severe impairment, it is already known who will develop which impairment, and a morally motivated stranger can either offer the one person treatment for the slight impairment or offer the other person treatment for the very severe impairment...an effective treatment for the slight impairment would do somewhat more good than an effective treatment for the very severe impairment.²⁹

Otsuka's conclusion of this pair of cases is that

In [Two-person Case Involving Certainty], by contrast to [Two-actual-person Case Involving Two Possible Persons], there is a strong competing claim to provision of treatment for the very severe impairment, but other moral considerations are equal. The *ex ante* prioritarian, however, registers no difference in the strength of reasons to provide

²⁸ Otsuka (2012, 378). Otsuka (2012, 373) writes that “[o]nly those who either do exist or stand a chance of existing have morally significant claims”. So whenever your choice between two different treatments ensures that someone will not come into existence, then, *after* your choice, that individual that never came into existence because of your choice cannot have a claim to the provision of that treatment that would or might have brought her into existence (see Otsuka 2012, 373). In other words, such an individual that *was* a possible future individual but that never came into existence cannot have claims on anything. It is controversial whether even such individuals that *are* possible future individuals can have any claims, as having a claim seems like having a property and it is doubtful whether possible future individuals can have any properties. I will discuss the question whether possible future individuals can have properties in the next sub-chapter. Nevertheless, it seems to me clear that such individuals that were possible future individuals and that never came into existence cannot have claims on anything.

²⁹ Otsuka (2012, 367).

treatment for the slight impairment versus the severe impairment in these two cases. He thereby ignores the moral significance of the separateness of persons.³⁰

I would like to point out that this pair of cases – or at least Otsuka’s discussion of this pair of cases – does not give any *additional* grounds for believing that competing ex post claims have moral significance. As we already saw, *before* presenting and discussing Two-actual-person Case Involving Two Possible Persons and comparing it to Two-person Case Involving Certainty, Otsuka argues for the moral significance of competing ex post claims by writing that

one must justify any claim on resources in light of the comparative strength of the claims of others. Those who are relatively worse off have stronger claims to a given increment of improvement simply by virtue of the fact that it is, other things equal, harder to justify improving the situation of someone who is better off rather than someone who is worse off.³¹

We also saw that before presenting and discussing Two-actual-person Case Involving Two Possible Persons and comparing it to Two-person Case Involving Certainty, Otsuka argues for the moral significance of competing ex post claims also by asking the following rhetorical question:

[If we are certain that in some particular two-person case one person will be worse off than the other,] how can one justify providing a benefit of a given size to someone who will already be better off in order to make him better off still, when one could instead provide nearly as large a benefit to someone else who will be worse off, and who would not even reach the (unimproved) level of the better off person if she (the worse off person) is benefited?³²

However, when Otsuka compares Two-actual-person Case Involving Two Possible Persons to Two-person Case Involving Certainty, he does not give any further considerations that might back the view that competing ex post claims have moral significance. Instead Otsuka shows that if in Two-actual-person Case Involving

³⁰ Otsuka (2012, 378-379).

³¹ Otsuka (2012, 371).

³² Otsuka (2012, 376).

Two Possible Persons you supply with the treatment for the slight impairment, there is no competing-claims-based ex post complaint that one even *could* appeal to and on the basis of which one *could* hold that you should (from a moral point of view) supply with the treatment for the very severe impairment unlike in Two-person Case Involving Certainty in which such a complaint could be appealed to and on the basis of which one could hold that you should (from a moral point of view) supply with the treatment for the very severe impairment. But the presence or absence of competing-claims-based ex post complaints may not make a difference from a moral point of view, because even when a competing-claims-based ex post complaint could be appealed to, it may be the case that it *should not* be appealed to. This is so because it may be the case that competing ex post claims do not have any moral significance. Related to this, an ex ante prioritarian could argue that

In Two-actual-person Case Involving Two Possible Persons – just like in Two-person Case Involving Certainty – *compassion* points towards the view that supplying with the treatment for the slight impairment is morally wrong, as in both cases compassion points towards the view that it is morally better to benefit a very severely impaired individual than to benefit a slightly impaired individual. I, as an ex ante prioritarian, am morally concerned with compassion and thus hold that in Two-actual-person Case Involving Two Possible Persons – just like in Two-person Case Involving Certainty – you should (from a moral point of view) supply with the treatment for the very severe impairment.

Otsuka fails to explain why ex ante prioritarianism should be rejected and why competing ex post claims should be considered to have any moral significance at all. In order to reject ex ante prioritarianism for the right reasons, we must look elsewhere.

3.3 Possible Future Individuals and the Scope of Ex Ante Prioritarianism

If it is plausible to believe – as I assume it is – that it is not just existing individuals, but also possible future individuals, that have intrinsic moral significance, it is

naturally also plausible to believe that any moral doctrine which does not include possible future individuals in its scope (i.e. any moral doctrine which does not give intrinsic moral significance to possible future individuals) is an untenable moral doctrine. Related to this, expected utility utilitarianism, ex post prioritarianism, utilitarian-adjusted ex post prioritarianism and ex ante prioritarianism face a metaphysical complication. Since a possible future individual does not exist (and has never existed), it arguably cannot have any properties. But if a possible future individual cannot have any properties, it also cannot have the property of expected utility, expected priority-weighted utility or priority-weighted expected utility.³³ How can expected utility utilitarians, ex post prioritarians and ex ante prioritarians then take possible future individuals into account in their moral calculations? In other words, how could possible future individuals be in the scope of these moral views? My discussion about Possible Future Individuals and Existing Individuals on pages 29-30 gives some indication that I believe that expected utility utilitarianism can overcome this metaphysical complication regarding possible future individuals, as I claimed that in Possible Future Individuals and Existing Individuals Alternative 2 rather than Alternative 1 maximizes total expected utility. Thus I believe that possible future individuals can be in the scope of expected utility utilitarianism. I also believe that ex post prioritarianism can overcome the meta-

³³ I believe the view that Holtug (2001, 370) calls “the No Properties of the Non-Existent Principle”, according to which “[a]n individual cannot have any properties if it does not exist”, is true. See also Holtug (1999, 23; 2001, 370-371; 2010, 138). If such individuals that do not exist and have never existed can have properties, then expected utility utilitarianism, ex post prioritarianism, utilitarian-adjusted ex post prioritarianism and ex ante prioritarianism do not face the metaphysical complication which I believe they face. However, it seems to me that possible future individuals might have properties only if the *reductionist* view of personal identity is *not* true. If such a *non-reductionist* view of personal identity, according to which there are souls which are somehow separate from any particular brain and body and which can in some sense exist without and even *before* the existence of some particular brain and body, is true, then the No Properties of the Non-Existent Principle might be false. However, I do not see good grounds for believing in such an extreme non-reductionist view of personal identity. (According to David W. Shoemaker (1999, 184), the reductionist view of personal identity claims that “the facts of personal identity simply consist in more particular facts about brains, bodies, and interrelated physical and mental events” and “non-reductionism is the position that the facts of personal identity consist in some further facts about Cartesian egos, souls, or other kinds of ‘separately existing entity.’”)

physical complication.³⁴ But it seems to me that ex ante prioritarianism cannot. Consider

Possible Future Individuals. You can choose between two alternatives.

Alternative 1: A 100 per cent probability that 10,000 individuals will come into existence and that their utility levels will be 4.

Alternative 2: A 100 per cent probability that 10,000 different individuals than in Alternative 1 will come into existence. A 100 per cent probability that the utility levels of 5,000 unspecific individuals (of these 10,000 individuals) will be 8 and a 100 per cent probability that the utility levels of 5,000 unspecific individuals (of these 10,000 individuals) will be 1.

Let's first look at the metaphysical complication in the context of expected utility utilitarianism. On the one hand, since in Possible Future Individuals you have information about how high utility levels various individuals would or might have as a result of each alternative and with what probabilities, it would not make sense to say that these *alternatives* do not have total expected utilities of 40,000 and 45,000. But on the other hand, if possible future individuals cannot have properties, it must also be the case that in Possible Future Individuals each of the possible future individuals is without expected utility regarding Alternative 1 and Alternative 2 when you are choosing between Alternative 1 and Alternative 2. These considerations imply that both Alternative 1 and Alternative 2 have total expected utilities which do *not* consist of *any* possible future individual's expected utility. In Possible Future Individuals all the possible future individuals are in the scope of expected utility utilitarianism *non-individually* (or indirectly), as although none of them *has* expected utility, some of them *contribute* to the total expected utility of Alternative 1 and some of them contribute to the total expected utility of Alternative 2.³⁵

³⁴ If both expected utility utilitarianism and ex post prioritarianism can overcome the metaphysical complication, then naturally also utilitarian-adjusted ex post prioritarianism can overcome it.

³⁵ Consider the calculation $5,000 \times (1 \times 8) + 5,000 \times (1 \times 1) = 45,000$. This calculation shows that possible future individuals can be in the scope of expected utility utilitarianism non-individually, as this calculation shows that it is possible to calculate the total expected utility of an alternative without first calculating anyone's expected utility. See also Footnote 33 on pages 47-48.

Let's now have a look at the metaphysical complication regarding possible future individuals in the context of ex post prioritarianism and ex ante prioritarianism. If possible future individuals cannot have properties, it must be the case that in Possible Future Individuals each of the possible future individuals is without expected priority-weighted utility and priority-weighted expected utility regarding Alternative 1 and Alternative 2 when you are choosing between Alternative 1 and Alternative 2. Can the same kind of reasoning be used in the context of ex post prioritarianism and ex ante prioritarianism as in the context of expected utility utilitarianism? Could it be argued plausibly that possible future individuals can be in the scope of ex post prioritarianism and ex ante prioritarianism non-individually?

If you are concerned with each alternative's total expected priority-weighted utility or total priority-weighted expected utility, you are concerned with the *distribution* of utility or expected utility. More specifically, you are then concerned with the distribution of utility or expected utility *between different individuals*. In order to distribute something, you must be able to distribute it between some individuals or entities (e.g. groups of individuals).³⁶ You cannot distribute something without distributing it between *some* individuals or entities. But if possible future individuals cannot have properties, then in Possible Future Individuals each of the possible future individuals is without expected utility regarding Alternative 1 and Alternative 2 when you are choosing between Alternative 1 and Alternative 2. In other words, if possible future individuals cannot have properties, you cannot distribute expected utility between different possible future individuals simply because no possible future individual can then have expected utility. This means that you cannot coherently be concerned with the distribution of expected utility between different possible future individuals. To put all this differently and simply, if possible future individuals cannot have properties, you cannot weight any possible future individual's expected utility simply because no possible future individual can then have expected utility. Thus possible future individuals cannot be in the scope of ex ante prioritarianism. This implies that in Possible Future Individuals

³⁶ One could argue that expected utility utilitarianism is an implausible moral view not because it does not give intrinsic moral significance to the distribution of utility or expected utility between different individuals, but because it does not give intrinsic moral significance to the distribution of utility or expected utility between certain kinds of groups of individuals (e.g. between women and men or different ethnic groups). These kinds of views are far beyond the scope of this book.

the moral betterness of Alternative 1 and Alternative 2 cannot be coherently evaluated in an ex ante prioritarian fashion on the basis of the information that you have about the alternatives.

It should be kept in mind, however, that ex post prioritarianism is not concerned with the distribution of expected utility between different individuals. Ex post prioritarianism is concerned with the distribution of utility between different individuals in foreseeable outcomes of different alternatives (and the probabilities of those outcomes). In Possible Future Individuals the total utility of each foreseeable outcome of each alternative consists of 10,000 individuals' utility who are *existing* individuals (rather than possible future individuals) *in those outcomes*, although none of those individuals whose utility that outcome's total utility would consist of exists when you are choosing between Alternative 1 and Alternative 2. Thus the utility of each of those individuals can be weighted and consequently each of them has the property of priority-weighted utility. However, it should be kept in mind that if possible future individuals cannot have properties, none of the *possible future* individuals can have the property of *expected* priority-weighted utility. It follows from these considerations that in Possible Future Individuals all the possible future individuals are in the scope of ex post prioritarianism but are in the scope of it non-individually in the same way as they are in the scope of expected utility utilitarianism non-individually.

One might suggest that ex ante prioritarianism should be combined with ex post prioritarianism so that an ex ante prioritarian principle is applied to existing individuals and an ex post prioritarian principle is applied to possible future individuals. I call this kind of view ex post-adjusted ex ante prioritarianism. However, although ex post-adjusted ex ante prioritarianism can include possible future individuals in its scope and can never violate the restricted desert-sensitive weak ex ante Pareto principle, I believe it is too problematic in the context of possible future individuals. Consider

Possible Future Individuals II. You can choose between two alternatives.

Alternative 1: A 100 per cent probability that 10,000 individuals will come into existence and that their utility levels will be 4.

Alternative 2: A 100 per cent probability that the same 10,000 individuals as in Alternative 1 will come into existence.³⁷ A 100 per cent probability that the utility levels of 5,000 unspecific individuals (of these 10,000 individuals) will be 8 and a 100 per cent probability that the utility levels of 5,000 unspecific individuals (of these 10,000 individuals) will be 1.

According to ex post-adjusted ex ante prioritarianism (combined with the square root weightings), Alternative 1 is a morally better alternative than Alternative 2. This is so because $10,000 \times (1 \times \sqrt{4}) = 20,000$ and $5,000 \times (1 \times \sqrt{8}) + 5,000 \times (1 \times \sqrt{1}) \approx 19,142$. If it is implausible to hold that possible future individuals can have properties, we cannot plausibly hold that by choosing Alternative 1 in Possible Future Individuals II you would deny the better prospect (in terms of expected utility) from each of the possible future individuals. If possible future individuals cannot have properties, they also cannot have the property of a prospect. On the other hand, in Multi-individual Case under Risk and Unspecificity (which I presented on page 47 and which is identical with Possible Future Individuals II except that the individuals in it are existing individuals rather than possible future individuals) you would deny the better prospect (in terms of expected utility) from each of the 10,000 individuals by choosing Alternative 1. But ex post-adjusted ex ante prioritarianism clearly cannot prescribe you to choose Alternative 1 in Multi-individual Case under Risk and Unspecificity, because ex post-adjusted ex ante prioritarianism says that an ex ante prioritarian principle rather than an ex post prioritarian principle should be applied to existing individuals.

However, even if in Possible Future Individuals II you do not deny the better prospect (in terms of expected utility) from any of the possible future individuals

³⁷ I assume that the identities of the possible future individuals are not fixed at the time you are choosing between the alternatives. So although the same individuals will come into existence regardless of which alternative you choose, it is not the case that it is already determined which individuals will come into existence. This implies that at the time you are choosing between the alternatives the possible future individuals that will in fact come into existence are merely possible future individuals rather than necessary future individuals that will necessarily come into existence. It seems to me that it would be more controversial to claim that necessary future individuals cannot have properties than that merely possible future individuals cannot have properties, as the ontological status of necessary future individuals seems to me to be something between possible future individuals and existing individuals.

by choosing Alternative 1, I believe it makes sense to say that it would be rational in self-interested terms (on the basis of the information that you have about the alternatives) for each of the possible future individuals to choose Alternative 2 (and irrational for each of them to choose Alternative 1) if it could make the choice instead of you. Claiming this does not seem to me metaphysically absurd or meaningless. On the basis of this, I believe it makes sense to say that by choosing Alternative 1 you would act unjustly and unfairly towards each of those possible future individuals even if none of them can have properties.³⁸ Thus it seems to me that ex post-adjusted ex ante prioritarianism is an implausible moral view.³⁹ On the basis of these considerations, I believe that we should accept what may be called

The Desert-sensitive Rationality-orientated Moral View: If in some choice situation your utility is at stake only trivially or is not at stake at all and it would be rational in self-interested terms (on the basis of the information that you have about the alternatives) for everyone else to choose some particular alternative and irrational in self-interested terms to choose any other alternative (where “everyone else” refers to each of those other individuals, including also possible future individuals, that contributes to the total expected utility of at least one of the alternatives in that choice situation) if she could make the choice instead of you, it must be so that you act morally wrong if you do not choose the alternative which it would be rational in self-interested terms for everyone else to choose, provided that there is no information available to you about the deservingness, or things related to evaluating the deservingness, of any of those individuals.

