

The Demandingness Objection to Peter Singer's Account of Our Obligations to the World's Poor

**By
Yuliya Kanygina**

Submitted to
Central European University
Department of Political Science

In partial fulfillment of the requirements for the degree of Master of Arts

Supervisor: Professor Zóltan Miklosi

Budapest, Hungary 2011

Abstract

In this thesis I argue that Peter Singer's account of our moral obligations to the world's poor in particular, and other act-consequentialist approaches to the problem of global poverty in general, are implausible due to their unreasonable demandingness. I examine three groups of act-consequentialists views: extremism of Peter Singer and Peter Unger, sub-maximizing act-consequentialism of Michael Slote and Samuel Scheffler, and Liam Murphy's collective act-consequentialism. I argue that each of these views, in addition to the problems of their own, faces the objection that the moral demands it imposes as morally required are overdemanding. I argue that the demandingness objection is sufficient to render each of them unsound and that our strongly embedded intuitive belief that morality cannot require us more than we can reasonably bear is not plausibly accounted for by any of these views.

Contents

INTRODUCTION	1
CHAPTER 1: THE DEMANDINGNESS OBJECTION	7
1.1 DEMANDS, IMPARTIALITY, AND SUPEREROGATION	7
1.2 THE DEMANDINGNESS OBJECTION AND THE INTEGRITY OBJECTION	12
CHAPTER 2: EXTREMISM	15
2.1 SINGER’S GENERAL ARGUMENT	15
2.2 THE ARGUMENT FROM CONFLICTING INTUITIONS	18
2.3 THE ARGUMENT FROM FMRI	25
CHAPTER 3: SUB-MAXIMIZING VIEWS.....	35
3.1 SLOTE’S SATISFICING CONSEQUENTIALISM	35
3.2 SCHEFFLER’S HYBRID VIEW	39
CHAPTER 4: COLLECTIVE ACT-CONSEQUENTIALISM	47
4.1 MURPHY’S COLLECTIVE PRINCIPLE OF BENEFICENCE	47
CONCLUSION	56
REFERENCE LIST.....	60

Introduction

All around the globe, people are dying from lack of food, shelter, malnutrition, treatable diseases, and other preventable causes. Nearly three billion people live on less than two dollars a day. Almost 25,000 children die every day due to malnutrition. About 3 billion children die each year from dehydrating diarrhea. More than a billion citizens of developing countries have no adequate access to safe water and basic sanitation. A quarter of humanity lives without electricity (UN Report on the World Social Situation 2010).

These figures are striking. Without any doubt, the majority of the world's population suffers enormous hardships. Contrary to this, most members of affluent countries are living comfortable lives and only a negligible minority of them think they have a moral duty to help the poor. These differences trigger important moral questions. What should our moral obligations towards the global poor be? What is the level of sacrifice that morality requires from the citizens of affluent countries? If we are not doing anything or doing very little to prevent famine and starvation in the third-world countries, do we live immoral lives?

Act-consequentialism provides straightforward answers to these questions. According to this moral theory, the morally right act is the one that brings about the best possible outcome in terms of the overall good (usually but not necessarily specified as well-being). Consequentialists claim that morality is exclusively about bringing about the best overall outcome. Their argument is the following. Since each and every person matters equally, the well-being of everyone should have the same moral value. My own well-being cannot matter more just because it is mine. The morally right act then will be the one that promotes the well-being of all affected parties, that is to say, that maximizes the well-being of everyone, taken as a collective (Singer 1993: 13; cf. Kymlicka 2002: 32-33). If we want to live moral lives

INTRODUCTION

according to act-consequentialist standards, we should strive to choose those actions among all the options available for us that will maximize the total net amount of goodness.

Suppose that act-consequentialism is the correct theory of morality. If this is the case, most of us, who are citizens of the affluent countries, are living highly immoral lives. Imagine that you have just received your monthly salary. Act-consequentialism suggests that you should spend it on bringing about the best outcome in terms of overall good. You might think that buying a new pair of shoes for yourself or a new toy for your child is such an outcome. After all, by doing so, you will improve the well-being of two people. However, there are much better ways to spend your money in order to bring the best overall state of affairs. You could benefit more people in need by transferring a part of your salary to Oxfam or any other world-renowned humanitarian organization that will use that money to provide food, water and shelter for people in need.

It was almost forty years ago that Peter Singer developed his, now widely acknowledged as classical, act-consequentialist argument for our obligations to the world's poor (see Singer 1972). Singer's argument is very straightforward. If he is right, then most of us live highly immoral lives. In one of his recent pieces (Singer 2009:15-6), Singer summarizes it in the following manner:

First premise. Suffering and death from lack of food, shelter, and medical care are bad.

Second premise. If it is in your power to prevent something bad from happening, without sacrificing anything nearly as important, it is wrong not to do so.

Third premise. By donating to aid agencies, you can prevent suffering and death from lack of food, shelter, and medical care, without sacrificing anything nearly as important.

Conclusion. Therefore, if you do not donate to aid agencies, you are doing something

wrong.

The first premise is a universally held belief. We all think that a life of absolute deprivation is unbearable. The third premise is an empirical claim. At least nowadays it is plausible to assume that prominent aid agencies such as Oxfam will perform their job properly, that is to say, they will use the money for the given purpose and in the most effective way.

The second premise, however, is problematic and may make Singer's argument unsound. If I sacrifice my new pair of shoes by jumping into the water in order to save a drowning child, then there is no doubt that I have to do it since the value of the child's life is incomparable to the value of my shoes (Singer 1972: 231). In the same way, why would there be any moral difference between saving the drowning child and one that dies from hunger or minor diseases in sub-Saharan Africa? That child could be easily saved by transferring a similar amount of money as needed for a pair of expensive shoes.

However, according to Singer, our obligations towards the poor do not stop here. It is not the case that we must only save one child but that we ought to prevent suffering and death whenever we can, or more precisely, whenever possible if by doing so we do not make our lives as bad as the ones we wanted to save. This would mean giving up many things we enjoy such as expensive clothes, technological devices, hobbies, travel, and much more. In his original paper, Singer defended a radical claim that we ought to work full time to relieve great suffering caused by famine and similar disasters (Singer 1972: 238). More recently, Singer wrote that he still holds this view on the level of ideal theory, though on the level of practical application he believes that we should give at least 5 percent of our annual income (Singer 2009: 152). Both of these views, however, have been criticized as unreasonably demanding (see, for example, Mulgan 2001; Murphy 2000).

Suppose that you become convinced by the act-consequentialist argument and regularly

INTRODUCTION

give some money to Oxfam. Your moral obligations towards the poor, however, still rest unfulfilled. You have done one morally right action but if you want to live up to act-consequentialist standards you should continue to bring about the best outcome. You should continue to spend your time, energy, and resources for promoting the overall well-being. As some act-consequentialists argue, morality requires that we work full time to relieve great suffering caused by famine and similar disasters (Singer 1972: 238).

To many people, these demands seem unreasonably high and counterintuitive. None of us lives up to these standards and we still believe that we are not acting immorally. Probably the most important objection to act-consequentialism is that it is an unreasonably demanding theory (see Mulgan 2001: 4). By the same token, Singer's act-consequentialist account of our obligations to the world's poor is subject to the demandingness objection.

Consequentialists react to this objection by saying that the fact that their standards are too demanding does not diminish their moral rightness. This means that, though giving up all or almost all of our personal projects and investing time, forces, and money in humanitarian missions is demanding, it is still something we should all strive to do. I claim that the demandingness objection has detrimental consequences for Singer's argument and act-consequentialism in general.

My argumentation will proceed in the following manner. In the first chapter of this thesis I will present the demandingness objection and its variations. I will explain three versions of the demandingness objection and show, by examples, how it threatens act-consequentialism in general and Singer's account in particular. I will also show why simple act-consequentialism does not allow for supererogation and how this is connected with the demandingness objection.

In the second chapter I will present and examine one of the act-consequentialist's responses to the demandingness objection. Following Shelly Kagan I shall call this response

INTRODUCTION

extremism and its followers extremists (Kagan 1989: 2). Extremists' strategy is to show that demands of act-consequentialism are not unreasonable. That is to say, the counterintuitiveness of act-consequentialists conclusions does not render them unreasonable. Extremists' main argument is the rejection of the reliability of our ordinary moral intuitions. I will argue against two versions of this argument. First, I will show that Singer's and Joshua Greene's rejection of moral intuitions on the basis of fMRI studies does not provide the normative argumentation necessary to undermine reliability of moral intuitions. Second, I will criticize Peter Unger's argument against the reliability of intuitions by showing that his method of several options and the examples he uses to back it up do not provide the basis to believe that moral intuitions are conflicting per se.

Having rejected Singer's classical act-consequentialist account, in the third chapter I show that there may be a way out for Singer if he accepts a version of act-consequentialism that does not always require constant maximization of the overall good, allowing agents to bring less than optimal outcome in circumstances where their own interests are in question. I will present two versions of this solution: Michael Slote's satisficing consequentialism and Samuel Scheffler's so-called hybrid view. I will show that the former is no less demanding than Singer's view and the latter has counterintuitive implications.

Finally, in the fourth chapter, I will examine an alternative view, defended by Liam Murphy and representing a specific form of collective consequentialism. Murphy's collective consequentialism supposedly offers a solution for unreasonable demandingness of act-consequentialism by suggesting that each of us has a moral obligation towards the poor as if all affluent people were engaged, as members of joint cooperative venture, in maximizing the overall good and has no moral obligation beyond that cooperative share. Murphy claims that one should sacrifice only optimal amount of personal good. This amount is calculated by dividing the good to be done in order to eliminate global poverty by the number of agents able

INTRODUCTION

to provide the good. I will undermine Murphy's view by showing its limitations and by criticizing his method of defining the optimal amount of sacrifice. I will demonstrate, that in some cases, it leads us to unreasonable demandingness, while in the others, it implies unreasonable underdemandingness.

The methodology that I use in this thesis relies heavily on reflective equilibrium (Rawls 1971: 46-53). When assessing the truth of some principle, theory, or claim, I often invoke our ordinary intuitions. If some principle or claim goes against our intuitions, I consider it to be mistaken. This method was criticised by Singer and others due to its reliance on ordinary moral intuitions which in its turn are supposedly unreliable because of their sensitivity to morally irrelevant features. In the second chapter, although with a different aim, I reject these arguments and argue that we should rely on our ordinary moral intuitions in moral philosophy. Another methodological device that I use are thought experiments. Although sometimes unrealistic, thought experiments are a widely used tool in philosophy and, apart from clarifying our intuitive understanding of some issue, their independent justification is unnecessary.

My thesis has two goals. On the surface, I will show that the demandingness objection is detrimental to most act-consequentialist arguments and that the forcefulness of this objection is especially vivid in matters of global poverty. In a broader way, my argument will suggest that ethical question of our obligations towards the world's poor can hardly be answered through act-consequentialist moral theory. Rather than debating whether being a moral saint is what morality really requires from us or not, we should try to find an alternative grounding for an intuitive view that something has to be done with respect to the absolute deprivation around the world.

Chapter 1: The Demandingness Objection

The demandingness objection roughly says that it is impossible or very unlikely that anyone could seriously follow the requirements proposed by act-consequentialism. To illustrate this, imagine the following scenario. Anna is a citizen of an affluent country and she is very active in reducing global poverty. She works for Oxfam and also gives half of her monthly income to this organization. She is also engaged in programs fighting illiteracy in poor countries. However, Anna is a great admirer of ballet. Now imagine that Anna has just received her monthly salary, half of which was already transferred to Oxfam's aid account. The rest of the money she will need for normal survival, though one small part of it, let us say 15 percent, will be an extra value. Anna decides to use this money to buy a ticket for the premiere of *The Nutcracker*. According to act-consequentialism, by buying the ticket, Anna is acting wrongly in a moral sense. She could have donated this money to Oxfam since watching ballet is something morally incomparable to saving someone's life. But how could it be that someone who has devoted her life to charity could be doing something wrong here?

In this chapter, I will examine the demandingness objection. I will also have a closer look at the impartiality requirement of act-consequentialism. I will show how it denies supererogation and generates two important implications which render act-consequentialism counterintuitive.

1.1 Demands, Impartiality, and Supererogation

Demandingness objection was developed as a reaction to the requirement of impartiality of beneficence and the requirement of maximization of the overall good that, conjointly taken,

make act-consequentialism a distinctive moral theory. Demandingness objection rests on the assumption that any moral theory with strikingly counterintuitive implications cannot be plausible. In other words, moral theory must be in coherence with our fundamental moral convictions (Hooker 2009: 148-50; Murphy 2000: 15).

The requirement of impartiality states that the only perspective from which we can judge the rightness and wrongness of acts is impartial perspective, where everyone's well-being has exactly the same moral weight. It means that every person's well-being should count as one. My well-being has exactly the same moral importance as the well-being of any other person in the world. Yet, should we not sometimes give more weight to our own well-being or to the well-being of people who are near and dear to us? Act-consequentialists think that we should not. The perspective of morality is not the perspective of each and every agent but the perspective of everyone taken as a whole. To judge the right thing in moral matters is, as Henry Sidgwick famously used to say, to look at the things from "the point of view of the universe" (see Singer 1993: 317; cf. Nagel 2010).

Though intuitive and seemingly fair, the impartiality requirement of treating everyone as having absolutely the same moral weight, in conjunction with the maximization principle, has some implications which go against our deepest moral convictions. In order to understand why, it will be useful here to introduce a distinction between agent-neutral and agent-relative theories (Portmore 2011). Agent-neutral theories give every agent the exact set of substantive moral aims while agent-relative ones do not. Since act-consequentialism sets as an ultimate moral aim to each and every agent to maximize the overall good whenever possible, it is an agent-neutral theory. Being an agent-neutral theory, act-consequentialism does not allow for agent-centered options, that is, the moral options to act so as to bring things better for oneself or near and dear to oneself but not the best overall. The reason why this is the case is because the impartiality requirement suggests that all people matter the same and that we should, thus,

strive to maximize the overall good, regardless of the means used for that maximization. In certain cases, that can bring unacceptable implications.

To illustrate such implications suppose that, in the previous example, Anna can choose either to spend the money from her salary by donating it to Oxfam or by buying an expensive medicine for her sick child. Suppose also that, if Anna decides to give money to Oxfam, five other children in Africa will live, but her own child will die. From an impartial perspective, the lives of five children outweigh the one life of Anna's child. Impartial maximizing act-consequentialism does nothing less than require Anna to give up the money she needs for her child's medicine. If it is true that we should always bring about as much overall good as possible and if it is true that every person's well-being has the same moral worth, then it has to follow that failing to save five children is worse than failing to save only one.

Nevertheless, almost everyone would agree that Anna should prioritize her child. It would be outrageous of her, most people would think, if she leaves her child to die only in order to save other children who she has never seen in her life. It would be unreasonably demanding to ask a mother to let her child die. It is a universally held belief that we can prioritize the members of our family, friends, and other near and dear people around us. Since the common-sense morality recognizes the existence of agent-centered options whereas act-consequentialism denies it, it faces a serious problem. So how can consequentialists reply?¹

One might say that whereas simple act-consequentialism will have to bite the bullet here, Singer's second premise need not. Singer contends that if the only way to prevent suffering and death from famine and similar preventable causes is that we sacrifice something of greater moral importance then we ought not to do it (Singer 1972: 231). The question then

¹ Common-sense morality also allows for agent-centered constraints, such as the constraints against murder of innocent. But act-consequentialism does not recognize such constraints since it determines the rightness and wrongness of actions solely by the overall consequences they produce or fail to produce. This theory implies that, for example, one should kill an innocent person if that is the only way for achieving better overall state of affairs, such as, for example, saving several other people. This raises another important objection to act-consequentialism, the so-called permissibility objection. Due to the fact that the principal aim of my thesis is to examine the forcefulness of the demandingness objection, I will not address the permissibility objection any further.

is whether by sacrificing her own child in order to save five distant children, Anna is sacrificing something of greater moral importance. The answer will depend on the point of view from which the comparison is conducted. As stated above, from the impersonal point of view in which each child counts as one, it is simply false that saving one child is more important than saving five. But from Anna's point of view, it is plausible to assume that she values her child more than five strangers' children. In his original paper, Singer did not tackle this issue at length.