Since ex post-adjusted ex ante prioritarianism violates the Desert-sensitive Rationality-orientated Moral View, one might argue that

³⁸ On the other hand, it would seem a bit odd to me to claim that by choosing Alternative 1 you would treat each of those possible future individuals unjustly and unfairly, as it seems to me that in order to treat some individual in some particular way, that individual must exist.

³⁹ Nevertheless, if possible future individuals cannot have properties, it seems to me clear that ex post-adjusted ex ante prioritarianism is a less implausible moral view than ex ante prioritarianism, because at least ex post-adjusted ex ante prioritarianism can include possible future individuals in its scope.

Because of the problems of ex post prioritarianism and ex post-adjusted ex ante prioritarianism, we should not be ex post prioritarians even regarding possible future individuals. Consequently, if we should be prioritarians, we should be ex ante prioritarians and not at all ex post prioritarians. But if it is true that a possible future individual cannot have the property of a prospect, the natural conclusion is that we should apply an ex ante prioritarian principle only to existing individuals and that we should apply a utilitarian principle to possible future individuals (in a non-individual fashion).

The form of prioritarianism, according to which an ex ante prioritarian principle should be applied to existing individuals and a utilitarian principle to possible future individuals, may be called utilitarian-adjusted ex ante prioritarianism. However, it seems to me that the following example gives good grounds for believing that utilitarian-adjusted ex ante prioritarianism should be rejected:

Possible Future Individuals and Existing Individuals II. You can choose between two alternatives. You are not one of the individuals mentioned in the information below that you have.

Alternative 1: A 100 per cent probability that 5,000 individuals will come into existence and that their utility levels will be 1. A 100 per cent probability that the utility levels of 5,000 specific (and already existing) individuals will be 9.

Alternative 2: A 100 per cent probability that the same 5,000 individuals as in Alternative 1 will come into existence and that their utility levels will be 9. A 100 per cent probability that the utility levels of the same 5,000 already existing individuals will be 1.

It seems to me that the alternatives are morally equally good, as the only foreseeable outcome of Alternative 2 (i.e. the outcome in which 5,000 individuals' utility levels are 9 and 5,000 individuals' utility levels are 1) differs from the only foreseeable outcome of Alternative 1 (i.e. the outcome in which 5,000 individuals' utility levels are 1 and 5,000 individuals' utility levels are 9) only in the sense that in the former there is a *permutation* of individuals' utility levels. In other words, since the *pattern* of the individuals' utility levels is exactly the same in the only foreseeable outcome of Alternative 2 as in the only foreseeable outcome of Alternative 1, and since there is no information available to you about the deservingness,

or things related to evaluating the deservingness, of any of the 5,000 existing individuals, it seems to me that there is no good reason to consider one of the alternatives morally better than the other. But rather than claiming that the alternatives are morally equally good, utilitarian-adjusted ex ante prioritarianism (combined with the square root weightings) claims that Alternative 2 is a morally better alternative than Alternative 1. This is so because $5,000 \times (1 \times 1) + 5,000 \times \sqrt{(1 \times 9)} = 20,000$ and $5,000 \times (1 \times 9) + 5,000 \times \sqrt{(1 \times 1)} = 50,000$. Thus utilitarian-adjusted ex ante prioritarianism is incompatible with what may be called

The desert-sensitive utility-based permutation principle: If the arrangement of individuals' utility levels in some foreseeable outcome of Alternative 2 is simply a permutation of individuals' utility levels in some foreseeable outcome of Alternative 1, and if the probabilities of both foreseeable outcomes are the same, it must be the case that the foreseeable outcome of Alternative 1 contributes as much to the total moral value of Alternative 1 as the foreseeable outcome of Alternative 2 contributes to the total moral value of Alternative 2, provided that there is no information available to a chooser about the deservingness, or things related to evaluating the deservingness, of any of the individuals.⁴⁰

The desert-sensitive utility-based permutation principle seems to me a plausible moral principle. Since utilitarian-adjusted ex ante prioritarianism is incompatible with the desert-sensitive utility-based permutation principle, I believe that utilitarian-adjusted ex ante prioritarianism should be rejected. It is also worth noticing that also ex post-adjusted ex ante prioritarianism is incompatible with the desert-sensitive utility-based permutation principle. Consider the following case in which the only foreseeable outcome of Alternative 2 (i.e. the outcome in which 2,500 individuals' utility levels are 7, 2,500 individuals' utility levels are 1 and 5,000 individuals' utility levels are 4) differs from the only foreseeable outcome of Alternative 1 (i.e. the outcome in which 5,000 individuals' utility levels are 4, 2,500 individuals' utility levels are 7 and 2,500 individuals' utility levels are 1) only in the sense that in the former there is a permutation of individuals' utility levels:

⁴⁰ I have developed this formulation partly on the basis of Adler (2012, 52).

Possible Future Individuals and Existing Individuals III. You can choose between two alternatives. You are not one of the individuals mentioned in the information below that you have.

Alternative 1: A 100 per cent probability that 5,000 individuals will come into existence and that their utility levels will be 4. A 100 per cent probability that the utility levels of 2,500 unspecific (and already existing) individuals of 5,000 specific individuals will be 7 and that the utility levels of 2,500 unspecific (and already existing) individuals of the same 5,000 specific individuals will be 1.

Alternative 2: A 100 per cent probability that the same 5,000 individuals as in Alternative 1 will come into existence. A 100 per cent probability that the utility levels of 2,500 unspecific individuals (of these 5,000 individuals) will be 7 and a 100 per cent probability that the utility levels of 2,500 unspecific individuals (of these 5,000 individuals) will be 1. A 100 per cent probability that the utility levels of the same 5,000 already existing individuals will be 4.

Rather than claiming that the alternatives are morally equally good, ex post-adjusted ex ante prioritarianism (combined with the square root weightings) claims that Alternative 1 is a morally better alternative than Alternative 2. This is so because $5,000 \times (1 \times \sqrt{4}) + 5,000 \times \sqrt{(0.5 \times 7 + 0.5 \times 1)} = 20,000$ and $2,500 \times (1 \times \sqrt{7}) + 2,500 \times (1 \times \sqrt{1}) + 5,000 \times \sqrt{(1 \times 4)} \approx 19,114$.

3.4 The Dominance Principle

Let's imagine that I am mistaken about whether possible future individuals can have properties and that possible future individuals can in fact have properties. If that is the case, it seems to me that ex ante prioritarianism is a more plausible moral view than ex post-adjusted ex ante prioritarianism and utilitarian-adjusted ex ante prioritarianism, as ex ante prioritarianism avoids those problems of ex post-adjusted ex ante prioritarianism and utilitarian-adjusted ex ante prioritarianism which I discussed in the previous sub-chapter. However, ex ante prioritarian views have a serious structural problem. Thus even if possible future individuals can have properties, ex ante prioritarianism should be rejected. Consider what may be called

The dominance principle: If in some moral choice situation every foreseeable outcome of some alternative is better than every foreseeable outcome of some other alternative, then the former alternative is a morally better alternative than the latter alternative.⁴¹

Ex ante prioritarianism (and consequently also ex post-adjusted ex ante prioritarianism and utilitarian-adjusted ex ante prioritarianism) is incompatible with the dominance principle. This incompatibility can be illustrated with the following example:

Multi-individual Case under Risk and Unspecificity II. You can choose between two alternatives. You are not one of the individuals mentioned in the information below that you have.

Alternative 1: A 100 per cent probability that the utility levels of 10,000 specific individuals will be 4. A 50 per cent probability that the utility levels of 100 specific individuals will be 9 and a 50 per cent probability that their utility levels will be 10.

Alternative 2: A 100 per cent probability that the utility levels of 5,000 unspecific individuals (of the same 10,000 individuals) will be 8 and a 100 per cent probability that the utility levels of 5,000 unspecific individuals (of the same 10,000 individuals) will be 1. A 50 per cent probability that the utility levels of the same 100 individuals will be 9 and a 50 per cent probability that their utility levels will be 5.

According to ex ante prioritarianism (combined with the square root weightings), Alternative 2 is a morally better alternative than Alternative 1. This is so because $10,000 \times \sqrt{1 \times 4} + 100 \times \sqrt{0.5 \times 9 + 0.5 \times 10} \approx 20,308$ and $10,000 \times \sqrt{0.5 \times 8 + 0.5 \times 1} + 100 \times \sqrt{0.5 \times 9 + 0.5 \times 5} \approx 21,478$. However, by holding that Alternative 2 is a morally better alternative and that the goodness of outcomes is also determined in a prioritarian fashion, an ex ante prioritarian must hold that the alternative whose every foreseeable outcome is (according to her) better than every foreseeable outcome of the other alternative is a morally worse alternative. Thus

⁴¹ See Adler (2012, 495); Rabinowicz (2001, 150; 2002, 11).

her view about the moral betterness of Alternative 1 and Alternative 2 is incompatible with the dominance principle. This is so because by holding that the goodness of outcomes is determined in a prioritarian fashion, the value of the first foreseeable outcome of Alternative 1 comes out as 20,300 ($10,000 \times \sqrt{4} + 100 \times \sqrt{9} = 20,300$) and the value of the second foreseeable outcome of Alternative 1 approximately as 20,316 ($10,000 \times \sqrt{4} + 100 \times \sqrt{10} \approx 20,316$), whereas the value of the first foreseeable outcome of Alternative 2 comes out approximately as 19,442 ($5,000 \times \sqrt{8} + 5,000 \times \sqrt{1} + 100 \times \sqrt{9} \approx 19,442$) and the value of the second foreseeable outcome of Alternative 2 approximately as 19,366 ($5,000 \times \sqrt{8} + 5,000 \times \sqrt{1} + 100 \times \sqrt{5} \approx 19,366$).⁴²

I argued in Chapter 2 that if the Morality for Mankind View is a plausible view, it implies that the restricted desert-sensitive weak ex ante Pareto principle should be accepted. But it also seems to me that if the Morality for Mankind View is a plausible view, a moral doctrine which is incompatible with the dominance principle must be an implausible moral doctrine. If morality exists to serve mankind, how

⁴² More generally, it can be shown that ex ante prioritarianism combined with any prioritarian priority function together with a prioritarian criterion of the goodness of outcomes is incompatible with the dominance principle (see Adler 2012, 509-510). See also Rabinowicz (2001, 150-151, 163; 2002, 11-12, 14-15).

Naturally also Multi-individual Case under Risk and Unspecificity can be used to illustrate that ex ante prioritarianism is incompatible with the dominance principle. Ex ante prioritarianism (combined with the square root weightings) says that in Multi-individual Case under Risk and Unspecificity Alternative 2 is a morally better alternative than Alternative 1. By holding that in Multi-individual Case under Risk and Unspecificity Alternative 2 is a morally better alternative than Alternative 1 and that the goodness of outcomes is also determined in a prioritarian fashion, an ex ante prioritarian must hold that the alternative whose every foreseeable outcome (in that particular case the only foreseeable outcome) is (according to her) better than every foreseeable outcome (in that particular case the only foreseeable outcome) of the other alternative is a morally worse alternative. This is so because $10,000 \times \sqrt{(1 \times 4)} = 20,000$ and $10,000 \times \sqrt{(0.5 \times 8 + 0.5 \times 1)} \approx 21,213$ and because $10,000 \times \sqrt{4} = 20,000$ and $5,000 \times \sqrt{8} + 5,000 \times \sqrt{1} \approx 19,142$. Multi-individual Case under Risk and Unspecificity II demonstrates that ex ante prioritarianism combined with the view that the goodness of outcomes is also determined in a prioritarian fashion is incompatible with the dominance principle not only in the context of such inter-personal trade-off choice situations which are inter-personal trade-off choice situations merely at the level of (foreseeable) outcomes but also in the context of such inter-personal trade-off choice situations which are inter-personal trade-off choice situations also at the level of prospects.

could such an alternative whose every foreseeable outcome is better than every foreseeable outcome of some other alternative be a morally worse alternative than the latter? I think it could not.

There is a dilemma for a prioritarian, as a prioritarian has to choose between (utilitarian-adjusted) ex post prioritarianism which is incompatible with the restricted desert-sensitive weak ex ante Pareto principle (and which is consequently incompatible also with the Desert-sensitive Rationality-orientated Moral View) and ex ante prioritarianism which is incompatible with the dominance principle. In other words, a prioritarian has to discard either the restricted desert-sensitive weak ex ante Pareto principle (and the Desert-sensitive Rationality-orientated Moral View) or the dominance principle, although none of them should be discarded.

One might suggest, however, that ex ante prioritarianism should be combined with the view that the goodness of outcomes is determined in a utilitarian rather than prioritarian fashion. Related to this, Parfit writes that

the Priority View can take either Telic or Deontic forms. It can be a view about which outcomes would be better, or a view that is only about what we ought to do.⁴³

Ex ante prioritarianism combined with a utilitarian criterion of the goodness of outcomes is an example of deontic prioritarianism, whereas ex ante prioritarianism combined with a prioritarian criterion of the goodness of outcomes is an example of telic prioritarianism. Ex ante prioritarianism combined with a utilitarian criterion of the goodness of outcomes does not violate the dominance principle in Multi-individual Case under Risk and Unspecificity II. This is so because according to a utilitarian criterion of the goodness of outcomes, in Multi-individual Case under Risk and Unspecificity II the value of the first foreseeable outcome of Alternative 1 is 40,900 ($10,000 \times 4 + 100 \times 9 = 40,900$), the value of the second foreseeable outcome of Alternative 1 is 41,000 ($10,000 \times 4 + 100 \times 10 = 41,000$), the value of the first foreseeable outcome of Alternative 2 is 45,900 ($5,000 \times 8 + 5,000 \times 1 + 100 \times 9 = 45,900$) and the value of the second foreseeable outcome of Alternative 2 is 45,500 ($5,000 \times 8 + 5,000 \times 1 + 100 \times 5 = 45,500$). By holding that Alternative 2 is a morally better alternative and that the goodness of outcomes is determined in a

⁴³ Parfit (1995, 101).

utilitarian fashion, deontic ex ante prioritarianism does not violate the dominance principle in Multi-individual Case under Risk and Unspecificity II.

It should be noticed, however, that deontic ex ante prioritarianism is not compatible with the dominance principle. This can be illustrated for example with Multi-individual Case under Certainty which I presented on page 14. According to ex ante prioritarianism (if the weightings are done with a square root), in Multi-individual Case under Certainty Alternative 1 is a morally better alternative than Alternative 2. This is so because $10,000 \times \sqrt{(1 \times 4)} = 20,000$ and $5,000 \times \sqrt{(1 \times 7)} + 5,000 \times \sqrt{(1 \times 1.5)} \approx 19,352$. But according to a utilitarian criterion of the goodness of outcomes, the value of the only foreseeable outcome of Alternative 1 is 40,000 ($10,000 \times 4 = 40,000$) and the value of the only foreseeable outcome of Alternative 2 is 42,500 ($5,000 \times 7 + 5,000 \times 1.5 = 42,500$). Thus by holding that Alternative 1 is a morally better alternative and that the goodness of outcomes is determined in a utilitarian fashion, deontic ex ante prioritarianism violates the dominance principle in Multi-individual Case under Certainty.

Moreover, we could question whether deontic prioritarianism makes much sense in the first place. Would it not be quite arbitrary to hold that the goodness of each outcome is determined by its total utility and that the correct criterion of the moral goodness of acts is prioritarian? If the distinction between individuals should not be given special importance in a prioritarian fashion in the evaluations of the goodness of outcomes, why should it actually be given special moral importance in a prioritarian fashion in the moral evaluations of acts? A deontic ex ante prioritarian needs to tell what is the relevant difference between outcomes and acts that justifies having a utilitarian attitude towards outcomes and a prioritarian attitude towards acts. If she is unable to do that, telic ex ante prioritarianism makes more sense than deontic ex ante prioritarianism.

3.5 Conclusion

The problems of ex ante prioritarianism, ex post-adjusted ex ante prioritarianism and utilitarian-adjusted ex ante prioritarianism that I identified and discussed in Chapter 3.3 and Chapter 3.4 naturally apply also to desert-sensitive versions of

them. In Chapter 3.3 I accepted the metaphysical view that possible future individuals cannot have properties, but I did not discuss it comprehensively. Instead I concentrated on the moral philosophical implications of this view regarding expected utility utilitarianism, ex post prioritarianism, utilitarian-adjusted ex post prioritarianism and ex ante prioritarianism. I believe that it has such implications regarding ex ante prioritarianism because of which ex ante prioritarianism should be rejected. But even if possible future individuals can have properties, I strongly believe that all kinds of ex ante prioritarian views should be rejected because of the dominance principle.

4 PRIORITARIANISM, IMPARTIALITY AND IMPERSONALITY

4.1 Introduction

Because of certain specific problems of ex post prioritarianism, utilitarian-adjusted ex post prioritarianism, ex ante prioritarianism, ex post-adjusted ex ante prioritarianism and utilitarian-adjusted ex ante prioritarianism that I have discussed (and which also apply to desert-sensitive versions of each of them), I believe that neither any kind of pure nor any kind of hybrid prioritarianism is the way to go. In this main chapter I discuss the question whether prioritarianism violates impartiality. If prioritarianism violates impartiality, there may be even more reason to believe that all pure and hybrid prioritarian views should be rejected.

Impartiality is a very tricky concept in moral philosophy. It is far from clear exactly what conditions are necessary or sufficient for a moral doctrine to fulfil in order to qualify as a moral doctrine which does not violate impartiality (i.e. in order to qualify as a moral doctrine which is completely impartial). Moreover, there is a very difficult normative question related to impartiality. Is impartiality a property which a moral doctrine must have in order to be a plausible moral doctrine? So if prioritarianism violates impartiality, should we reject prioritarianism because of that?

My own impression is that most moral philosophers think that impartiality is a property which a moral doctrine must have in order to be a plausible moral doctrine.¹ Here I cannot discuss in detail whether they are right. Thus I will not discuss much the question whether prioritarianism should be rejected if it is true that it violates impartiality. I will also not discuss the very difficult question how exactly impartiality should be understood in the context of morality. What makes that question so difficult is that there are numerous structurally totally different kinds of moral doctrines whose moral frameworks differ from each other completely. Instead I will concentrate on the question whether prioritarianism violates impartiality or not, which I believe can be found out without having an idea how exactly

¹ See also e.g. Boran (2004, 333); Singer (1979/1993, 10-11; 1979/2011, 10-11).

impartiality should be understood in the context of morality. I will argue that prioritarianism violates impartiality. My own intuitions say that prioritarianism should be rejected because of that, but I will not refer to or give any argument to back these intuitions of mine, because I am not aware of any good argument for impartiality in morality² and because I personally have also not managed to come up with such an argument. But my intuitions say that such an argument could be developed. Nevertheless, even if my intuitions are mistaken and it is not relevant regarding the plausibility of prioritarianism whether prioritarianism violates impartiality or not, I find this question interesting in its own right. By exploring this question, we can understand prioritarianism better on a structural level, even if it is normatively insignificant what the correct answer to that question is.

4.2 Prioritarianism Violates Impartiality: The First Argument

Although it is far from clear exactly what conditions are necessary or sufficient for a moral doctrine to fulfil in order to qualify as a moral doctrine which does not violate impartiality, it is nevertheless clear that utilitarianism does not violate impartiality. It is clear that a purely utility-based moral doctrine³ is a completely impartial moral doctrine if it claims that *every* individual's utility, expected utility and/or utility chances count(s) morally and count(s) morally *equally*. Actual consequence utilitarianism and expected utility utilitarianism are examples of such moral doctrines.⁴ It also seems to me clear that any such moral doctrine which is partly

² However, it seems to me that such arguments have been developed that give good grounds for believing that *impersonality* is a property which a moral doctrine must have in order to be a plausible moral doctrine. I will discuss such an argument in Chapter 6. In the next sub-chapter I will distinguish impartiality from impersonality.

³ By a purely utility-based moral doctrine I refer to any such moral doctrine whose "currency" or currencies, and only currency or currencies, is/are utility, expected utility and/or utility chances. For example, rights-based and deontological moral doctrines are not purely utility-based moral doctrines. In the context of some such moral doctrines which are not purely utility-based or which are not at all utility-based impartiality may come out as something quite different than in the context of purely utility-based moral doctrines. I will discuss utility chances in Chapter 5.

⁴ If it is true that possible future individuals cannot have expected utility, then naturally no possible future individual can have expected utility which counts morally. However, according to expected

utility-based and partly desert-based is completely impartial if it claims that every individual's utility, expected utility and/or utility chances, and every individual's deservingness, count morally and count morally equally. Thus desert-sensitive versions of utilitarianism also qualify as completely impartial moral doctrines.

It should also be noticed that since utilitarianism claims that every individual's utility or expected utility counts morally and counts morally equally, a utilitarian criterion of moral rightness and goodness does not include proper names or other expressions that refer to particular individuals. It seems to me clear that any purely utility-based moral doctrine (and any partly utility-based and partly desert-based moral doctrine), whose criterion of moral rightness and goodness includes proper names or other expressions that refer to particular individuals, violates impartiality. For example, how could such an egoistic moral doctrine, according to which a moral agent's act is morally right if and only if it maximizes *her* expected utility, qualify as an impartial moral doctrine?⁵

It is much less clear whether a purely utility-based moral doctrine is completely impartial if *and only if* it claims that every individual's utility, expected utility and/or utility chances count(s) morally and count(s) morally equally (and whether a partly utility-based and partly desert-based moral doctrine is completely impartial if and only if it claims that every individual's utility, expected utility and/or utility

utility utilitarianism, the utility of each individual in some foreseeable outcome of an alternative counts morally together with its probability and counts morally together with its probability as much, regardless of whether that individual exists when a moral agent is making a choice between different alternatives.

It should also be noticed that by impartiality and partiality I refer here only to *formal* impartiality and partiality. These should be distinguished from what may be called *substantial* impartiality and partiality. According to utilitarianism, one should not put extra weight on anyone's, not even on her own or her loved ones' utility or expected utility (or interests, as some would prefer to say) on the formal level of moral reasoning. But of course concentrating on taking care of one's own and one's loved ones' interests may in many cases be justified from the point of view of utilitarianism (see Gruen 1999, 139-140). For example, as it is commonly acknowledged, one usually knows better how to help her loved ones effectively than how to help some complete strangers effectively. It is also likely to be the case that concentrating more on helping some sick children of strangers than satisfying the interests of one's own children would cause guilt and anxiety in the parent (see Singer 1999, 300). There are arguably also many other reasons why in many cases the formal impartiality of utilitarianism permits moral agents to be substantially partial.

⁵ This paragraph is partly based on Adler (2012, 23, 52, 57).

chances, and every individual's deservingness, count morally and count morally equally). But if that is what it takes to be completely impartial, then no form of prioritarianism qualifies as completely impartial. Although prioritarianism, as I understand it, says that *every* individual's utility or expected utility counts morally,⁶ it clearly does not say that every individual's utility or expected utility counts morally *equally*. This is so because prioritarianism is morally more concerned with equal changes or differences in an individual's utility or expected utility, the lower her/his/its utility or expected utility is. This means that prioritarianism, unlike utilitarianism, is not committed to what may be called *equal consideration of individuals*. Because of these considerations, it seems to me that prioritarianism is impartial to certain extent but not completely. I think it makes sense to say that prioritarianism is impartial in *scope* (except for ex ante prioritarianism if it is true that it cannot include possible future individuals in its scope) but not in *consideration*.⁷

One could argue that prioritarianism is completely impartial because although prioritarianism is morally more concerned with equal changes or differences in an individual's utility or expected utility, the lower her/his/its utility or expected utility is, prioritarianism claims that the moral significance of an individual's utility or expected utility does *not* depend on *whose* or *which particular individual's* utility or expected utility it happens to be.⁸ Along the same lines, one might argue that prioritarianism does not violate impartiality, because a prioritarian criterion of moral rightness and goodness does not include proper names or other expressions that refer to particular individuals. One could also argue that of different forms of prioritarianism only those forms of prioritarianism violate impartiality that are incompatible with the desert-sensitive utility-based permutation principle, and that consequently ex post prioritarianism, utilitarian-adjusted ex post prioritarianism and ex ante prioritarianism (unlike ex post-adjusted ex ante prioritarianism and utilitarian-adjusted ex ante prioritarianism) do not violate impartiality.

⁶ Of course, one could believe in such *restricted* prioritarianism according to which only persons' or people's utility or expected utility counts morally, though that would seem quite arbitrary to me.

⁷ Since I believe that ex ante prioritarianism cannot include possible future individuals in its scope, it seems to me that ex ante prioritarianism is not impartial even in scope.

⁸ See Hooker (2000, 60-61); McCarthy (2008, 6). Hooker (2000, 61-62), however, is inclined to believe that prioritarianism violates impartiality. I will present his argument in the next sub-chapter.

I believe, however, that the arguments in the previous paragraph confuse impartiality with *impersonality*. On the basis of those arguments, there may be good grounds for believing that prioritarianism or certain forms of prioritarianism is/are completely impersonal, but I do not think that any of those arguments gives good grounds for believing that prioritarianism or certain forms of prioritarianism is/are completely impartial. Let me explain by starting from Bykvist who writes that

One problem with prioritarianism is that, on the face of it, it does not seem to square with impartiality, one of the cornerstones of utilitarianism. If I give more weight to worse off people, I seem to be biased in favour of some people over others, and that does not look like impartiality. Compare: if I give more weight to the well-being of the rich and famous, I seem to...show bias towards some people over others.⁹

But Bykvist continues that

it is not clear that this is a good objection. Strictly speaking, it is not true that prioritarianism favours worse off people over better off people. Prioritarianism is only concerned with the absolute well-being levels of individuals; it is not concerned with whether people are better or worse off than others...[For a prioritarian] the value of a life does not depend on how it fares in comparison to other lives; it only depends on the absolute well-being level of the life. I can be better off than you in one situation, and worse off than you in another, but if my absolute level stays the same in both situations, its weight will also be the same.¹⁰

I believe, however, that this argument does not give much grounds for believing that prioritarianism does not violate impartiality. Instead it may give good grounds for believing that prioritarianism does not violate impartiality as *directly* as certain comparative moral views do. It seems to me that prioritarianism does not favour worse-off individuals (worse-off individuals in terms of utility or expected utility) over better-off individuals (better-off individuals in terms of utility or expected utility) as directly as such comparative moral views like the rank-weighted priority

⁹ Bykvist (2010, 71).

¹⁰ Bykvist (2010, 71).

view¹¹ and leximin (of which the former favours worse-off individuals over better-off individuals non-lexically and the latter lexically) do. I will come back to the partiality of prioritarianism very shortly. Before that, I illustrate the partiality of the rank-weighted priority view and leximin and the difference between impartiality and impersonality with the following example:

Strange Referee. There is a football referee who, when she is refereeing, always consciously favours to some extent the team that is behind in goals. The way how she adjudicates is affected by how many goals the teams have scored compared to each other rather than how many goals each team has scored in an absolute sense. Whenever the score is even, she does not favour either team. She also never favours any team on the basis that it is some particular team.

It seems to me clear that this strange referee adjudicates in an impersonal way, as she never favours any team on the basis that it is some particular team. However, *from the point of view of adjudicating*, this referee clearly is a partial referee, as she always favours, on the basis of comparative considerations, the team that is behind in goals (i.e. the team that has scored less goals in the match than its opponent). It seems to me that in a similar fashion, *from a moral point of view*, the rank-weighted priority view and leximin violate impartiality,¹² because although they do not favour any individuals on the basis that they are some particular individuals, they favour (either non-lexically or lexically), on the basis of comparative considerations, those individuals who are “behind” in terms of utility or expected utility (i.e. those individuals whose utility or expected utility is lower). If the strange

¹¹ As I already said on page 11, according to the rank-weighted priority view, the moral significance of an individual’s utility or expected utility depends non-lexically on how high it is in an ordinal sense. See Footnote 13 on pages 91-92 which clarifies what this means.

¹² I do not think there is any perspective-independent impartiality. In other words, if something or someone is impartial, it/she/he is impartial *from some point of view* (e.g. from a moral point of view or from the point of view of adjudicating) rather than in some more general, perspective-independent, way.

referee is not impartial *as an adjudicator*, how could we consistently claim that the rank-weighted priority view and leximin are impartial *as moral views*?¹³

¹³ Even if a moral doctrine does not have to be a completely impartial moral doctrine in order to be a plausible moral doctrine, I strongly believe that the rank-weighted priority view and leximin are highly implausible moral views. I already rejected leximin on pages 19-20. Related to the implausibility of the rank-weighted priority view, consider

Multi-individual Case under Certainty IV. You can choose between two alternatives. You are not one of the individuals mentioned in the information below that you have.

Alternative 1: A 100 per cent probability that the utility levels of 5,000 specific individuals will be 6 and a 100 per cent probability that the utility levels of 5,000 specific individuals will be 4.

Alternative 2: A 100 per cent probability that the utility levels of 5,000 specific individuals (those individuals whose utility levels would end up being 6 if you chose Alternative 1) will be 10 and a 100 per cent probability that the utility levels of 5,000 specific individuals (those individuals whose utility levels would end up being 4 if you chose Alternative 1) will be 1.