However, in his most recent book, Singer writes that "no principle of obligation is going to be widely accepted unless it recognizes that parents will and should love their own children more than the children of strangers, and, for that reason, will meet the basic needs of their children before they meet the needs of strangers" (Singer 2009: 139). This, however, does not mean, Singer contends, that parents can provide luxuries for their children ahead of the basic needs of others. Anna is, then, morally allowed to prioritize her child paying for much needed expensive medicine though she is not allowed to pay for, for example, piano lessons since there is a much more morally urgent way to spend this money.

However, this solution is obviously ad-hoc. Singer proposes that our moral obligations towards the global poor that he inferred from his argument should be changed so that they comply with what people would generally accept. However, we cannot change the requirements of morality simply because there are many people who would not accept it. Such a solution would be especially problematic for act-consequentialists who believe that once we have worked out the valid principles of morality it does not matter whether people will in general comply with them or not.

Another specific feature of impartial maximizing act-consequentialism is that it does not allow for supererogatory acts. Supererogatory acts are the acts that we can do but we are not morally required to do. Failing to do such acts does not signify that we are doing something

morally wrong. However, according to act-consequentialism, any act which does not bring about the best possible outcome is morally wrong. It means then that impartial maximizing act-consequentialism denies the possibility of supererogation.

Our common-sense intuitions tell us that we should help in reducing global poverty but not at the price of living a completely ascetic life. Doing more than usual is praiseworthy but it is not what morality requires of us. In other words, unordinary sacrifices are supererogatory (Kagan 1984: 239). We can be moral heroes if we choose, but that is not what morality requires from us. The problem with act-consequentialism is that it does not recognize the existence of supererogatory acts. All acts that maximize overall good are morally required and all that fail to do so are simply morally wrong. Whenever we fail to bring about the best overall result, we are acting wrongly. Since we are failing to do so most of the time, we are living highly immoral lives.

To illustrate this, suppose once again that our Anna, who spends almost all her income to fight global poverty, has some spare money. Instead of giving it to Oxfam she decides to buy a book and give it to her poor neighbor's child, who she knows loves to read but does not have any money to buy books. Our moral convictions would tell us that Anna's act is praiseworthy but not morally required. Act-consequentialists, however, would say that her act was simply morally wrong because it failed to bring about the best possible outcome.

The demandingness objection does not mean to suggest that significant sacrifice in favor of strangers is not a morally right act. It aims to draw a distinction between acts which are morally required and the acts which are praiseworthy but morality does not require them. To put it in a different way, due to the fact that act-consequentialist impartial maximization principle does not give space for supererogation, it implies demands on agents that are unreasonably demanding.

There is an indirect way for act-consequentialists to allow for supererogatory acts and

possibly get around the demandingness objection. Singer, for example, distinguishes between what morality requires us to do and what we have decisive reasons to do (Singer 1993: 326; cf. Nagel 2010). Singer sticks to his old act-consequentialist opinion that an action that does not bring the best outcome in terms of overall goodness and, in particular, the one that does not forgo something for the sake of achieving some morally more important goal, is morally wrong. He claims, however, that this is only what morality suggests we should do, not necessarily what we have most reason to do. We might sometimes have good reasons not to do what morality requires us to do and we might have reasons not to act in line with Singer's second premise.

It is difficult to see what would be the cases in which reasons for actions deviate from the requirements of morality. Singer can say that if Anna spends the money to save her sick child when she could otherwise spend it to save several other distant children she is doing a morally wrong thing but not the thing she does not have decisive reason to do. But then, Singer owes us an account of reasons for action and, more importantly, an explanation when these reasons are decisive and when they are not. Since they reject supererogation, act-consequentialists cannot say that decisiveness is, in these cases, determined by presence or absences of those who are near and dear to us. Until the plausible and characteristically consequentialist account of decisive reasons is provided, this cannot be a way out for Singer and others.

1.2 The Demandingness Objection and the Integrity Objection

In order for us to understand better the counterintuitive implications of impartial maximizing act-consequentialism it is useful to compare the demandingness objection with another one very similar objection, the so-called integrity objection (see Mulgan 2001: 15–6). According

to this objection, the life of full compliance with the requirements of act-consequentialism would be very hard, if not unbearable, to live. It would undermine the integrity of the agent's life, alienating her from her personal projects and commitments. Bernard Williams famously notices: "how can a man, as a utilitarian agent, come to regard as one satisfaction among others, and a dispensable one, a project or attitude round which he has built his life, just because someone else's project have so structured the causal scene that that is how the utilitarian sum comes out?" (Williams 1973: 116) The integrity objection rests on a particular view of a meaningful life as a life of active engagement in projects of worth. Due to the fact that no agent following act-consequentialism would be able to pursue her own projects, no such agent could live a meaningful life.

Mulgan distinguishes between the two functions of the integrity objection (Mulgan 1999: 16). First, the objection serves as an example of unreasonable demands of act-consequentialism. Imagine that each of us starts doing what consequentialist morality requires of us right now. Our lives would become changed so rapidly and to the extent that neither we nor anyone else who knows us could perceive them as characteristically ours. Second, the integrity objection provides an explanation of this unreasonableness. The reason why act-consequentialism is so demanding is that it ignores the moral significance of personal projects and relations.

Imagine now that Anna becomes convinced in the simple truth of act-consequentialism, deciding to promote the overall good whenever possible. She decides to live an ascetic life, sacrificing almost all of the activities she enjoyed in the past. Instead of buying ballet tickets and paying for piano lessons for her child, she transfers the money to Oxfam's aid account. If Anna is to continue living her life in the way described, it is almost certain that her life will stop flourishing and turn into a very miserable one. The integrity objection rests on a particular view of a meaningful life as a life of personal active engagement of projects of

worth. Since no agent who followed consequentialism would be able to identify with her own projects, no such agent could live a meaningful life (Mulgan 2001: 15).

Remember that Singer's principle of moral comparison leaves open whether and how in problematic cases one's own interests can outweigh the impartial ones. However, in cases with roughly comparable moral values, the principle seems straightforwardly action-guiding: whenever a value of something is not nearly as important as the one we ought to promote, morality requires from us to sacrifice it (see Singer 1972: 231; Singer 2009: 17-8). But suppose that someone replies saying that a meaningful life consists of fulfilling the projects and desires that are vital to our personality. If Anna stops going to ballet she will not be the same Anna she used to be. If she acts according to Singer's ideals, her life will transform into a single-minded pursuit of preventing the badness and suffering of global poverty. That cannot be the good life by any reasonable account.

In this chapter, I have closely examined the demandingness objection, its basis and its implications. The variations of this objection, also known as the near and dear objection and the integrity objection, demonstrated how demands of act-consequentialism in general and Singer's account in particular conflict with our moral intuitions, and hence, undermines the plausibility of the theory itself. In the next chapter, I will turn to the first of the three act-consequentialist strategies of refuting the demandingness objection. This strategy takes a revisionist line. It rejects the demandingness objection by denying the relevance of our common-sense intuitions.