The most obvious priority function for the (probabilistic) rank-weighted priority view – and perhaps the only non-arbitrary priority function for the rank-weighted priority view – is to give the weighting 1 (i.e. no weighting) to the highest utility level of an individual in the foreseeable outcomes of an alternative (or to the highest expected utility level of an individual in an alternative), the weighting 2 to the second highest utility level of an individual in the foreseeable outcomes of that alternative (or to the second highest expected utility level of an individual in that alternative), the weighting 3 to the third highest utility level of an individual in the foreseeable outcomes of that alternative (or to the third highest expected utility level of an individual in that alternative), and so on. (For example giving the weighting 1 to the highest utility level of an individual in the foreseeable outcomes of an alternative, the weighting 3 to the second highest utility level of an individual in the foreseeable outcomes of that alternative, the weighting 5 to the third highest utility level of an individual in the foreseeable outcomes of that alternative, and so on, would seem quite arbitrary to me.) The rank-weighted priority view combined with the most obvious rank-weighted priority function says that in Multi-individual Case under Certainty IV Alternative 1 is a morally better alternative than Alternative 2. This is so because $5,000 \times (1 \times 6 \times 1) + 5,000 \times (1 \times 4 \times 2) = 70,000$ and $5,000 \times (1 \times 10 \times 1) + 5,000 \times (1 \times 1 \times 2) = 60,000$.

It should be noticed, however, that the rank-weighted priority view implies the view that some individual's utility in a foreseeable outcome of an alternative (or some individual's expected utility in an alternative) may contribute more to that alternative's total moral value than some other individual's *higher* utility in that foreseeable outcome of that alternative (or some other individual's higher expected utility in that alternative). This can be illustrated with Multi-individual Case under Certainty IV in which those individuals whose utility levels would end up being 4 as a result of you choosing Alternative 1 would each, according to the rank-weighted priority view, contribute more to

Due to the fact that the way how the strange referee adjudicates is affected by comparative considerations, the previous case does not give grounds for believing that prioritarianism violates impartiality. It could be said that the rank-weighted priority view and leximin favour worse-off individuals over better-off individuals because they quite *explicitly* claim that worse-off individuals should be favoured over better-off individuals. Prioritarianism, on the other hand, does not explicitly claim that worse-off individuals should be favoured over better-off individuals. However, I believe that prioritarianism violates impartiality because it seems to me that prioritarianism *indirectly* favours worse-off individuals over better-off individuals due to its view that an individual's utility or expected utility has diminishing marginal moral significance.¹⁴ The partiality of prioritarianism can be illustrated with the following example:

Strange Judge. The final of a slam dunk contest is about to start. Sam and Mick have reached the final. In the final both competitors have two attempts. There are three judges. Each of them gives points to each competitor from each attempt from 0 to 10 rounded to the nearest tenth. The competitor who gets a higher total of points from two attempts from the judges wins the contest. Sam's first attempt is extremely good, perhaps fantastic. Sam gets 9.5 from Judge A, 9.6 from Judge

the total moral value of Alternative 1 than each of those individuals whose utility level would end up being 6 as a result of you choosing Alternative 1. This is so because $1 \times 4 \times 2 = 8$ and $1 \times 6 \times 1 = 6$. I find this absurd. Prioritarianism, of course, does not imply such a view and is not even compatible with such a view. According to prioritarianism, an individual's utility or expected utility has diminishing marginal moral significance. But prioritarianism also says that someone's higher utility in a foreseeable outcome of an alternative or someone's higher expected utility in an alternative nevertheless always contributes more to that alternative's total moral value than someone else's lower utility in that foreseeable outcome of that alternative or someone else's lower expected utility in that alternative.

¹⁴ It should be noticed that although egalitarianism is a comparative moral view in the sense that it is concerned with how high utility or expected utility levels different individuals have compared to each other, it does not favour worse-off individuals over better-off individuals as directly as the rank-weighted priority view and leximin do. This is so because egalitarianism, unlike the rank-weighted priority view and leximin, does not explicitly claim that worse-off individuals should be favoured over better-off individuals. However, it seems to me that egalitarianism indirectly favours worse-off individuals over better-off individuals due to its view that it is intrinsically morally significant how equally utility or expected utility is distributed between different individuals.

B but only 3.1 from Judge C. Sam's second attempt goes completely wrong and he gets 1.0 from Judge A, 1.1 from Judge B and 1.0 from Judge C. Both of Mick's attempts are mediocre. From his first attempt Mick gets 5.0 from Judge A, 4.9 from Judge B but only 2.2 from Judge C. From his second attempt Mick gets 4.9 from Judge A, 5.0 from Judge B but only 2.2 from Judge C. Both Judge A and Judge B give those points to each competitor from each attempt which they think best reflect the athletic excellence shown by the competitor. Judge C, however, gives those points to each competitor from each attempt which is the square root (rounded to the nearest tenth) of the points she thinks would best reflect the athletic excellence shown by the competitor. Judge C thought that in Sam's first attempt 9.6 would best reflect the athletic excellence shown by him and that in Sam's second attempt 1.0 would best reflect the athletic excellence shown by him. Judge C thought that in both of Mick's attempts 4.9 would best reflect the athletic excellence shown by Mick.

If Judge C had given each competitor those points from each attempt which she thought would best reflect the athletic excellence shown, she would have given a higher total of points to Sam, as she would have given 10.6 points in total to Sam and 9.8 points in total to Mick. But since Judge C gave more marginal value (in terms of points) to athletic excellence on lower levels of athletic excellence than on higher levels of athletic excellence, she ended up giving more points in total to Mick (4.4 points in total) than to Sam (4.1 points in total).

On the one hand, it seems to me clear that Judge C adjudicated in an impersonal way, as she did not favour either athlete on the basis that he was some particular individual. But on the other hand, it seems to me that Judge C was a partial judge (from the point of view of adjudicating), because the way how she adjudicated (i.e. giving diminishing marginal value to athletic excellence in terms of points) was that kind of adjudicating which is biased towards consistent athletes and against inconsistent athletes. In other words, the way how Judge C adjudicated was that kind of adjudicating which can favour consistent athletes over inconsistent athletes.¹⁵ In a similar fashion, it seems to me that prioritarianism violates impartiality

¹⁵ Let's imagine that in the final between Sam and Mick both competitors would have had only one attempt and that Judge C would have adjudicated in the same absurd fashion. Would Judge C have


(from a moral point of view), as it is biased towards worse-off individuals and against better-off individuals due to its view that an individual's utility or expected utility has diminishing marginal moral significance. Because of this kind of bias, prioritarianism favours some individuals over some other individuals at some level and in certain kinds of choice situations (in the case of ex post prioritarianism at the level of foreseeable outcomes in inter-personal trade-off choice situations and in the case of ex ante prioritarianism at the level of prospects in those kinds of inter-personal trade-off choice situations which are inter-personal trade-off choice situations also at the level of prospects and not merely at the level of outcomes), although it never does so because of comparative considerations and never favours any individuals on the basis that they are some particular individuals.

Figure 4 illustrates three different judges who give points to athletes in a strictly concave fashion (i.e. three different judges who give diminishing marginal value to athletic excellence in terms of points). Figure 5, on the other hand, illustrates a judge who gives points to athletes in a linear fashion.

been partial as an adjudicator in that case? I would be inclined to say "yes", because although in that case Judge C could not have favoured the athlete who was more consistent in the final (simply because in that case neither competitor could have been more consistent than the other in the final nor could they have been equally consistent in the final, as it does not make sense to talk about one's consistency and inconsistency regarding slam dunks in the context of one slam dunk), she nevertheless would have adjudicated in a way that is formally or inherently biased towards consistent athletes and against inconsistent athletes.

FIGURE 4. Giving points in a strictly concave fashion.


Points given by a judge



Athletic excellence shown by an athlete
as perceived by a judge

FIGURE 5. Giving points in a linear fashion.

Points given by a judge



Athletic excellence shown by an athlete
as perceived by a judge

If it is true that Judge C was not completely impartial as an adjudicator, how could we consistently claim that prioritarianism is completely impartial as a moral view? I think there are strong grounds for believing that Judge C did not adjudicate impartially (from the point of view of adjudicating) and that prioritarianism consequently violates impartiality (from a moral point of view). One might claim that Judge C was *incompetent* rather than partial. But since it looks like that Judge C

understood what kinds of points would best reflect the athletic excellence shown by the athletes, it seems to me that she was not incompetent. I think the way how Judge C adjudicated could be characterized well by saying that she judged in an impersonal but partial way.

4.3 Prioritarianism Violates Impartiality: The Second Argument

I think there is also another good reason to believe that prioritarianism violates impartiality. Hooker writes that


Some people hold in Nietzschean fashion that the priority should be not towards the worst off but rather towards the most accomplished and powerful. On this view, the interests of the worst off and indeed of people who are living only average lives should be sacrificed whenever this would help the great artists, writers, and generals to ascend even higher pinnacles. Those who hold this view might claim that it is an impartial view. They might point out that they are not partial towards any particular individual, but rather give priority to whichever one can achieve most. And which one can achieve most changes as people's powers develop or deteriorate... Take an even more familiar form of elitism. Some people want even greater amounts going to the richest royalty, movie stars, singers, and athletes. This is in effect the view that the very richest should get richer, even at the cost to the rest of us... such elitist principles... certainly are *not* impartial. The fact that elitist views are not impartial presents a clear problem for the priority for the worst off. If various forms of priority for those at the top is a rejection of impartiality, then how can priority for those at the bottom be impartial?¹⁶

Related to this passage, we could consider such a utility-based moral view whose moral significance function leads to a *convex* moral significance line regarding an individual's utility or expected utility (i.e. such a moral view according to which an individual's utility or expected utility has *increasing* marginal moral significance, see Figure 6). According to my intuitions, such a moral view, call it *inverse prioritarianism*, clearly violates impartiality. It actually seems to me intuitively clearer that inverse prioritarianism violates impartiality than that regular prioritarianism

¹⁶ Hooker (2000, 61-62).

violates impartiality. However, if inverse prioritarianism violates impartiality, how could we consistently claim that regular prioritarianism does *not* violate impartiality? It seems to me clear that we could not. Since inverse prioritarianism violates impartiality, also regular prioritarianism violates impartiality.¹⁷

FIGURE 6. Convex moral significance lines.



4.4 Normative Considerations regarding Impartiality

It is arguably much less controversial to claim that in each sports contest each adjudicator (including each referee, judge and jury member) should (from the point of view of minimizing the odds that the official result of that sports contest will not reflect well the athletic excellence shown by the athletes or teams in that contest) be completely impartial than it is to claim that a moral doctrine should be rejected if it violates impartiality. It could be argued (plausibly, I believe) that one of the structural goals or built-in ideals of each sports contest is to provide an official result which reflects at least reasonably well the athletic excellence shown by those athletes or teams that participated in that contest.¹⁸ By using Nicholas Dixon's

¹⁷ I am not aware of any moral philosopher who has argued for inverse prioritarianism, but that does not matter here.

¹⁸ See e.g. Berman (2011) and Dixon (1999, 116) who think along these lines. The term "structural goal" comes from Sigmund Loland (2002, 10).

terminology, we could call those sports contests which fail to provide such results “failed athletic contests”.¹⁹ It is obvious that in order to minimize the occurrence of failed athletic contests, adjudicators in sports contests should be completely impartial as adjudicators. Whether a moral doctrine must be completely impartial in order to be a plausible moral doctrine is less clear to me, although it seems to me (on the basis of my intuitions) that it must.

4.5 Conclusion

In this main chapter I have given two arguments (one fairly complex argument and one very straightforward argument) that prioritarianism violates impartiality. I consider both arguments plausible. If a moral doctrine must be completely impartial in order to be a plausible moral doctrine and if either of the arguments that prioritarianism violates impartiality is plausible, prioritarianism should be rejected.

¹⁹ Dixon (1999, 116).

5 GIVING FAIR CHANCES TO INDIVIDUALS

On the face of it, I find the general idea that either an individual's utility (in foreseeable outcomes of different alternatives) or an individual's expected utility (in different alternatives) has diminishing marginal moral significance intuitively quite attractive. Nevertheless, on the basis of chapters 2-4, I find this idea so problematic that I believe that it should be rejected. My rejection of this idea implies not only the rejection of prioritarianism but also the rejection of sufficientarianism, as sufficientarianism is identical with prioritarianism up to some specific utility or expected utility level of an individual.

I have also rejected egalitarianism (including also prigualitarianism), the rank-weighted priority view and leximin. I have also rejected Otsuka & Voorhoeve's Competing Claims View and argued that Otsuka fails to explain why competing ex post claims should be considered to have any moral significance at all. However, it should not be thought that I am not prepared to give intrinsic moral significance to the distribution of utility or expected utility between different individuals, as I have made it clear that I am inclined to believe that desert should be given intrinsic moral significance. It seems to me that how utility or expected utility should (from a moral point of view) be distributed between different individuals also depends intrinsically on desert considerations rather than only on which alternative maximizes total expected utility. Thus I am inclined to believe that there is more plausibility in some kind of desert-sensitive expected utility utilitarianism than in pure expected utility utilitarianism. However, there are very different ways of giving intrinsic moral significance to desert and it is unclear to me in what way(s) we should do that.

Moreover, I am prepared to give intrinsic moral significance to the distribution of *utility chances* between different individuals.¹ I feel compelled to do that, because I find expected utility utilitarianism – and different versions of desert-sensitive expected utility utilitarianism – extremely problematic in one particular type of *utilitarian tie*. I also find ex post prioritarianism and utilitarian-adjusted ex post

¹ On chance-orientated (or lottery-based) moral views which give intrinsic moral significance to the distribution of utility chances between different individuals, see e.g. Hirose (2004; 2007); Timmermann (2004).

prioritarianism (and desert-sensitive version of them) equally problematic in that particular type of utilitarian tie. Thus I think the following example gives additional grounds for believing that those forms of prioritarianism should be rejected:

*Multi-individual Utilitarian Tie under Risk.*² You can choose between two alternatives. You are not one of the individuals mentioned in the information below that you have.

Alternative 1: A 100 per cent probability that the utility levels of 5,000 specific individuals will be 7 and a 100 per cent probability that the utility levels of 5,000 specific individuals will be 1.

Alternative 2: A 50 per cent probability that the utility levels of 5,000 specific individuals (those individuals whose utility levels would end up being 7 if you chose Alternative 1) will be 7 and that the utility levels of 5,000 specific individuals (those individuals whose utility levels would end up being 1 if you chose Alternative 1) will be 1. A 50 per cent probability that the utility levels of 5,000 specific individuals (those individuals whose utility levels would end up being 7 if you chose Alternative 1) will be 1 and that the utility levels of 5,000 specific individuals (those individuals whose utility levels would end up being 1 if you chose Alternative 1) will be 7.

According to ex ante prioritarianism (regardless of the priority function), ex post-adjusted ex ante prioritarianism (regardless of the priority function) and utilitarian-adjusted ex ante prioritarianism (regardless of the priority function), Alternative 2 is a morally better alternative than Alternative 1. If the weightings are done with a square root, the total moral value of Alternative 1 comes out approximately as 18,229 and the total moral value of Alternative 2 as 20,000. This is so because $5,000 \times \sqrt{1 \times 7} + 5,000 \times \sqrt{1 \times 1} \approx 18,229$ and $10,000 \times \sqrt{0.5 \times 7 + 0.5 \times 1} = 20,000$. But according to expected utility utilitarianism, ex post prioritarianism (regardless of the priority function) and utilitarian-adjusted ex post prioritarianism (regardless of the priority function), the alternatives are morally equally good. Thus they are indifferent between the alternatives. According to expected utility utilitarianism, the total moral values of both alternatives are 40,000. This is so because

² This example is based on Peter A. Diamond's (1967) famous example against utilitarianism.

$5,000 \times (1 \times 7) + 5,000 \times (1 \times 1) = 40,000$ and $10,000 \times (0.5 \times 7 + 0.5 \times 1) = 40,000$. According to ex post prioritarianism (combined with the square root weightings) and utilitarian-adjusted ex post prioritarianism (combined with the square root weightings), the total moral values of both alternatives come out approximately as 18,229. This is so because $5,000 \times (1 \times \sqrt{7}) + 5,000 \times (1 \times \sqrt{1}) \approx 18,229$ and $10,000 \times (0.5 \times \sqrt{7} + 0.5 \times \sqrt{1}) \approx 18,229$.

It seems to me quite clear that in Multi-individual Utilitarian Tie under Risk Alternative 2 is a morally better alternative than Alternative 1. But why do I think so? I have rejected the view that an individual's expected utility has diminishing marginal moral significance. I have also rejected the view that it is intrinsically morally significant how equally expected utility is distributed between different individuals. I have also rejected the ex ante versions of the rank-weighted priority view and leximin. I believe that in Multi-individual Utilitarian Tie under Risk Alternative 2 is a morally better alternative than Alternative 1, because the former seems to me to have a fairer distribution of utility chances between the 10,000 individuals. Unlike Alternative 1, Alternative 2 gives each of the 10,000 individuals a chance to the highest possible utility level that she could (on the basis of the information that is available to you) get as a result of your choice. Related to this, consider also

Multi-individual Utilitarian Tie under Certainty. You can choose between two alternatives. You are not one of the individuals mentioned in the information below that you have.

Alternative 1: A 100 per cent probability that the utility levels of 10,000 specific individuals will be 4.

Alternative 2: A 100 per cent probability that the utility levels of 5,000 specific individuals (of the same 10,000 individuals) will be 7 and a 100 per cent probability that the utility levels of 5,000 specific individuals (of the same 10,000 individuals) will be 1.

According to ex post prioritarianism (regardless of the priority function) and utilitarian-adjusted ex post prioritarianism (regardless of the priority function), Alternative 1 is a morally better alternative than Alternative 2. If the weightings are done with a square root, the total moral value of Alternative 1 comes out as 20,000

and the total moral value of Alternative 2 approximately as 18,229. This is so because $10,000 \times (1 \times \sqrt{4}) = 20,000$ and $5,000 \times (1 \times \sqrt{7}) + 5,000 \times (1 \times \sqrt{1}) \approx 18,229$. Also according to ex ante prioritarianism (regardless of the priority function), ex post-adjusted ex ante prioritarianism (regardless of the priority function) and utilitarian-adjusted ex ante prioritarianism (regardless of the priority function), Alternative 1 is a morally better alternative than Alternative 2. If the weightings are done with a square root, the total moral value of Alternative 1 comes out as 20,000 and the total moral value of Alternative 2 approximately as 18,229. This is so because $10,000 \times \sqrt{(1 \times 4)} = 20,000$ and $5,000 \times \sqrt{(1 \times 7)} + 5,000 \times \sqrt{(1 \times 1)} \approx 18,229$. According to expected utility utilitarianism, on the other hand, Alternative 1 and Alternative 2 are morally equally good. Thus it is indifferent between the alternatives. According to expected utility utilitarianism, the total moral values of both alternatives are 40,000. This is so because $10,000 \times (1 \times 4) = 40,000$ and $5,000 \times (1 \times 7) + 5,000 \times (1 \times 1) = 40,000$.

I believe, however, that the morally best thing to do in Multi-individual Utilitarian Tie under Certainty would be to toss a coin about which alternative to choose rather than to choose directly one of the alternatives, except if you do not have enough time for tossing a coin or do not have a coin and cannot arrange a lottery between the alternatives in any other way. By tossing a coin you would give each of the 10,000 individuals a chance to the highest possible utility level that she could get as a result of your choice. That seems to me a fairer thing to do than choosing directly one of the alternatives.

I am inclined to believe that some sort of *desert-sensitive fair chance expected utility utilitarianism* is the way to go. However, I believe that total expected utility has so much more moral significance than a fair distribution of utility chances between different individuals that the latter should come into play as a moral consideration (in terms of determining what one should morally do) only in utilitarian ties. (More specifically, I believe that it should come into play only in such utilitarian ties in which there is such an inter-personal trade-off choice situation which is an inter-personal trade-off choice situation at the level of prospects.) I can illustrate this with the following example:

Multi-individual Case under Certainty V. You can choose between two alternatives. You are not one of the individuals mentioned in the information below that you have.

Alternative 1: A 100 per cent probability that the utility levels of 9,999 specific individuals will be 7 and a 100 per cent probability that the utility level of one specific individual will be 1.

Alternative 2: A 100 per cent probability that the utility levels of the same 9,999 individuals will be 1 and a 100 per cent probability that the utility level of the same single individual will be 7.

It seems to me that from the point of view of the fairness of the distribution of utility chances between the 10,000 individuals, you should toss a coin about which alternative to choose and thus give each of the individuals a 50 per cent chance to the highest possible utility level that she could get as a result of your choice. On the other hand, if one thinks that what a moral agent should (from a moral point of view) do in different moral choice situations is a function of the total expected utility of each alternative and the fairness of the distribution of utility chances between different individuals, she might think that what you should (from a moral point of view) do in Multi-individual Case under Certainty V is to arrange such a lottery (provided that you can) in which you give one of the individuals a 0.01 per cent chance to the highest possible utility level that she could get as a result of your choice and in which you give each of the 9,999 individuals a 99.99 per cent chance to the highest possible utility level that each of them could get as a result of your choice.

I believe, however, that the morally best thing to do in Multi-individual Case under Certainty V would be to choose directly Alternative 1 without organizing any lottery. It seems to me clear that in Multi-individual Case under Certainty V every foreseeable outcome of Alternative 1 (i.e. the only foreseeable outcome of Alternative 1) is better than every foreseeable outcome of Alternative 2 (i.e. the only foreseeable outcome of Alternative 2), regardless of whether you arrange a lottery or not and what the result of that lottery would be. Because of this it seems to me that it cannot be morally wrong to choose Alternative 1 even if choosing it is in conflict with the result of the lottery that you arranged. In order not to endorse a moral doctrine which is incompatible with the dominance principle, it seems to me

that a fair distribution of utility chances between different individuals should come into play as a moral consideration (in terms of determining what one should morally do) only in utilitarian ties.

I think an important question is whether desert-sensitive fair chance expected utility utilitarianism is a completely impartial moral view. Obviously I do not believe that desert-sensitive fair chance expected utility utilitarianism violates impartiality, as otherwise I would not support it. I cannot see how Multi-individual Utilitarian Tie under Certainty could give grounds for believing that desert-sensitive fair chance expected utility utilitarianism violates impartiality, as by doing in Multi-individual Utilitarian Tie under Certainty what desert-sensitive fair chance expected utility utilitarianism prescribes you to do (i.e. tossing a coin), you clearly avoid favouring anyone. However, one might argue that

Multi-individual Utilitarian Tie under Risk gives good grounds for believing that desert-sensitive fair chance expected utility utilitarianism violates impartiality, as in Multi-individual Utilitarian Tie under Risk it prescribes you to choose Alternative 2 and not to choose Alternative 1. Thus in Multi-individual Utilitarian Tie under Risk desert-sensitive fair chance expected utility utilitarianism – just like ex ante prioritarianism, ex post-adjusted ex ante prioritarianism and utilitarian-adjusted ex ante prioritarianism – in fact favours those individuals whose utility level will definitely end up being 1 if you choose Alternative 1. Only such a moral doctrine which prescribes you to arrange a lottery regarding which alternative to choose both in Multi-individual Utilitarian Tie under Risk and Multi-individual Utilitarian Tie under Certainty or which is indifferent between the alternatives both in Multi-individual Utilitarian Tie under Risk and Multi-individual Utilitarian Tie under Certainty can be a completely impartial moral doctrine, as only such a moral doctrine avoids favouring some individuals over some other individuals in *both* cases.

It seems to me, however, that Multi-individual Utilitarian Tie under Risk does not give grounds for believing that desert-sensitive fair chance expected utility utilitarianism violates impartiality, because in Multi-individual Utilitarian Tie under Risk desert-sensitive fair chance expected utility utilitarianism does *not* favour any individuals *in terms of giving them chances to the highest possible utility level* that each of them could get as a result of your choice. The weighting element of prioritarianism is concerned with the utility level or the expected utility level of each

individual. The fair chance element of desert-sensitive fair chance expected utility utilitarianism is, however, concerned with different individuals' *chances* to their highest possible utility levels. It seems to me that the *outlook* of prioritarianism on utility or expected utility is partial, as it favours worse-off individuals (worse-off individuals in terms of utility or expected utility) over better-off individuals (better-off individuals in terms of utility or expected utility) in inter-personal trade-off choice situations, or in the case of ex ante prioritarian views, in certain kinds of inter-personal trade-off choice situations. But since the fair chance element of desert-sensitive fair chance expected utility utilitarianism is concerned with different individuals' chances to their highest possible utility levels, what determines whether desert-sensitive fair chance expected utility utilitarianism violates impartiality is whether its outlook on giving different individuals chances to their highest possible utility levels is impartial. Since desert-sensitive fair chance expected utility utilitarianism holds that every individual's chances to her highest possible utility level count morally equally, I do not see any grounds for believing that the outlook of desert-sensitive fair chance expected utility utilitarianism on giving different individuals chances to their highest possible utility levels is somehow partial. On the other hand, such an outlook on utility chances, that claims that in Multi-individual Utilitarian Tie under Risk you should (from the point of view of the fairness of the distribution of utility chances and from a moral point of view) toss a coin about which alternative to choose, is partial, as instead of giving each of the 10,000 individuals a 50 per cent chance to the highest possible utility level that each of them could get as a result of your choice, it gives a 75 per cent chance to 5,000 individuals to the highest possible utility level that each of them could get as a result of your choice and a 25 per cent chance to 5,000 individuals to the highest possible utility level that each of them could get as a result of your choice. According to such a chance-orientated view, not every individual's utility chances count equally.³

³ Such fair chance expected utility utilitarianism according to which the fair chance element should come into play (in terms of determining what one should morally do) not merely in utilitarian ties but also in such moral choice situations like the one in Multi-individual Case under Certainty V, and which claims that in Multi-individual Case under Certainty V you should (from a moral point of view) arrange such a lottery (provided that you can) in which you give one of the individuals a 0.01 per cent chance to the highest possible utility level that she could get as a result of your choice and

One might argue against desert-sensitive fair chance expected utility utilitarianism in the following way that consists of three parts. First, one could argue that it is not the case that impartiality is a property which a moral doctrine must have in order to be a plausible moral doctrine. Secondly, one could argue that a moral doctrine cannot be plausible if it is incompatible with the following principle that states a necessary condition for a moral doctrine to be such which gives priority to worse-off individuals (worse-off individuals in terms of expected utility) over better-off individuals (better-off individuals in terms of expected utility):

The desert-sensitive ex ante Pigou-Dalton principle: If in some moral choice situation individual I's expected utility in Alternative 1 is lower than individual J's expected utility in Alternative 1, if in that choice situation I's expected utility in Alternative 2 is lower than or equal to J's expected utility in Alternative 2, if in that choice situation I's expected utility is higher in Alternative 2 than in Alternative 1, if in that choice situation J's expected utility is lower in Alternative 2 than in Alternative 1, and if in that choice situation the difference in I's expected utility in Alternative 2 and in Alternative 1 is the same as the difference in J's expected utility in Alternative 1 and in Alternative 2, it must be the case that Alternative 2 is a morally better alternative than Alternative 1, provided that the total expected utility of each of these alternatives consists only of I's and J's ex-

in which you give 9,999 individuals a 99.99 per cent chance to the highest possible utility level that each of them could get as a result of your choice, is not incompatible with the view that every individual's utility chances count equally and that the correct outlook on utility chances is impartial. Since in Multi-individual Case under Certainty V there is *no* utilitarian tie, one could hold that the chances should be 0.01 per cent to one individual and 99.99 per cent to 9,999 individuals rather than 50 per cent to all of them without giving up the view that every individual's utility chances count equally and that the correct outlook on utility chances is impartial, because if one believes that both utility chances and the total amount of expected utility count equally *compared to each other*, the obvious conclusion is that the chances should be 0.01 per cent and 99.99 per cent rather than 50 per cent. In other words, the question whether every individual's utility chances count equally and the question whether every individual's expected utility counts equally are related to impartiality, whereas the question whether both utility chances and the total amount of expected utility count equally compared to each other is not related to impartiality but instead to the question which of these moral considerations is a more important moral consideration or whether they are equally important moral considerations.

pected utility and there is no information available to a chooser about I's or J's deservingness or things related to evaluating I's or J's deservingness.⁴

Thirdly, one could argue that in order to avoid the problems of such forms of prioritarianism, the rank-weighted priority view and leximin which are compatible with the desert-sensitive ex ante Pigou-Dalton principle, and in order not to endorse such a moral doctrine (e.g. desert-sensitive fair chance expected utility utilitarianism) which is incompatible with the desert-sensitive ex ante Pigou-Dalton principle,⁵ one should endorse what may be called *desert-sensitive priority-catering expected utility utilitarianism*. According to desert-sensitive priority-catering expected utility utilitarianism, there is a good moral reason to give priority to worse-off individuals (worse-off individuals in terms of expected utility) over better-off individuals (better-off individuals in terms of expected utility). But desert-sensitive priority-catering expected utility utilitarianism also says that there is even a better or stronger moral reason to maximize total expected utility. On the basis of these views, desert-sensitive priority-catering expected utility utilitarianism says that priority should be given to worse-off individuals (worse-off individuals in terms of expected utility) over better-off individuals (better-off individuals in terms of expected utility) in utilitarian ties (but only in utilitarian ties) in the same fashion as ex ante leximin gives priority to worse-off individuals (worse-off individuals in terms of expected utility) over better-off individuals (better-off individuals in terms of expected utility) regardless of whether there is a utilitarian tie or not.

One problem of desert-sensitive priority-catering expected utility utilitarianism is, I believe, that it violates impartiality in the same way in such inter-personal trade-off choice situations which are inter-personal trade-off choice situations at the level of prospects and in which there is a utilitarian tie as ex ante leximin violates impartiality in all such inter-personal trade-off choice situations which are inter-personal trade-off choice situations at the level of prospects. But what if impartiality is not a property which a moral doctrine necessarily must have in order to be a plausible moral doctrine?

⁴ This definition is based on Adler's (2012, 502) definition of the ex ante Pigou-Dalton principle.

⁵ Desert-sensitive fair chance expected utility utilitarianism clearly violates the desert-sensitive ex ante Pigou-Dalton principle in Multi-individual Utilitarian Tie under Certainty.