Chapter 2: Extremism

In this chapter, I will present two extremist arguments aimed at refuting the demandingness objection. By extremism I understand the version of act-consequentialism, defended by Peter Singer and Peter Unger, which rejects the soundness of the demandingness objection (cf. Kagan 1991). What is common to extremist arguments is that they take it as given that act-consequentialism implies great demands, but they then try to show that such demands are not unreasonable. They do so by denying the relevance of our moral intuitions. I will start by outlying Singer's and Unger's general argument and then present two different strategies that these authors respectively use in order to pounce on the reliability of our moral intuitions.

2.1 Singer's General Argument

Following Richard Hare (see Hare 1981), Singer formulates his view of act-consequentialism in accordance with the general two-level theory about the content of morality. The higher level defines normative content of morality and provides guidance for moral behavior based on strict impartial act-consequentialist maximization principle. It is according to this guidance that we are supposed to decide whether an act is morally wrong or morally right. The lower level takes into consideration empirical grounds and defines rules of compliances with the higher level. These rules need not be the ones of impartial act-consequentialist maximization.

Singer believes that, in order to bring the best possible overall consequences, we must never infringe the rules even though, in some particular occasions, by breaking a rule, we would bring the best possible outcome contrary to the outcome possibly brought about by our complying with this rule. That is why the lower level of the theory prescribes people to

develop specific character traits which would make breaking the rules for them something impermissible even to deliberate about (Singer 1993: x; see also Kamm 2007: 401-2).

There is another two-level theory in which Singer develops the view on people's motivations to lead a moral life. Singer makes a distinction between the two levels of moral reasoning: intuitive (lower) level and the more reflective or critical (higher) level of moral thinking (Singer 1993: x). At the higher level, he claims, moral motivation comes from self-interest. Leading a moral life is in our own interests and it is the right way to achieve a meaningful life.² At the lower level, people should develop moral conscientiousness, that is to say, they should be moral for the sake of being moral and not for the sake of their own interests (Kamm 2007: 402). We can put Singer's theory of moral motivation aside here.

Singer justifies the first two-level theory by the classical act-consequentialist argument which states that strict impartiality is the most obvious moral point of view. The partial point of view is morally permissible only at the lower level of normative theory and only if it serves to produce the best outcome impartially understood. Morality, then, is all about equal consideration of the relevant interests. And the natural way to take everyone's interests into account is to maximize their overall satisfaction (see Singer 1993: 13).

How is this maximization to be achieved? In general, act-consequentialism can be spelled out in both negative and positive ways. Positive act-consequentialism suggests that we should always strive to maximize the total net amount of value. In its negative variant, it proposes that we should always strive to minimize suffering. Singer's second premise of his argument about our obligations towards the poor is an instance of the latter version of act-consequentialism. The principle of beneficence states that "if it is in our power to prevent something bad from happening, without thereby sacrificing anything of comparable moral importance, we ought, morally, to do it" (Singer 1972: 231). Similarly, Unger argues that "if

² Singer believes that the meaningful life is an independent good which every person should strive for. Happiness, understood as preference satisfaction, is a less significant good comparing to meaningful life.

we are able to provide assistance to an innocent person in great need, at negligible cost to ourselves, then we ought to do so” (Unger 1996: 7). We should, therefore, only minimize the suffering around the world without trying to maximize the overall happiness or benefits.

Singer defends this principle by a thought experiment. He asks you to imagine that, while walking to work, you notice a child drowning in a shallow pond. There is no one else who could save the child and, unless you jump into the water, the child will drown. However, if you jump, you will ruin your new expensive shoes. Now Singer concludes that there is no doubt that you have to save the child even if you will ruin your new shoes since the value of child’s life is morally incomparable to the value of your shoes (Singer 1972: 231). Any other response would be morally unacceptable. In the same way, Singer asks, why would there be any moral difference between saving the drowning child and the one who dies from hunger or minor diseases in a third-world country? That child could be easily saved by transferring similar amount of money as needed for a pair of expensive shoes. We then have obligation to prevent suffering and death happening around the world whenever we can, or more precisely, whenever possible if by doing so we do not make our lives as bad as the ones we wanted to save. In other words, the required demand cannot be unreasonable.

One might respond that Singer’s argument is too hasty. There is a difference between a child dying in front of you and a child dying several thousand kilometers from you. Whether this difference is morally relevant is a difficult issue, which I will not discuss here (see Singer 1993: 232–4). Suppose that distance is not morally relevant and that the mere fact that some people around the world are dying from preventable causes is sufficient to generate our moral obligations towards them. If that is true, then the extremist conclusion seems to follow naturally. However, once established the plausibility of their principles, extremists invoke a further argument in order to discard all intuitions, judgments, and motivations that depart from or conflict with their principles. Most people remain unconvinced by the extremists’

initial argument because they firmly believe that its implications are unreasonably demanding. This belief is, at least, backed up by our common sense moral intuitions. The most popular extremist argument, therefore, is to argue against the reliability of our ordinary intuitions.

2.2 The Argument from Conflicting Intuitions

There are two versions of this additional argument. The first version holds that our intuitions are unreliable because they are mutually conflicting. Peter Unger claims that a straightforward way to demonstrate this is to apply what he calls the method of several options (Unger 1996: 88–91). This method demonstrates that our intuitions in the same or analogous moral cases can be conflicting. The method of several options was supported as successful in debunking our moral intuitions by other extremists, most notably by Singer (see Singer 1999). Unger demonstrates this conflictedness by enriching standard two-option scenarios of moral dilemmas with two or more middle options. After the introduction of additional options, our moral intuitions about some cases supposedly come out to be diametric to those we had about the same case before other options were introduced. Since the number of options or cases is by itself a morally irrelevant thing, Unger concludes that there is a conflict among ordinary moral intuitions. In order to understand the method better, let us briefly reproduce one of his four-option scenarios, called Swiches and Skates (Unger 1996: 89-90).

Suppose an empty trolley is rolling down the main track. The first option is the passive one. The agent does nothing and the trolley rolls over and inflicts death on six people who, through no fault of their own, are trapped on the track. The second option gives agent the possibility to switch the remote control and thus redirect the empty runaway trolley to another track where three people are trapped, through no fault of their own. The third option offers agent to switch the remote control toggle and another runaway trolley with two heavy people

trapped in it will be redirected from the feed track to the main track. As a result, these two trolleys will collide which will make them stop but both persons trapped in one of the trolleys will die. Finally, there is a fourth option, which allows us to take the life of only one innocent person. According to this option, a heavy skater skates down the path which goes right across the main track. The agent turns the remote control lever and thus starts up the skates of the skater who, as a result, collides with an empty trolley and dies. Nevertheless, this will stop the trolley and six innocent people trapped on down on the main track will live (Unger 1996: 90).

Now, Unger plausibly says that most of our so-called trolley problems give us only two options, the passive one, analogous to the first option in the Switches and Skates scenario, and the active one, analogous to the fourth option in the same scenario. In other words, an agent can either be passive and six innocent people will die or he can start up the skates of the heavy skater by turning the remote control lever and thus kill one person for the sake of saving six.³ Exposed to such two-option cases, the moral intuitions of the majority of people would support the belief that turning the remote control and sentencing an innocent heavy skater to death is wrong. Unger, however, claims that if we enrich the same scenario with two middle options such as, for example, the second and the third option in the later Switches and Skates example, people tend to reply contrary to their initial intuitions. Unger gives the names to these intuitions as first-order intuitions and the second-order intuitions accordingly.

Unger believes that, since the cases with two or more options are analogous, there is absolutely no reason for people to change their intuitions on moral permissibility of actions. The question that Unger asks is why adding or deleting alternative options should affect our intuitive judgments in the previous example.

Unger is himself trying to find an answer to this question by working out two hypotheses: the Liberation Hypothesis and the Fanaticism Hypothesis (Unger 1996: 94).

³ In case of two-option scenarios there are neither feed track, nor another trolley with two people trapped in it accordingly.

According to the Fanaticism Hypothesis, our basic moral values, which require from us act so as to lessen the suffering of innocent people, are morally disoriented by other, not so central values, which require us not to harm seriously other people. Unger, however, concludes that the Liberation Hypothesis is the plausible one. According to it, middle options help us to liberate our basic moral values from contrary psychological factors which in two-option cases constrain us from making the choice in accordance with our basic moral values.