Even if I leave impartiality considerations aside, I find the idea of giving fair chances to individuals in utilitarian ties more plausible than the idea of giving priority to worse-off individuals over better-off individuals in utilitarian ties. I think the fair chance element of desert-sensitive fair chance expected utility utilitarianism gives the most natural explanation for why in Multi-individual Utilitarian Tie under Risk Alternative 2 is a morally better alternative than Alternative 1. It seems to me that the moral betterness of Alternative 2 in Multi-individual Utilitarian Tie under Risk is directly connected to the fact that by choosing Alternative 1 you would fail to give half of the 10,000 individuals a chance to the highest possible utility level that they could get as a result of your choice and to the fact that you could give such a chance to each of the 10,000 individuals without doing that at the expense of total expected utility. A supporter of the desert-sensitive ex ante Pigou-Dalton principle is unable to explain the moral betterness of Alternative 2 in Multi-individual Utilitarian Tie under Risk in this way, as if she explains the rationale of the desert-sensitive ex ante Pigou-Dalton principle through the idea that in such inter-personal trade-off choice situations which are inter-personal trade-off choice situations at the level of prospects and in which there is a utilitarian tie all individuals should (from a moral point of view) be given a chance to the highest possible utility level that they could (on the basis of the information that is available to a chooser) get as a result of a chooser's choice, she contradicts herself. This is so because in Multi-individual Utilitarian Tie under Certainty the desert-sensitive ex ante Pigou-Dalton principle prescribes you to choose Alternative 1 rather than to give each of the 10,000 individuals a chance to the highest possible utility level that she could get as a result of your choice.

I find it somewhat strange that Matthew D. Adler, who is attracted to the strong ex ante Pareto principle and very attracted to the dominance principle and the general idea of giving priority to worse-off individuals over better-off individuals, and who recognizes the incompatibility of ex post prioritarianism with the strong ex ante Pareto principle and the incompatibility of ex ante prioritarianism with the dominance principle, nevertheless ends up endorsing (responsibility-sensitive) ex post prioritarianism without even considering any such moral view that gives priority to worse-off individuals over better-off individuals only in utilitarian ties.⁶

⁶ Adler (2012).

However, having said this, it is actually very hard to find such arguments in the literature according to which some specific distribution-sensitive principle should come into play (in terms of determining what one should morally do) only in utilitarian ties. The only one that I am able to point out here comes from Henry Sidgwick who wrote that

It is evident that there may be many different ways of distributing the same quantum of happiness among the same number of persons; in order, therefore, that the Utilitarian criterion of right conduct may be as complete as possible, we ought to know which of these ways is to be preferred. This question is often ignored in expositions of Utilitarianism. It has perhaps seemed somewhat idle, as suggesting a purely abstract and theoretical perplexity, that could have no practical exemplification; and no doubt, if all the consequences of actions were capable of being estimated and summed up with mathematical precision, we should probably never find the excess of pleasure over pain exactly equal in the case of two competing alternatives of conduct. But the very indefiniteness of all hedonistic calculations...renders it by no means unlikely that there may be no *cognisable* difference between the quantities of happiness involved in two sets of consequences respectively; the more rough our estimates necessarily are, the less likely we shall be to come to any clear decision between two apparently balanced alternatives. In all such cases, therefore, it becomes practically important to ask whether any mode of distributing a given quantum of happiness is better than any other. Now the Utilitarian formula seems to supply no answer to this question: at least we have to supplement the principle of seeking the greatest happiness on the whole by some principle of Just or Right distribution of this happiness.⁷

A partial explanation for why it is so hard to find such arguments in the literature according to which some distribution-sensitive principle should come into play (in terms of determining what one should morally do) only in utilitarian ties might be that if one is very attracted to some particular type of distributive ideal, it would be a big concession for her to hold that it has practical relevance (in terms of determining what one should morally do) only in utilitarian ties. It may also be the case that many moral philosophers think that it would be somehow incoherent and/or *ad hoc* (or incoherent because it is *ad hoc*) to hold that some particular distributive ideal has practical relevance only in utilitarian ties. But why would it be incoherent

⁷ Sidgwick (1874/1907, 416-417).

or *ad hoc*? Let's have a look at what may be called *goal-based team ball games* (i.e. those team ball games in which goals are scored). In series (e.g. leagues) and group stages in many such sports not only points gathered by teams from different matches but also the goal difference of each team matter in terms of determining the standings of the teams. However, in such series and group stages points lexically dominate goal difference in terms of determining the standings of the teams. There is nothing incoherent or *ad hoc* about such a system for ranking teams, as there are good grounds for believing that a tenable system for ranking teams in series and group stages in goal-based team ball games must be, on the one hand, goal difference-sensitive⁸ and, on the other hand, such in which points lexically dominate goal difference in terms of determining the standings of the teams. (Such a ranking system in a series or a group stage in which goal difference has a bigger role in determining the standings of the teams than the role of a tie-breaker would have various problematic implications which I cannot discuss here.)⁹ Thus why could some distribution-sensitive principle not work as a tie-breaker in such moral choice situations in which there is a utilitarian tie (and which are also inter-personal trade-off choice situations at the level of prospects) in the same fashion as goal difference works as a tie-breaker in series and group stages in such situations in which two or more teams have equally many points? If some distribution-sensitive principle has a lot of plausibility but also has highly problematic implications in many such moral choice situations in which there is not a utilitarian tie, there may be at least some grounds for believing that instead of rejecting such a distribution-sensitive principle, it should be used as a tie-breaker in utilitarian ties. Thus per-

⁸ Dixon (1992, 89) writes that

Throughout the world of soccer, including the World Cup, "goal difference" is used as a tie-breaker. Of all the methods of tie-breaking, this is the least controversial, since people recognize that the ability to score and prevent goals over a series of games, especially over a whole season, is a reliable measure of excellence in soccer.

Since the ability to score and prevent goals over a series of matches is a reliable measure of athletic excellence in goal-based team ball games, it seems to me clear that in goal-based team ball games goal difference should matter in terms of determining the standings of the teams.

⁹ My intention is to discuss them elsewhere.

haps some distribution-sensitive principle is a plausible criterion of moral betterness in utilitarian ties but only in utilitarian ties.

An obvious question then is what distribution-sensitive principle? As I have indicated, I believe that giving intrinsic moral significance to the distribution of utility chances between different individuals, call it *the fair chance principle*, is a very good candidate for such a principle. Goal difference has a lot of plausibility as one criterion for ranking teams in goal-based team ball games, but in series and group stages it can justifiably be used for determining the standings of the teams probably only in such situations in which two or more teams have equally many points. If the fair chance principle has a lot of plausibility as a moral principle but has highly problematic implications in many such moral choice situations in which there is not a utilitarian tie, it may be an indication that, on the one hand, the fair chance principle should not be discarded completely and that, on the other, it should be invoked only in such moral choice situations in which there is a utilitarian tie.

6 THE ARBITRARINESS OF PRIORITARIAN WEIGHTINGS

Let's imagine that at least some of the arguments that I have given against (desert-sensitive versions of) ex post prioritarianism, utilitarian-adjusted ex post prioritarianism, ex ante prioritarianism, ex post-adjusted ex ante prioritarianism and utilitarian-adjusted ex ante prioritarianism are implausible arguments. Let's also imagine that it is not true that prioritarianism violates impartiality or that it is not a problem regarding the plausibility of prioritarianism that it violates impartiality. In other words, let's imagine that at least one form of prioritarianism survives from my arguments against it and that it is not true that prioritarianism violates impartiality or that it is not a problem regarding the plausibility of prioritarianism that it violates impartiality. Even in that case there may be good grounds for believing that prioritarianism is not the way to go if it is true (as I believe it is) that prioritarians face a problem regarding determining the amount of prioritarian weighting on each utility or expected utility level of an individual. Related to that problem, Hirose writes that

[My] criticism [regarding Prioritarianism] is concerned with whether or not there exists a strict concave function that is independent of distributions of people's well-being. The strictly concave function of Prioritarianism determines how much each person's well-being counts in the overall goodness of a distribution. How do we decide how much an individual's well-being counts? How do we discover the shape of such a function? Prioritarians would claim that there is a moral scale of how much a person's well-being counts...I, however, find it hard to believe that such a moral scale [that stipulates how much each person's well-being counts] exists independently of distributions of people's well-being. Prioritarianism's claim about the law of diminishing moral goodness is *intuitively* appealing. But it requires a moral scale that determines how much a person's well-being counts. Arguably, some scientist or sociologist may find that the goodness of a person's well-being is, for example, a square-root function of his well-being, or another strictly concave function. But it is dubious if there really exists such an absolute moral scale.¹

¹ Hirose (2011, 95-96).

To put it more epistemically and less ontologically, I believe that prioritaricians face what may be called

The Arbitrariness Problem of Prioritarian Weightings: Even if there are good grounds for believing that the most plausible form of prioritarianism (e.g. ex ante prioritarianism or some kind of desert-sensitive ex ante prioritarianism) is a more plausible moral view than the most plausible form of utilitarianism, how could it be determined *non-arbitrarily* how much prioritarian weighting (where this means more than no prioritarian weighting at all) is the correct amount of prioritarian weighting on each utility or expected utility level of an individual?²

We have already seen that although all prioritarian moral significance lines slope upwards and bend downwards, prioritarian moral significance lines vary considerably (see Figure 1 on page 9). In order for the amount of prioritarian weighting on each utility or expected utility level of an individual to be determined non-arbitrarily, a prioritarian needs to give a *good* reason for believing in some particular amount of prioritarian weighting (where this means more than no prioritarian weighting at all) on each utility or expected utility level of an individual. But what might be a good reason? There are two very different kinds of positions on this question that one could take.

The Intuitionistic Position about Prioritarian Weightings: The correct amount of prioritarian weighting (where this means more than no prioritarian weighting at all) on each utility or expected utility level of an individual may be determined non-arbitrarily by relying on moral intuitions about the correct amount of prioritarian weighting and some argument for the moral authority of moral intuitions in determining the correct amount of prioritarian weighting on each utility or expected utility level of an individual. An argument is not needed about the correct amount of prioritarian weighting (where this means more than no prioritarian weighting at all) as such. Instead an argument is needed for the moral authority of moral intuitions in determining the correct amount of prioritarian weighting on

² For the sake of clarity, I will repeat the expression “where this means more than no prioritarian weighting at all” many times in this chapter.

each utility or expected utility level of an individual. This kind of *indirect* (and intuitionistic) way of arguing about the correct amount of prioritarian weighting on each utility or expected utility level of an individual is the correct way of arguing by a prioritarian.

The Non-intuitionistic Position about Prioritarian Weightings: The correct amount of prioritarian weighting (where this means more than no prioritarian weighting at all) on each utility or expected utility level of an individual cannot be determined non-arbitrarily by relying on moral intuitions about the correct amount of prioritarian weighting and some argument for the moral authority of moral intuitions in determining the correct amount of prioritarian weighting on each utility or expected utility level of an individual. The correct amount of prioritarian weighting (where this means more than no prioritarian weighting at all) on each utility or expected utility level of an individual might be determined non-arbitrarily only by relying on some *direct* (and non-intuitionistic) argument about the correct amount of prioritarian weighting. In other words, the correct amount of prioritarian weighting (where this means more than no prioritarian weighting at all) on each utility or expected utility level of an individual might be determined non-arbitrarily only by relying on some argument about the correct amount of prioritarian weighting and thus could not be determined non-arbitrarily by relying on *an argument about something else* (i.e. an argument for the moral authority of moral intuitions in determining the correct amount of prioritarian weighting).

Parfit seems to assume that the Intuitionistic Position about Prioritarian Weightings is a plausible position, as he writes that

[prioritarianism] is, in Rawls's sense, intuitionist. It does not tell us how much priority we should give to those who are worse off. On this view, benefits to the worse-off could be morally outweighed by sufficient benefits to the better off. To decide what would be sufficient, we must simply use our judgement.³

³ Parfit (1995, 101).

If the correct amount of prioritarian weighting (where this means more than no prioritarian weighting at all) on each utility or expected utility level of an individual can be determined non-arbitrarily by relying on moral intuitions and some argument for the moral authority of moral intuitions in determining the correct amount of prioritarian weighting on each utility or expected utility level of an individual, the Arbitrariness Problem of Prioritarian Weightings is not much of a problem for prioritarians. But if the Non-intuitionistic Position about Prioritarian Weightings is a more plausible position than the Intuitionistic Position about Prioritarian Weightings, the Arbitrariness Problem of Prioritarian Weightings may be a big problem for prioritarians, as arguably *at least so far* no one has given a good non-intuitionistic argument for any specific amount of prioritarian weighting (where this means more than no prioritarian weighting at all) on each utility or expected utility level of an individual, or to put it differently, for any particular kind of prioritarian moral significance line. In fact, I am not aware of any non-intuitionistic argument about how much prioritarian weighting (where this means more than no prioritarian weighting at all) – even roughly speaking – on each utility or expected utility level of an individual is the correct amount of prioritarian weighting. So from the point of view of the Non-intuitionistic Position about Prioritarian Weightings, there are probably no grounds for believing that the Arbitrariness Problem of Prioritarian Weightings has been solved. Related to this, the following excerpt by a well-known advocate of responsibility-catering prioritarianism, Richard J. Arneson, is telling:

“Prioritarianism” names a type of position, not a specific principle. To get a specific principle one needs a function that determines, for each increment on an absolute benefit scale, the moral value of securing a small benefit for a person at that benefit level. At one end of the prioritarian scale one gets virtually no weighting and an identification of maximal moral value with maximization of benefits; at the other end of the scale one gets leximin. When I contrast prioritarianism with other views I have in mind not the generic position but a restricted family of priority weightings in the middle of the range, but I have nothing useful to say about how to identify a specific principle.⁴

⁴ Arneson (1999, 238).

Arneson seems to think that the most plausible version of prioritarianism is such prioritarianism which is “halfway” between utilitarianism and leximin, but he does not tell what that actually means. He does not tell what that kind of prioritarianism in fact looks like. In other words, Arneson does not have anything concrete to say about how much prioritarian weighting (where this means more than no prioritarian weighting at all) on each utility or expected utility level of an individual is the correct amount of prioritarian weighting. It should also be noticed that Arneson does not provide an *argument* for the view that the most plausible version of prioritarianism is halfway between utilitarianism and leximin. So even if that kind of prioritarianism which is halfway between utilitarianism and leximin is identified, it would be unclear why it should be thought that the most plausible version of prioritarianism is such which is halfway between utilitarianism and leximin.

The question whether it is the Intuitionistic Position about Prioritarian Weightings or the Non-intuitionistic Position about Prioritarian Weightings which is the more plausible one is a complex metaethical question which is tightly connected to a broader metaethical question about the role of intuitions in moral reasoning. I cannot discuss these questions here in detail, but I explain briefly why I find the Non-intuitionistic Position about Prioritarian Weightings a more plausible position.

If one thinks that the Intuitionistic Position about Prioritarian Weightings is the correct position, then an obvious question which arises is *whose* intuitions might determine the correct amount of prioritarian weighting (where this means more than no prioritarian weighting at all) on each utility or expected utility level of an individual? Related to this, my worry is that the Intuitionistic Position about Prioritarian Weightings makes prioritarianism collapse into some sort of relativism. If the correct prioritarian approach to the question how much prioritarian weighting on each utility or expected utility level of an individual is the correct amount of prioritarian weighting is to rely on people’s intuitions rather than on some argument for some particular prioritarian priority function, it seems to imply that the correct prioritarian priority function might be determined by what kinds of prioritarian intuitions (some) people *in fact happen to have*. But this kind of relativistic

view seems implausible to me.⁵ Related to these considerations, Peter Singer writes that

In *A Theory of Justice*, John Rawls...argued that the test of a sound moral theory is that it can achieve a “reflective equilibrium” with our considered moral judgments. By “reflective equilibrium” Rawls meant that, where there is no inherently plausible theory that perfectly matches our initial moral judgments, we should modify either the theory, or the judgments, until we have an equilibrium between the two. The model here is the testing of a scientific theory. In science, we generally accept the theory that best fits the data, but sometimes, if the theory is inherently plausible, we may be prepared to accept it even if it does not fit all the data. We might assume that the outlying data are erroneous, or that there are still undiscovered factors at work in that particular situation. In the case of a normative theory of ethics, Rawls assumes, the raw data is our prior moral judgments. We try to match them with a plausible theory, but if we cannot, we reject some of the judgments, and modify the theory so that it matches others. Eventually the plausibility of the theory and of the surviving judgments reach an equilibrium, and we then have the best possible theory. On this view the acceptability of a moral theory is not determined by the internal coherence and plausibility of the theory itself, but, to a significant extent, by its agreement with those of our prior moral judgments that we are unwilling to revise or abandon.⁶

Singer makes the following conclusion:

The model of reflective equilibrium has always struck me as dubious. The analogy between the role of a normative moral theory and a scientific theory is fundamentally misconceived. A scientific theory seeks to explain the existence of data that are about a world “out there” that we are trying to explain. Granted, the data may have been affected by errors in measurement or interpretation, but unless we can give some account of what the errors might have been, it is not up to us to choose or reject the observations. A normative ethical theory, however, is not trying to explain our common moral intuitions. It might reject all of them, and still be superior to other normative theories that better matched our moral judgments. For a normative moral theory is not an attempt to answer the question “Why do we think as we do about moral questions?”...it is obvious that the

⁵ Cf. this problem with Adler (2012, 399-404) who discusses some metaethical questions related to determining the amount of prioritarian weighting but who does not consider this problem.

⁶ Singer (2005, 344-345).

question “Why do we think as we do about moral questions?” may require a historical, rather than a philosophical, investigation. On abortion, suicide, and voluntary euthanasia, for instance, we may think as we do because we have grown up in a society that was, for nearly 2000 years, dominated by the Christian religion. We may no longer believe in Christianity as a moral authority, but we may find it difficult to rid ourselves of moral intuitions shaped by our parents and our teachers, who were either themselves believers, or were shaped by others who were.⁷

Thus prioritarianism should not be considered as an attempt to explain people’s intuitions about giving priority to worse-off individuals over better-off individuals. It is obvious that answering the question “Why do people tend to have certain kinds of intuitions about how much priority should be given to people at some specific utility or expected utility level over some other people at some other specific, higher, utility or expected utility level?” requires historical and sociological, rather than philosophical, investigation. It is hard for me to see how such historical and sociological investigation could be helpful for determining non-arbitrarily how much prioritarian weighting (where this means more than no prioritarian weighting at all) is the correct amount of prioritarian weighting on each utility or expected utility level of an individual. I strongly believe that in order to determine it non-arbitrarily (if it can be determined non-arbitrarily), we need philosophical investigation that looks for a plausible non-intuitionistic argument about the correct amount of prioritarian weighting (where this means more than no prioritarian weighting at all) on each utility or expected utility level of an individual.

If there are good grounds for believing that the most plausible form of prioritarianism is a more plausible moral view than the most plausible form of utilitarianism (which I doubt) but we are nevertheless unable (at least for now) to go beyond utilitarianism in a non-arbitrary prioritarian way due to the implausibility of the Intuitionistic Position about Prioritarian Weightings and the lack of an argument in favour of any particular prioritarian priority function, what does it tell us about prioritarianism? Should we conclude that there is *only* a practical epistemic problem regarding prioritarianism in the sense that we are unable (at least for now) to have justified beliefs about how much prioritarian weighting (where this means more than no prioritarian weighting at all) on each utility or expected utility level

⁷ Singer (2005, 345).

of an individual is the correct amount of weighting? Or should we go further and conclude that if we are unable to have justified beliefs about how much prioritarian weighting (where this means more than no prioritarian weighting at all) on each utility or expected utility level of an individual is the correct amount of weighting, it implies that there are *better* grounds for believing that the correct amount of prioritarian weighting is no prioritarian weighting at all than there are for believing that the most plausible form of prioritarianism is a more plausible moral view than the most plausible form of utilitarianism, *even if* there are *good* grounds for believing that the most plausible form of prioritarianism is a more plausible moral view than the most plausible form of utilitarianism?⁸ One might argue that there is not merely a practical epistemic problem and answer the latter question affirmatively on the basis of the view that (1) the arbitrariness of some moral view, or the

⁸ Bykvist (2010, 72) writes that

Another, more pressing, problem for prioritarianism [than the problem that, on the face of it, prioritarianism does not seem to square with impartiality] is that it is not clear how the weights should be determined. Exactly how much weight should we give to a person at a certain absolute well-being level?

Nevertheless, Bykvist (2010, 159-160, 163) is inclined to endorse some form of prioritarianism. Bykvist (2010, 159-160) writes that

What is radical about utilitarianism is that it cares *equally* about everyone's well-being. This, in part, is what explains why utilitarianism sometimes gives counter-intuitive moral prescriptions. Among the most counter-intuitive ones are, (a) any sacrifice no matter how great can be justified if we make sufficiently many other people slightly better off, (b) inequality of well-being does not matter...I argued that problems (a) and (b) can be dealt with at least to some extent, if we adopt a prioritarian aggregation method, according to which the lower well-being levels of worse off people are assigned more weight...Since it is not clear exactly how much weight should be given to worse off people, sum-ranking still stands out as the simpler and more straightforward method of aggregation. But in the light of its more counter-intuitive implications, I am inclined to reject it in favour of some form of prioritarianism.

It is worth pointing out, however, that Bykvist (2010) does not discuss prioritarianism at all in the context of risk. Thus Bykvist's assessment of the relative plausibility of utilitarianism and prioritarianism does not take into account any of the problems that different forms of prioritarianism can face in such moral choice situations which involve risk.

arbitrariness of some aspect of some moral view, is something that counts, *at least a little bit*, against that moral view, and on the basis of the view that (2) there is nothing arbitrary about utilitarianism if we approach arbitrariness in terms of steps of moral reasoning.

Let's imagine that a prioritarian is unable to provide a good reason for believing in any particular prioritarian priority function but that there are *slightly* better grounds, *if the arbitrariness considerations are excluded*, for believing that the most plausible form of prioritarianism is a more plausible moral view than the most plausible form of utilitarianism than there are for believing that the most plausible form of utilitarianism is a more plausible moral view than the most plausible form of prioritarianism. In that case taking the arbitrariness considerations into consideration might tip the balance in favour of utilitarianism, provided that both (1) and (2) are plausible views. However, I have no answer to how much the arbitrariness of some moral view, or the arbitrariness of some aspect of some moral view, should be considered to count against it.

It seems to me quite clear that (1) is a plausible view. Surely there is something negative about arbitrariness. Imagine that Lisa thinks that the case for the most plausible form of utilitarianism is *equally* strong as the case for the most plausible form of prioritarianism. Lisa thinks so because she thinks that the distribution-insensitivity of utilitarianism counts against utilitarianism, because she thinks that prioritarianism violates impartiality and that that counts equally heavily against prioritarianism, and because she thinks that the most plausible form of utilitarianism and the most plausible form of prioritarianism are not problematic in any other ways. However, Lisa has not thought about any arbitrariness considerations at all. But then Richard tells her that it is unlikely that it could be determined non-arbitrarily how much prioritarian weighting (where this means more than no prioritarian weighting at all) is the correct amount of prioritarian weighting on each utility or expected utility level of an individual. Would it then not be reasonable for Lisa to make the conclusion that utilitarianism is a more plausible moral view than prioritarianism, provided that Richard is right and (2) is a plausible view?

It seems to me that (2) is a much more controversial view than (1). I believe, however, that by invoking Singer I can give a good argument for (2). Singer writes that

the notion of ethics carries with it the idea of something bigger than the individual. If I am to defend my conduct on ethical grounds, I cannot point only to the benefits it brings me. I must address myself to a larger audience...Ethics takes a universal point of view...in making ethical judgments we go beyond our own likes and dislikes...Ethics requires us to go beyond 'I' and 'you' to the universal law, the universalisable judgment, the standpoint of the impartial spectator or ideal observer, or whatever we choose to call it...The universal aspect of ethics, I suggest, does provide a persuasive, although not conclusive, reason for taking a broadly utilitarian position...In accepting that ethical judgments must be made from a universal point of view, I am accepting that my own interests cannot, simply because they are my interests, count more than the interests of anyone else. Thus my very natural concern that my own interests be looked after must, when I think ethically, be extended to the interests of others...The utilitarian position is a minimal one, a first base that we reach by universalising self-interested decision making. We cannot, if we are to think ethically, refuse to take this step. If we are to be persuaded that we should go beyond utilitarianism and accept non-utilitarian moral rules or ideals, we need to be provided with good reasons for taking this further step.⁹

Accepting that my own interests cannot, simply because they are my interests, count more than the interests of anyone else, and extending to the interests of others my very natural concern that my own interests be looked after, is what may be called the universalisation of self-interested decision-making. There is nothing arbitrary about the universalisation of self-interested decision-making, which can and I believe should be seen as the first step of theoretical moral reasoning. Of course, prioritarianism is not incompatible with the universalisation of self-interested decision-making. However, the *mere* universalisation of self-interested decision-making leads to a utilitarian rather than prioritarian position. After the

⁹ Singer (1979/1993, 10-14). Singer provides here what may be characterized as a positive argument (rather than a negative argument) for utilitarianism. Apparently Singer understands his argument to be also an argument for impartiality in moral reasoning, as he writes "the standpoint of the impartial spectator". However, since I believe that there is a difference between impartiality and impersonality and that impersonality does not imply impartiality, I see Singer's argument as an argument (and a good argument) for impersonality rather than impartiality. Thus I think Singer should have written "the standpoint of the impersonal spectator" rather than "the standpoint of the impartial spectator". I strongly believe that plausible moral reasoning implies impersonality, but it is less clear to me whether plausible moral reasoning implies impartiality, although, as I already made it clear in Chapter 4, I believe it does.

universalisation of self-interested decision-making, a prioritarian would have to take a further step which is adding the weightings. But any particular step beyond utilitarianism to prioritarianism is arbitrary if a prioritarian is unable to provide a good reason for believing in some particular amount of prioritarian weighting on each utility or expected utility level of an individual. If a prioritarian is unable to provide such a reason, it perhaps gives us a good reason – besides the other reasons I have given in this book – not to go beyond utilitarianism to prioritarianism.¹⁰ Nevertheless, as I have already made it clear, I do believe that we should go beyond pure utilitarianism and accept non-utilitarian ideals. But those non-utilitarian ideals are non-prioritarian.

¹⁰ I think going beyond utilitarianism to a deontological moral doctrine faces similar problems as going beyond utilitarianism to prioritarianism. Walter Sinnott-Armstrong (2003/2006) writes that

Moderate deontologists...often judge that it is morally wrong to kill one person to save five but not morally wrong to kill one person to save a million. They never specify *the line* between what is morally wrong and what is not morally wrong, and it is hard to imagine any *non-arbitrary* way for deontologists to justify *a cutoff point*. In contrast, consequentialists can simply say that the line belongs wherever the benefits outweigh the costs (including any bad side effects). [Italics added by me.]

By “consequentialists” Sinnott-Armstrong clearly refers to utilitarians. It should be noticed, however, that due to the epistemic problems of estimating utility and probabilities, utilitarian prescriptions are often based on such utility and probability judgments which may be considered quite arbitrary. But it should also be noticed that prioritarian prescriptions are also often based on such utility and probability judgments which may be considered quite arbitrary. So even if it is true that utilitarianism involves arbitrariness, it seems to me that prioritarianism involves more arbitrariness. We may call this “double arbitrariness”. In the case of utilitarianism, it is at most the *application* of utilitarianism which involves arbitrariness, at least if we approach arbitrariness in terms of steps of moral reasoning. This implies that utilitarianism *itself as a moral view* is not arbitrary from a theoretical point of view unlike any particular prioritarian priority function may be. I owe the term “double arbitrariness” to Frej Klem Thomsen.

7 CONCLUSION

In this book I have rejected prioritarianism on the basis of various targeted objections to various specific forms of prioritarianism. All these targeted objections are in one way or another connected to risk or possible future individuals. Besides rejecting prioritarianism on the basis of those targeted objections, I have argued that prioritarianism violates impartiality and discussed moral reasoning and the arbitrariness of prioritarian weightings. Since I believe that impartiality is a property which a moral doctrine must have in order to be a plausible moral doctrine, I believe that prioritarianism should be rejected. Moreover, I believe that there are considerations related to moral reasoning and the arbitrariness of prioritarian weightings that count against prioritarianism. However, I have not presented any (good) *argument* for the view that impartiality is a property which a moral doctrine must have in order to be a plausible moral doctrine.¹ I have also not had anything to say about *how much* the considerations related to moral reasoning and the arbitrariness of prioritarian weightings count against prioritarianism. These are clearly shortcomings and I believe that future research on prioritarianism should address them.

Besides rejecting prioritarianism, I have rejected various other distribution-sensitive moral views. However, I have made it clear that I am in favour of distribution-sensitivity in morality, as I have argued for desert-sensitive fair chance expected utility utilitarianism. What I have left open is in what way(s) (e.g. in some meritocraticness-orientated and/or responsibility-orientated way) we should give intrinsic moral significance to desert. I have relied on my intuition that desert should be given intrinsic moral significance *at least in some way*. In order to develop a complete moral theory, I would have to explain in what way(s) desert should be given intrinsic moral significance, how big (or small) of a role desert should play in moral evaluations,² why desert should be given intrinsic moral significance in the

¹ As I already indicated in Footnote 9 on page 121, I consider Singer's argument a good argument for impersonality rather than a good argument for impartiality.

² One possibility is that desert should come into play as a moral consideration (in terms of determining what one should morally do) only in utilitarian ties. In that case desert should have a small role in moral evaluations compared to the maximization of total expected utility but should have as big a role in moral evaluations as what I think giving fair chances to different individuals should have.

way(s) it should be given, and why desert should play as big (or small) of a role in moral evaluations as it should.

APPENDIX A: RABINOWICZ ON THE SCOPE OF THE DIMINISHING MARGINAL MORAL SIGNIFICANCE OF AN INDIVIDUAL'S UTILITY

Also Wlodek Rabinowicz considers a limitation of the scope of the diminishing marginal moral significance of an individual's utility.¹ Rabinowicz writes that "if prioritarianism is driven by a concern for the distinctness of persons, then the priority weights should only be used in the interpersonal but not in the intrapersonal balancing of benefits and losses."² Rabinowicz also writes that

If...the only function of moral weights were to hinder unacceptable interpersonal compensations, then it would be natural to argue for the coincidence of prudence and morality in one-person cases. This would require that such cases should be treated differently from those that involve several persons. In the determination of overall goodness, moral weights would need to be ignored as long as only one person is involved, and they would only be brought into play in many-person cases.³

This passage puzzles me a lot. Why does Rabinowicz switch from talking about intra-personal and inter-personal balancing and compensations to talking about one-person and multi-person cases? As we have already seen, there are multi-individual cases which involve intra-personal rather than inter-personal trade-offs. Moreover, Rabinowicz continues that

this special treatment of one-person cases would result in implausible consequences. On such a 'mixed' view, as long as Robinson is alone on his island, prioritarian morality would give him purely prudential recommendations. Or perhaps it would give him no recommendations at all, if we suppose that morality is silent as long as there are no other persons that the agent needs to consider...But things change as soon as Man Friday enters the picture. Then moral weights are brought into play. Consequently, with Friday present, Robinson is no longer morally allowed to go for his better prospect (the risky

¹ Rabinowicz (2001; 2002).