But what are those psychological and at the same time morally irrelevant factors which prevent us from giving the consequentialist answer? Supposedly, influenced by one of the strong critiques of utilitarianism, more precisely, by the fact that utilitarians do not give justice to the separateness of persons (see Rawls 1971: 27), Unger claims that it is projective grouping that conduces us to change our initially non-consequentialist intuitions into the consequentialist ones. His argument can be summarized as follows:

- (1) Our intuitions change in Switches and Skates from the two-option case with respect to pushing someone into the trolley, so that we think it is permissible to push the one person into the trolley;
- (2) It is projective grouping of the one person with the six that leads us to think that it is permissible to do anything necessary to him to help minimize mortal loss to the six;
- (3) Projective separating leads us to believe we may not involve someone in this way;
- (4) If it is permissible to act when people are grouped, then it is permissible to act in the same way when they are separate (Kamm 1998: 299).

Let me briefly explain Unger's projective grouping argument. Unger believes that in such cases as when we are confronted with classical moral dilemmas, that is to say, when the trolley would roll over and kill three persons or, if the track is switched, it would roll over and kill the one, three and one are grouped because they share the trolley tracks, connected with each other. The trolley tracks is not the only fact which can group the one and the six. The common threat, for example, can also be such a factor. When grouped, people are involved with each other and thus it becomes fair to harm one person in order to save six (Kamm 1998: 296).

There are cases, however, where all the parties are not grouped. The classical example of such cases is the footbridge moral dilemma. Unger's two-option version of Switches and Skates case is analogous to it. Thus, being the case where initially one skater and the six people trapped on the main track are not grouped, Unger creates two middle options in order for them to become a bridge to the situation when the parties are grouped. Thus, the method of several options helps us to overcome projective separating by employing two middle options as a sort of a bridge between the case in which we cannot look at one person and six people as grouped to the same case but in which we can.

Before starting to elaborate on the details of bridging, another important point is worth mentioning. Unger believes that redirection of things (people) which are already in motion is more permissible than redirection of those which are motionless. Now, the bridging, as it is, proceeds in the following way. It is permissible to redirect the trolley with two people trapped in it from the feed track to the main one just because this moral dilemma has many features in common with the classical trolley case.⁴ Since starting up the skates of the skater has some characteristics of redirecting the trolley with two people trapped in it, more specifically, the

⁴ This moral dilemma is composed of two options: letting the empty trolley to roll over and kill six people trapped on the main track or redirecting another trolley with two people trapped in it, and is analogous to the classical trolley case when the choice we are supposed to make is between the letting the trolley roll over and kill three people or redirecting the trolley on another track where one person would be killed instead.

fact that the skater would be in motion once we start up her skates, then we are supposed to find it permissible to do so.

Simply put, the argument proceed as follows. According to previous research (see Green 2008), many people find it morally justified to redirect the trolley so that it would roll over and kill less people than it would otherwise do. So the Switches and Skates four-option case, if we look at it in detail, consists of three moral dilemmas: (1) to leave six people to die or to redirect the empty trolley and kill three instead; (2) to leave six people to die or to redirect the trolley with two people trapped in it which will result in their death; or (3) to leave six people to die or kill one skater through starting up his skates. Since (2) has some common properties with (1) and since (1) is analogous to the classical trolley problem and therefore permissible, (2) becomes permissible too. Then (3), by sharing some common characteristics with (2), which is permissible, becomes permissible too.

There are several problems with Unger's method of several options which make his argument from conflicting intuitions unsound. First of all, it is far from being true that our intuitions change when we are exposed to several-options cases. Kamm, Mulgan and I included did not share conflict of intuitions in the case of Switches and Skates (cf. Kamm 1998: 299-300; Mulgan 2001: 30-1). The experimental data are not transparent. Unger writes:

Though it certainly doesn't obtain with all ethically sensitive respondents, a fair number, myself included, find these fluctuating reactions to obtain, time after time, despite the effects of presentational order. In this respect, our experience here is quite like what we've had with many (only) moderately effective optical illusions (Unger 1996: 92, n.6).

Moreover, Mulgan's multiple attempts to reproduce the experiment were unsuccessful.⁵ However, let us, for the sake of argument, assume that Unger's Switches and Skates four-option case does demonstrate the conflictedness of our common sense moral intuitions. Is there any other way than the appeal to intuitions to refute his argument?

A more substantial critique of Unger's conclusion about unreliability of moral intuitions is the following. Unger believes that people change their moral intuitions on the basis of morally irrelevant factors. He believes that if the parties in moral dilemmas are grouped, then the mortal harm we inflict on one person in order to save other six is morally irrelevant. That is why Unger believes that, if he succeeds in grouping the parties, which he does using two middle options, it would do the trick. People would change their moral intuitions regarding permissibility of actions on the basis of a morally irrelevant factor, which is to harm one person mortally, of course, only if she is grouped with six others, hence is a fair game.

Kamm plausibly points out an obvious flaw in Unger's argument by challenging his belief that parties in moral dilemmas, if grouped, are involved in a fair game and, therefore, it is fair that one party should be sacrificed for the sake of another one. Hence grouping and especially, a projective one, is far from being sufficient for obtaining the permissibility to harm one person in order to save several others. Besides, Unger also erroneously suggests that the only factor which induces us to resolve a moral dilemma in favour of non-consequentialist' solution is the separateness of persons. Therefore, as it was said above, grouping people is not enough for switching our moral intuitions (Kamm 1998: 299).

This means, that even if Unger's method of several options does lead to the switch of our moral intuitions, it is far from being obvious that the switch was inflicted by morally irrelevant factors. As Mulgan suggests, this might simply mean that our common sense morality is messier than we are prone to think and without knowing the truth we cannot really

⁵ Mulgan claims that he was presenting Unger's scenario for several years to his undergraduate students and that majority of them did not share Unger's intuitions (Mulgan 2001: 31).

know whether the case alteration leads to the positive distortion of values or to the negative one (Mulgan 2000: 400). Maybe we found (1) intuitively permissible and (3) not exactly because of some properties which they do not share together even through linking by (2). The difference between (1), (2) and (3) is morally significant and their linkage in a way in which Unger does it, is illegitimate.

As it was mentioned above, Unger uses the separateness of persons objection as a psychological factor which prevents people from giving the ‘correct’ consequentialists answers. Kamm plausibly argues that he misuses this objection which in the first place aimed to show that the mortal sacrifice made by one person produces the good which obviously cannot be benefited by the very same person but only by those others for the sake of whom this sacrifice had been made. By grouping people, Unger believes we can solve the problem because the sacrifice made by one member of the group benefits other members of the same group and thus, he claims, makes the ‘game’ be fair. The separateness of persons objection, however, stipulates intra-personal harm for benefit and not intra-group-al one which, in itself, constitutes in inter-personal harm for benefit exchange.

There is another way in which Unger’s projective grouping for generating permissibility of action is unsound (see Kamm 1998: 302). For him, people are separate if they do not share some common features. If they do, what supposedly can be achieved through projective grouping, then mortal harm to one for the benefit of others is morally irrelevant. In order to show that Unger is wrong, Kamm challenges Unger’s claim that, by grouping parties in moral dilemmas, people can change their intuitions.

Suppose that A, B, and C are swimming in the pond. Suppose also that not far from them but at the same time near the ground there is D standing in the water and collecting leeches for his ill grandmother. Now, I assume that Unger would consider all these people to be grouped by the water of the pond they are in. Imagine then that you, while fishing, spotted

a crocodile fastly approaching A, B, and C. You realise that, if you tell D that there are many more leeches on the other side of the pond where she would immediately swim to, D would be eaten by the crocodile first since she would be closer to it than the other three. Meanwhile, you would alert A, B and C who would have enough time to come out of the water while the crocodile is eating D. Do you think it is permissible to kill D in order to save A, B and C? I think that it is not. Justice to be told, I should mention that not everyone would share my intuition. Some people with consequentialist-friendly intuitions would find it permissible to send D to be eaten by crocodile. However, the important point of this thought experiment is to show that, the fact that all four of these people are grouped, according to Unger, by the pond does not add to the permissibility of sending D to die (cf. Kamm 1998: 303).

The same objection can be applied to Unger's claim that people intuitively feel it to be more permissible to redirect something (or someone) which is already in motion than when it is motionless. If we think that in the classical trolley case it is permissible to redirect the trolley so that it would roll over and kill three innocent people instead of six, my feeling is that the fact that the trolley is already in motion has nothing to do with people's consequentialist intuitions.

Unger's method of several options in general and his methodology of projective grouping in particular do not provide sound proof of unreliability of our moral intuitions. Let us turn now to the second version of extremist attack on moral intuitions.

2.3 The Argument from fMRI

Another version of the extremist argument appeals to the recent studies on moral intuitions by cognitive psychologists using the functional magnetic resonance imagining (fMRI). These studies show that there is a statistically significant relation between moral intuitions people

have and the particular parts of the brain being active when people express judgments relying on these intuitions (see Greene 2008).

The working assumption of the research conducted by Joshua Green and his colleagues was that we can distinguish between two types of moral judgments. The first are ones where we think that what is important for morality is bringing about the best possible outcome ('typically consequentialist judgments'). The second are ones where we believe that there are reasons not to bring the best outcome for the sake of achieving some other desirable moral aim, such as, for example, not using someone merely as a means or respecting the rights of that person ('typically deontological judgments'). The fMRI screenings showed that each time that the examinees expressed consequentialist judgments about certain moral dilemmas,⁶ the parts of the brain associated with working memory and rationality demonstrated activity and whenever these people were giving deontological judgments, the emotional parts of the brain lighted up on the screen (Greene 2008: 43–4).

The findings were interpreted as showing that the brain regions associated with what is considered to be 'higher cognition' are responsible for our ordinary consequentialist judgments, while the deontological judgments were under the responsibility of the emotional parts of the brain, that is, the parts which are responsible for triggering our intuitions. What is more, Greene concludes, that these findings give us good grounds to discard deontological judgments in favor of consequentialist ones because they are based on moral intuitions which, in turn, are based on our emotions (Greene 2008: 233). If this conclusion is sound, act-consequentialists have a strong argument against the reliability of our ordinary moral intuitions. As a matter of fact, Greene notices that, if his argument from the fMRI research is sound, Singer's argument for our obligations to the world's poor gets additional support because the common-sense intuition that we should help the starving children but not for the

⁶ The trolley problem and the footbridge dilemma (see Foot 1967; Thomson 1985).

sake of overriding some personal priorities becomes no longer credible (Greene 2008: 70). Similarly, in one of his recent papers, Singer quotes approvingly Green's research, stating that the research casts doubt on the usefulness of reflective equilibrium as a methodological tool widely shared among normative ethicists (Singer 2005: 339ff).

Green's argument proceeds as follows. As mentioned before, Greene makes two relevant distinctions: between two types of moral judgments and two types of psychological process underlying these judgments. The first two are 'characteristically' consequentialist and deontological judgments (see Greene 2008: 38–9). Deontologists would say that we are not permitted to push the large person in order to save the five workmen. Consequentialists would not care about the rights of either the trapped workmen or the one on the footbridge, as long as by killing him, we save the other five. Our common intuitions suggest that we should not push the large person in the second case, so Greene names the judgments relying on these intuitions deontological. In general, deontological judgments would include all those judgments that express the belief that some actions are just wrong despite the benefit they produce (Greene 2008: 39). These judgments concord the most with what we typically call the common sense morality, principles of which are in its turn almost always grounded on our moral intuitions. On the other hand, our moral intuitions also suggest that it is permissible to redirect the trolley and thereby cause the death of one person. These intuitions support the judgments that Greene calls consequentialist.⁷

Another distinction that Greene introduces is between the types of psychological processes that underlie our judgments in moral dilemmas. This distinction is rather uncontroversial. Among neuroscientists there is almost a universal consensus that certain parts of our brain are associated with abstract reasoning, problem solving, self-control,

⁷ Frances Kamm criticizes this distinction on the basis that deontological theories may in some cases allow for killing one person in order to save the five (the exceptions is absolute deontology which is nowadays rarely argued for). What deontologists claim, Kamm writes, is that we are not morally required always to save the greater number of people by killing some but that we are sometimes permitted to do so (Kamm 2009: 337).

planning and manipulating information in working memory while other parts are associated with our emotions (Greene 2008: 40–1). Green calls the former processes ‘cognitive’ and the latter processes ‘emotional’. Cognitive processes tend to be slower, flexible, based on sufficient amount of information, and responding to abstract and more general subject matters. Emotional processes are instead fast, frugal, based on limited amount of information, and responding to particular subject matters (Berker 2009: 300).

The working hypothesis, confirmed by the fMRI results, was that the deontological judgments are triggered by the emotional processes, while the consequentialist judgments are triggered by the cognitive processes. This claim is usually called the Dual-Process Hypothesis (henceforth the DPH). The DPH predicts that when a person is presented with the trolley case, the cognitive parts of the brain should show themselves as active on the fMRI machine while when one is presented with the footbridge case, the emotional parts should light up. As Greene and others tell us, this is what actually happened in several experiments (Green 2008: 43–6).

Greene confirmed these findings by testing people’s intuitions in moral dilemmas similar to the classic trolley problems. Dilemmas analogous to the footbridge case were named ‘personal’ because they involved bringing harm in an ‘up-close and personal way’ and dilemmas similar to the initial trolley case were named ‘impersonal’ because of the impersonal nature of the harm present in them. By ‘personal’ harm, Greene refers to harms involving direct violation of bodily integrity of a particular person and not being a consequence of deflecting an existing threat on another party (Greene 2001: 2107). Harms falling outside these criteria were named ‘impersonal’ (e.g. you are redirecting the trolley by pulling the lever and only indirectly killing one workman). Moreover, it was assumed that deontological judgments are in general concerned with personal interactions, while consequentialist judgments are more impersonal in their nature.

Examinees were then presented with personal and impersonal dilemmas and their responses were monitored via fMRI. In personal moral dilemmas, such as the footbridge case, examinees were mostly giving deontological judgments. This was confirmed by increasing activity of the ‘emotional’ parts of their brains on the screen. In impersonal moral dilemmas, such as the initial trolley case, examinees’ judgments were consequentialist in a sense that, when responded to the dilemmas, they exhibited excessive activity in the ‘cognitive’ regions of the brain (Greene 2001: 2106; Greene 2008: 46). In other words, the results confirmed the DPH.

This is all that was demonstrated by the fMRI findings of Greene and his colleagues. Greene interpreted these findings as showing that characteristically deontological moral intuitions should be taken as a rather untrustworthy basis for our moral judgments (see Greene 2008: 66–72; Singer 2005). If his claim is true, it gives an indirect support to act-consequentialist moral theory, which was often accused as having counterintuitive outcomes.

Greene’s argument is the following. He claims that, in difficult moral dilemmas such as the trolley problem, the correct moral answer must somehow be the one concerned with impersonal rather than personal factors and the consequentialist answers are by definition impersonal. The reason why Greene believes moral judgments involving personal considerations are dubious is because they reflect the influence of morally irrelevant factors (Greene 2008: 75). One of these morally irrelevant factors, according to Greene, is distance between an agent and a person whose life depends on the outcome of moral judgment made by the agent. Singer echoes Greene’s thought by saying that there is no morally relevant difference between the trolley case and the footbridge case (Singer 2005: 347-48).

If we formulate it this way, the central claim seems to lack substantive argumentation. It is obvious that fMRI findings themselves cannot give us an answer to the question why it is the case that deontological intuitions stand on morally irrelevant factors. For that conclusion,

we need substantive normative argumentation and neither Greene nor Singer are offering one, assuming something what they have to prove yet.