² Rabinowicz (2002, 10-11).

³ Rabinowicz (2002, 17).

one), even if...his choice would not affect Friday in any way. If Friday's welfare given each state is the same whatever prospect Robinson chooses and if Robinson's larger gain in one state weighs less, morally, than his smaller loss in the other (equiprobable) state, then, with Friday present, the risky prospect is overall worse than the riskless one. Such an extreme sensitivity to 'other persons being present' is counter-intuitive. Surely, if Robinson's choice cannot affect Friday, bringing the latter into the picture should not morally matter according to prioritarianism. If the riskless prospect is morally preferable to the risky one with Friday present, and the outcome for Friday is not affected by the choice between these prospects, then it is reasonable to require that the riskless prospect be morally preferable also when Friday is absent. We get this desirable implication if we treat one-person cases in the same way as situations that involve several persons.⁴

This passage is also very puzzling. Clearly Man Friday's presence would not make the case a multi-individual case if his utility is not at stake in it. Whether some case is a one-individual case or a multi-individual case is determined by how many individuals' utility is at stake in it, or more accurately, I believe, by how many individuals' utility is at stake in it non-trivially.

⁴ Rabinowicz (2002, 17).

APPENDIX B: THE DOMINANCE PRINCIPLE AND CONSEQUENTIALISM

It may be wondered whether such a moral doctrine which is incompatible with the dominance principle can be a form of *consequentialism*. If consequentialism is understood so that a moral doctrine is a form of consequentialism if and only if it evaluates the moral goodness of acts solely in terms of *the goodness of* outcomes (or consequences) or solely in terms of the goodness and probabilities of outcomes, then such a moral doctrine which is incompatible with the dominance principle clearly cannot be a form of consequentialism and must be a form of *non-consequentialism*. On the other hand, if consequentialism is understood so that a moral doctrine is a form of consequentialism if and only if it evaluates the moral goodness of acts solely in terms of outcomes or solely in terms of outcomes and probabilities of outcomes, then such a moral doctrine which is incompatible with the dominance principle is not necessarily a form of non-consequentialism. If consequentialism is understood in this latter and wider way, *ex ante* prioritarianism qualifies as a form of consequentialism, as *ex ante* prioritarianism does not claim that the moral goodness or badness of an act can be affected by something else than some kinds of outcomes or consequences and their probabilities. However, a deontological moral doctrine, which claims that certain kinds of acts (e.g. lying, cheating, stealing and/or murdering another human being) are *intrinsically* morally bad acts, is a form of non-consequentialism also according to this latter and wider understanding of consequentialism, because an intrinsic moral badness of an act is not any kind of outcome, consequence or probability of an act.

BIBLIOGRAPHY

Adler, Matthew D. (2012): *Well-Being and Fair Distribution: Beyond Cost-Benefit Analysis*. Oxford University Press, New York.

Arneson, Richard J. (1999): "Egalitarianism and Responsibility", *The Journal of Ethics*, Vol. 3, 225-247.

Arneson, Richard J. (2000): "Luck Egalitarianism and Prioritarianism", *Ethics*, Vol. 110, 339-349.

Arneson, Richard J. (2001): "Against Rights", *Philosophical Issues*, Vol. 11, *Social, Political, and Legal Philosophy*, supp. to *Noûs*, 172-201.

Benbaji, Yitzhak (2005): "The Doctrine of Sufficiency: A Defence", *Utilitas*, Vol. 17, 310-332.

Bentham (1789/1823): *The Principles of Morals and Legislation*. Second edition. Reprinted in 1948 by Hafner Press, New York.

Berman, Mitchell N. (2011): "On Interpretivism and Formalism in Sports Officiating: From General to Particular Jurisprudence", *Journal of the Philosophy of Sport*, Vol. 38, 177-196.

Bognar, Greg (2012): "Empirical and Armchair Ethics", *Utilitas*, Vol. 24, 467-482.

Boran, Idil (2004): "On Distinguishing Between Types of Impartiality", *The Journal of Value Inquiry*, Vol. 38, 333-339.

Broome, John (1991): *Weighing Goods*. Basil Blackwell, Oxford.

Broome, John (2004): *Weighing Lives*. Oxford University Press, Oxford.

Brown, Campbell (2005): "Priority or Sufficiency...or Both?", *Economics and Philosophy*, Vol. 21, 199-220.

Brown, Campbell (2007): "Prioritarianism for Variable Populations", *Philosophical Studies*, Vol. 134, 325-361.

Bykvist, Krister (2010): *Utilitarianism: A Guide for the Perplexed*. Continuum International Publishing Group, London.

Casal, Paula (2007): "Why Sufficiency Is Not Enough", *Ethics*, Vol. 117, 296-326.

Crisp, Roger (2003): "Equality, Priority and Compassion", *Ethics*, Vol. 113, 745-763.

Crisp, Roger (2006): *Reasons and the Good*. Oxford University Press, Oxford.

Crisp, Roger (2011): "In Defence of the Priority View: A Response to Otsuka and Voorhoeve", *Utilitas*, Vol. 23, 105-108.

Diamond, Peter A. (1967): "Cardinal Welfare, Individualistic Ethics, and Interpersonal Comparison of Utility: Comment", *Journal of Political Economy*, Vol. 75, 765-766.

Dixon, Nicholas (1992): "On Sportsmanship and 'Running Up the Score'". Reprinted in Jan Boxill (ed.) *Sports Ethics: An Anthology* (2003). Blackwell Publishing, Malden, 84-94.

Dixon, Nicholas (1999): "On Winning and Athletic Superiority". Reprinted in Jan Boxill (ed.) *Sports Ethics: An Anthology* (2003). Blackwell Publishing, Malden, 116-129.

Driver, Julia (2012): *Consequentialism*. Routledge, Abingdon.

Feldman, Fred (1995): "Adjusting Utility for Justice: A Consequentialist Reply to the Objection from Justice", *Philosophy and Phenomenological Research*, Vol. 55, 567-585.

Frankena, William K. (1963/1973): *Ethics*. Second edition. Prentice-Hall, Englewood Cliffs.

Gibbard, Allan (1984): "Health Care and the Prospective Pareto Principle", *Ethics*, Vol. 94, 261-282.

Gruen, Lori (1999): "Must Utilitarians Be Impartial?". In Dale Jamieson (ed.) *Singer and His Critics*. Blackwell, Oxford, 129-149.

Hare, R.M. (1981): *Moral Thinking: Its Levels, Method, and Point*. Oxford University Press, Oxford.

Hirose, Iwao (2004): "Aggregation and Numbers", *Utilitas*, Vol. 16, 62-79.

Hirose, Iwao (2005): "Intertemporal Distributive Judgement", *Ethical Theory and Moral Practice*, Vol. 8, 371-386.

Hirose, Iwao (2007): "Aggregation and Non-Utilitarian Moral Theories", *Journal of Moral Philosophy*, Vol. 4, 273-284.

Hirose, Iwao (2011): "Egalitarianism and Prioritarianism". In Thom Brooks (ed.) *New Waves in Ethics*. Palgrave Macmillan, Basingstoke, 88-107.

Holtug, Nils (1999): "Utility, Priority and Possible People", *Utilitas*, Vol. 11, 16-36.

Holtug, Nils (2001): "On the Value of Coming into Existence", *The Journal of Ethics*, Vol. 5, 361-384.

Holtug, Nils (2007a): "A Note on Conditional Egalitarianism", *Economics and Philosophy*, Vol. 23, 45-63.

Holtug, Nils (2007b): "Equality for Animals". In Jesper Ryberg, Thomas S. Petersen & Clark Wolf (eds.) *New Waves in Applied Ethics*. Palgrave Macmillan, Basingstoke, 1-24.

Holtug, Nils (2007c): "On Giving Priority to Possible Future People". In Toni Rønnow-Rasmussen, Björn Petersson, Jonas Josefsson & Dan Egonsson (eds.) *Hommage à Wlodek. Philosophical Papers Dedicated to Wlodek Rabinowicz*. URL = <<http://www.fil.lu.se/hommageawlodekhttp://plato/>>.

Holtug, Nils (2007d): "Prioritarianism". In Nils Holtug & Kasper Lippert-Rasmussen (eds.) *Egalitarianism: New Essays on the Nature and Value of Equality*. Oxford University Press, Oxford, 125-156.

Holtug, Nils (2010): *Persons, Interests, and Justice*. Oxford University Press, Oxford.

Hooker, Brad (2000): *Ideal Code, Real World: A Rule-Consequentialist Theory of Morality*. Oxford University Press, Oxford.

Hooker, Brad (2003/2008): "Rule Consequentialism". In Edward N. Zalta (ed.) *The Stanford Encyclopedia of Philosophy* (Winter 2008 Edition). URL = <<http://plato.stanford.edu/archives/win2008/entries/consequentialism-rule/>>.

Loland, Sigmund (2002): *Fair Play in Sport*. Routledge, London.

Mason, Andrew (2006): *Levelling the Playing Field: The Idea of Equal Opportunity and its Place in Egalitarian Thought*. Oxford University Press, Oxford.

McCarthy, David (1998): "Actions, Beliefs, and Consequences", *Philosophical Studies*, Vol. 90, 57-77.

McCarthy, David (2006): "Utilitarianism and Prioritarianism I", *Economics and Philosophy*, Vol. 22, 335-363.

McCarthy, David (2008): "Utilitarianism and Prioritarianism II", *Economics and Philosophy*, Vol. 24, 1-33.

McCarthy, David (2012): "Risk-Free Approaches to the Priority View", *Erkenntnis*, Vol. 78, 421-449.

McKerlie, Dennis (1984): "Egalitarianism", *Dialogue*, Vol. 23, 223-237.

McKerlie, Dennis (1989): "Equality and Time", *Ethics*, Vol. 99, 475-491.

Mill, John Stuart (1863/1871): "Utilitarianism". Fourth edition. Reprinted in *The Basic Writings of John Stuart Mill* (2002). Modern Library, New York, 231-301.

Miller, Dale E. (2003): "Actual-consequence Act Utilitarianism and the Best Possible Humans", *Ratio*, Vol. 16, 49-62.

Mulgan, Tim (2007): *Understanding Utilitarianism*. Acumen, Stocksfield.

O'Neill, Martin (2008): "What Should Egalitarians Believe?", *Philosophy & Public Affairs*, Vol. 36, 119-156.

O'Neill, Martin (2012): "Priority, Preference and Value", *Utilitas*, Vol. 24, 332-348.

Otsuka, Michael (2012): "Prioritarianism and the Separateness of Persons", *Utilitas*, Vol. 24, 365-380.

Otsuka, Michael & Alex Voorhoeve (2009): "Why It Matters That Some Are Worse Off Than Others: An Argument against the Priority View", *Philosophy & Public Affairs*, Vol. 37, 171-199.

Otsuka, Michael & Alex Voorhoeve (2011): “Reply to Crisp”, *Utilitas*, Vol. 23, 109-114.

Parfit, Derek (1995): “Equality or Priority?”. Reprinted in Matthew Clayton & Andrew Williams (eds.) *The Ideal of Equality* (2000). Palgrave Macmillan, Basingstoke, 81-125.

Parfit, Derek (1997): “Equality and Priority”, *Ratio*, Vol. 10, 202-221.

Parfit, Derek (2012): “Another Defence of the Priority View”, *Utilitas*, Vol. 24, 399-440.

Peterson, Martin & Sven Ove Hansson (2005): “Equality and Priority”, *Utilitas*, Vol. 17, 299-309.

Porter, Thomas (2012): “In Defence of the Priority View”, *Utilitas*, Vol. 24, 349-364.

Rabinowicz, Wlodek (2001): “Prioritarianism and Uncertainty: On the Interpersonal Addition Theorem and the Priority View”. In Dan Egonsson, Jonas Josefsson, Björn Petersson & Toni Rønnow-Rasmussen (eds.) *Exploring Practical Philosophy: From Action to Values*. Ashgate, Aldershot, 139-165.

Rabinowicz, Wlodek (2002): “Prioritarianism for Prospects”, *Utilitas*, Vol. 14, 2-21.

Rawls, John (1971/1999): *Theory of Justice*. Revised edition. Belknap Press of Harvard University Press, Cambridge.

Scheffler, Samuel (1988): “Introduction”. In Samuel Scheffler (ed.) *Consequentialism and Its Critics*. Oxford University Press, Oxford, 1-13.

Segall, Shlomi (2011): "Luck Prioritarian Justice in Health". In Carl Knight & Zofia Stemplowska (eds.) *Responsibility and Distributive Justice*. Oxford University Press, Oxford, 246-265.

Shoemaker, David W. (1999): "Utilitarianism and Personal Identity", *The Journal of Value Inquiry*, Vol. 33, 183-199.

Sidgwick, Henry (1874/1907): *The Methods of Ethics*. Seventh edition. Reprinted in 1981 by Hackett Publishing Company, Indianapolis.

Singer, Peter (1979/1993): *Practical Ethics*. Second edition. Cambridge University Press, Cambridge.

Singer, Peter (1999): "A Response". In Dale Jamieson (ed.) *Singer and His Critics*. Blackwell, Oxford, 269-335.

Singer, Peter (2005): "Ethics and Intuitions", *The Journal of Ethics*, Vol. 9, 331-352.

Singer, Peter (1979/2011): *Practical Ethics*. Third edition. Cambridge University Press, Cambridge.

Sinnott-Armstrong, Walter (2003/2006): "Consequentialism". In Edward N. Zalta (ed.) *The Stanford Encyclopedia of Philosophy* (Fall 2008 Edition), URL = <http://plato.stanford.edu/archives/fall2008/entries/consequentialism/>.

Smart, J.J.C. (1973): "An Outline of a System of Utilitarian Ethics". In J.J.C. Smart & Bernard Williams *Utilitarianism: For and Against*. Cambridge University Press, Cambridge, 1-74.

Temkin, Larry S. (1993): *Inequality*. Oxford University Press, New York.

Temkin, Larry S. (2001): "Worries about Continuity, Expected Utility Theory, and Practical Reasoning". In Dan Egonsson, Jonas Josefsson, Björn Petersson & Toni

Rønnow-Rasmussen (eds.) *Exploring Practical Philosophy: From Action to Values*. Ashgate, Aldershot, 95-108.

Temkin, Larry S. (2003): "Equality, Priority or What?", *Economics and Philosophy*, Vol. 19, 61-87.

Temkin, Larry (2011): "Justice, Equality, Fairness, Desert, Rights, Free Will, Responsibility, and Luck". In Carl Knight & Zofia Stemplowska (eds.) *Responsibility and Distributive Justice*. Oxford University Press, Oxford, 51-76.

Timmermann, Jens (2004): "The Individualist Lottery: How People Count, but not Their Numbers", *Analysis*, Vol. 64, 106-112.

Tungodden, Bertil (2003): "The Value of Equality", *Economics and Philosophy*, Vol. 19, 1-44.

Voorhoeve, Alex & Marc Fleurbaey (2012): "Egalitarianism and the Separateness of Persons", *Utilitas*, Vol. 24, 381-398.

Weirich, Paul (1983): "Utility Tempered with Equality", *Noûs*, Vol. 17, 423-439.

Williams, Andrew (2012): "The Priority View Bites the Dust?", *Utilitas*, Vol. 24, 315-331.