In one of his recent lectures (Greene 2010), Greene claims that, although no normative ‘ought’ can be derived from a scientific ‘is’, scientific data can have normative implications in a sense that it can conduce people to change values that they hold (Greene 2010: 9). Greene gives an example of a normative claim that juries in capital cases often make unfair decisions. This claim can be inferred from an evidence-supported descriptive claim that juries’ decisions are affected by the race of the defendant and a further normative claim that defendant’s race is an irrelevant factor in making decisions in capital cases. Although the descriptive claim alone does not give us a reason to infer the normative conclusion, it can, in Greene’s words, ‘do work’ in determining which values we should have (Greene 2010: 9). That is to say, the descriptive claim helps us to understand why the normative premise is true, that is, why race is a morally irrelevant factor in the particular case.

Similarly to this example, the results of brain imaging that Greene and his colleagues have conducted help us in understanding the relevant normative premises that lead to the normative conclusion that deontological judgments are problematic. Recall that, according to Greene’s argument, deontological judgments, unlike consequentialist judgments, do not have any normative force because they are tied to personal considerations (such as, in the footbridge example, directly pushing the fat guy onto the tracks). This follows from the discovered relationship between people giving deontological judgments while being presented with a personal dilemma. However, since factors that make a dilemma personal rather than impersonal are morally irrelevant, emotions respond to factors that are irrelevant from the moral point of view. Deontological judgments are, therefore, founded on morally irrelevant factors (Berker 2008: 321).

Selim Berker finds the argument unsound because of the problematic normative premise about the moral irrelevance of personal-impersonal factors in moral dilemmas. Berker claims that this premise “appeals to a substantive normative intuition, which presumably one must arrive at from the armchair, rather than directly read off from any experimental results; this is why the argument does not derive an ‘ought’ from an ‘is’” (Berker 2008: 322; cf. Kamm 2009: 334–6). In other words, it is impossible to derive a conclusion about the unreliability of deontological judgments in a sound way without arguing for all the premises. Since among Greene’s premises there is a premise which is not empirical but rather normative, Greene owes us a philosophical argument in order to support it. No fMRI data can provide a support for moral irrelevance of personal-impersonal factors.

Surprisingly enough, Greene agrees with Berker to a certain degree. He believes that it is true that the data gained from fMRI cannot give us good grounds for accepting the truth of the problematic premise. But the data can help us understand better why the premise is true (see Greene 2010: 14). The problematic normative premise claims that it is irrelevant from a moral point of view whether a moral dilemma is personal or impersonal. This means that the factor of ‘personalness’ (that is, the spatial distance and the recognition of the victim as an individual that will be harmed from the action) does not play any significant role in moral dilemmas.

Greene provides the following reason for why he thinks ‘personalness’ is a morally irrelevant factor. First of all, Greene claims that his initial argument can be rephrased so that it mirrors the jury example mentioned above. Descriptive claim (first premise) of the argument says that people’s judgments in trolley cases are strongly influenced by the presence of absence of personalness (or ‘personal force’). This is a claim confirmed by the described experiment, so there is no reason not to take it as true. The second and normative premise, however, is more problematic. It claims that the presence or absence of personal force is

morally irrelevant to the moral permissibility or impermissibility of actions such as pushing the fat guy or pulling the switch. From these two premises, Greene infers the normative conclusion that people's judgments in response to trolley problems are influenced by the factor of personal force, which is morally irrelevant. This, in turn, makes such judgments unreliable (Greene 2011: 16).

Greene then invokes two very contestable claims. First, he believes that somehow the descriptive premise 'does work' in supporting the normative premise. This is, I believe, false in most of the cases presented by Greene, including the jury case. In the jury case, the fact that the jury often makes an unfair decision because it is biased towards the race of the defendant, may serve, in Greene's opinion, to help us understand a normative claim that race itself is a morally irrelevant factor. The latter claim is obviously true. However, it is hard to see how an empirical claim as the one above can give us any reason to see the truth of the normative premise. Greene says that this difficulty arises from the fact that we all accept the normative claim that racism is wrong. Suppose that we do not. Those of us who would question this belief, Greene writes, will have only to consult the descriptive claim, which would challenge their values (Greene 2008: 9). But this is evidently false. If someone is a racist, he will see the fact that juries' decisions are often made on the basis of race of the defendant as a good thing. The racist will believe that juries' decisions are actually fair. The only way for a racist to be proven wrong is, however, to offer him a normative argument showing that race is a morally irrelevant factor.

Similarly, it is hard to see how the cited research can highlight the idea that personalness is a morally irrelevant thing in determining the permissibility of some action. Suppose that someone is presented with the final data of Greene's research and then asked whether she believes that spatial distance and causal closeness should play any role in determining what to do in situations where we can harm someone. The natural response of a

person who has no knowledge of moral philosophy will be that she knows only that personal force actually plays the role in determining people's reactions in footbridge cases. Whether personal force should be relevant or not, she has no idea. Greene, therefore, needs a further argument in order to reach the desirable inference.

Second, Greene makes a remarkable claim that there is a universal consensus among moral philosophers, regardless of whether they are consequentialists or deontologists, that personal factors are irrelevant for permissibility of our actions (Greene 2010: 17). This may be true in the two footbridge cases mentioned above, where you are either pushing the fat guy or hitting a remote switch that will kill him. It does not really matter whether you are killing someone through one or two causal chains (in the second case, you first pull the switch and then the person dies, while in the first, you only push him).

However, it does not follow that personal force is never relevant or that philosophers agree that this is the case. Greene mentions the example of global poverty. He claims that deontological thinking is a major obstacle to ending poverty now (Greene 2010: 24). Act-consequentialists say that it really does not matter whether we should help an African child dying from cholera or a child in our own town dying from cancer. They matter equally and if you can bring more overall good by transferring your money to aid agencies for the sake of helping the world's poor than you should do so whenever possible. Many people, however, stand by their intuitions that you should prioritize those who are near and dear to you. These people as well as some moral philosophers make deontological judgments (see Kamm 2007: ch. 12). At this point, it does not really matter whether deontological philosophers are right about this or not. As long as there is a debate around the question of moral relevance of personal force and proximity, Greene needs to provide a normative argument for supporting his claim.

Thus, the results of the fMRI research do show that people react differently in personal and impersonal dilemmas, but not that the former rest on morally irrelevant factors. To reach this conclusion, we need normative arguments which Greene never offered. The argument from fMRI, therefore, does not support the extremist contention about the unreliability of our moral intuitions.

One of the implications of the unreasonable demandingness of act-consequentialism is its counterintuitiveness. In order to defend their view, extremists among act-consequentialists often appeal to arguments intended to undermine the reliability of moral intuitions. In this chapter, I have examined two of such arguments: Peter Unger's argument from conflicting moral intuitions and Joshua Green's argument from fMRI, supported by Peter Singer. Both arguments were proved to be unsound.

Trying to demonstrate the unreliability of our moral intuitions, Unger created a method of several options. He claims that persons, when exposed to a moral dilemma with more than two options, change their moral intuitions for opposite ones in exactly the same cases. Unger also claims that this happens because middle options create a sort of linkage between the parties in moral dilemmas, thus grouping them into a fair game, which, in its turn, makes mortal harm inflicted on one in order to save six, morally irrelevant. As it was shown by Kamm, Unger's projective grouping is not the kind of grouping that would involve the parties into a fair game and thus, make the mortal harm to the one person morally irrelevant.

Greene's argument against the reliability of moral intuitions, based on fMRI research, fails to provide the evidence that choices people make in personal moral dilemmas rest on morally irrelevant factors. This is a normative assumption present throughout Greene's and Singer's argument but we find no further argumentation that would support it in their writings. I, therefore, conclude that extremist attack on the reliability of moral intuitions is unsuccessful and that we can and should rely on our intuitions in moral reasoning.

Chapter 3: Sub-maximizing Views

In this chapter, I will examine two famous strategies to reconcile act-consequentialism with common-sense morality. They present act-consequentialism as a moral theory that does not require bringing about optimal but rather sub-optimal consequences. If these views are correct, they provide a version of act-consequentialism invulnerable to the demandingness objection. I will first focus on the view defended by Michael Slote (see Slote 1984; Slote 1991), the so-called satisficing consequentialism.

3.1 Slote's Satisficing Consequentialism

The general idea of satisficing consequentialism is that an act is morally right if it brings about the consequences that are 'good enough'. 'Satisficing' is a portmanteau word from 'satisfy' and 'suffice'. It is often used in economics to describe the idea that rational agents may sometimes choose to act sub-optimally, that is, the acts that are good enough but not the best one from the available alternatives (Slote 1984: 141). For example, suppose that Anna decides to sell her house because she wants to move to another location. The house is on the market and someone offers a reasonable price for it. Anna finds the price satisfactory and decides to sell the house. Anna could have waited for some time and achieved a better price but she is satisfied with the offered price because she thinks it is a good enough deal. She knows that she could get more, but nevertheless decides to sell the house. Anna's decision is an example of a sub-optimal but still rational action. Although Anna fails to maximize the expected utility, she performs an action that is perfectly rational.

Similarly to economic satisficing, Slote believes, we may choose an action that is sub-optimal but justified from a moral point of view. Consider the following examples. You are a manager of a resort hotel, who discovers, late one evening, that a car has broken down near the hotel. In the car there is a family of four who cannot afford a room to stay for the night. Since you believe that it would be wrong not to offer these people a room free of charge, you decide to do it. However, you choose a room that is good enough to satisfy the needs of a family of four but not the best one available in the hotel. As a matter of fact, there are several other unoccupied rooms that are better than the one you have chosen (Slote 1984: 149–50). Recall that, according to simple act-consequentialism, the morally required act is the one that brings about the best possible consequences in terms of maximizing the overall good. In the described case, you could have produced a better overall outcome if you had chosen the best room available. Thus, according to simple act-consequentialism your act was wrong. However, as Slote argues, our intuitions tell us that there is nothing wrong with your action. Your act of benevolence, although sub-optimal in terms of net goodness, is morally justified.

In the second case, a great Viking warrior who fought meritoriously during his lifetime, dies and arrives at Valhalla. The gods reward him by offering a fulfillment of a single wish for those he left behind in the previous life. He asks the gods to make his family and his descendants to be ‘comfortably well off’ for the rest of their lives. The gods ask him in reply whether he would prefer his family and descendants to be made as well off as possible. He answers negatively. He thinks that making them ‘comfortably well off’ will be good enough (Slote 1984: 150–1; Slote 1991: 27). According to simple act-consequentialism the Viking has acted wrongly. Slote thinks that the example shows the deep counterintuitiveness of this view. We think that there is nothing wrong in Viking’s choice. It seems, then, that agents are sometimes morally justified to choose an action that is not the best one possible in terms of the overall good it produces.

Slote's argument has the following structure. Slote claims that economic satisficing is, in certain circumstances, perfectly rational. Since there is an analogy between moral and economic satisficing, moral satisficing should be permitted in similar circumstances. Slote's satisficing consequentialism, therefore, claims that agents are morally justified in acting in the way that brings about an outcome which is good enough but not the best on the overall scale. Acting sub-optimally is not just permitted but it is the morally right thing to do (Slote 1984: 157).

Satisficing consequentialism is appealing because it accommodates our common-sense intuition that certain actions are supererogatory, that is to say, that they are good to do but not morally required. Satisficing consequentialism complies with the intuition that some actions are beyond the call of duty. Recall that one of the features of the demandingness objection to act-consequentialism was that this view does not allow for supererogatory actions. If an action fails to bring the best possible outcome in terms of the good it brings about, act-consequentialism considers it wrong. If the action achieves this aim, it is right. In many situations, the right action, according to this view, will be the one that is very hard, if not impossible, to achieve. On the other hand, satisficing consequentialism avoids this problem by allowing sub-optimal acts to comply with the criterion of rightness.

Even though Slote's view is an attractive alternative to simple act-consequentialism, it has several flaws. One way to argue against it is to reject its crucial premise about rationality of economic satisficing (see Pettit 1984). I will not go in this direction. I will assume that Slote's argument for satisficing consequentialism is sound and then examine its implications.

The problem with this view is the following. Satisficing consequentialism entails that it is sometimes morally right to do things that will produce much less good than one could, only on the basis that one has produced enough good already. In order to assess this claim it will be first necessary to decipher what does producing enough good for others mean. This

formulation is ambiguous between two different interpretations (see Mulgan 2001: 132). On the first reading, Slote might be claiming that one does enough good for others if and only if she produces a large enough amount of good for each individual who is in a position to benefit from her actions.

However, if this is the right interpretation, Slote's example does not support it sufficiently. Consider the Viking case once again. The Viking requests from the gods to make the members of his family and his descendants comfortably well off. The members of Viking's family, however, are not the only people who can benefit from his decision. There are other people around the world who do not live comfortable lives. If good enough means good enough overall, that is, good enough for each individual who can be benefited, then the Viking must ask from the gods to benefit all those people. This, of course, imposes unreasonable demands on the benefactor. In the particular example, the benefactors are the gods and their resources and time are, supposedly, unlimited. But in any other more realistic case (e.g. the one of the hotel manager), the benefactors will be ordinary people. They will have their own lives, projects, and priorities. However, once one accepts and acts upon this interpretation of satisficing consequentialism, she will have to discard her life plans and priorities for the greater good she might do by making everyone comfortably well off. If this interpretation is true, satisficing consequentialism loses its most appealing feature: the feature of not laying down unreasonable demands.

On the other reading, one does good enough for others if the total amount of good she produces is above a certain threshold (Mulgan 2001: 132). In this case one is allowed not to provide any good to some persons as long as the total amount of good she brought about to others is sufficiently great. Satisficing consequentialism, therefore, does not necessarily impose unreasonable demands providing that the threshold of beneficence is not too high (the threshold needs to be specified in the amount of good to be produced and amount of

individuals that are to receive the good). However, it seems that in situations where an agent can relatively costlessly bring about the greatest possible good overall, she should do it. For example, suppose that the Viking was asked by the gods whether he would prefer all people in the world to be made as well off as possible or that only a certain number of people receive the benefits such that the desired threshold is reached. If the Viking chooses the latter option, he is certainly doing something wrong. He could have saved many other people from poverty and misery through no expense to himself. Although he brought about a good enough outcome, he could have done better. At least in this case, morality requires more than doing just enough.

Satisficing consequentialists, thus, face a dilemma. If they accept the first interpretation of what makes consequences ‘good enough’, their view will be no less unreasonably demanding than simple act-consequentialism. If, instead, they accept the second interpretation, their view will entail morally unacceptable results.

3.2 Scheffler’s Hybrid View

The second version of sub-optimizing act-consequentialism I will look at is Samuel Scheffler’s ‘hybrid view’ (see Scheffler 1982; Scheffler 1992). Similarly to satisficing consequentialism, the hybrid view tries to reconcile simple act-consequentialism with the common-sense morality by making room for agent’s personal projects, needs, and relations. In other words, the hybrid view tries to avoid the demandingness objection by claiming that agent is sometimes allowed to devote her time and efforts to personal projects and activities even if that would mean bringing about a sub-optimal outcome from an impartial perspective.

Scheffler argues for the hybrid view by introducing the notion of the agent-centered prerogative. The prerogative permits the agent to choose to perform the action that will be

sub-optimal from an impartial perspective, that is, that will bring about less than the best state of affairs (Scheffler 1982: 5). According to the hybrid view, agents are not always morally required to bring about the greatest overall good. They are permitted to give greater moral weight to their own interests than to the interests of other people taken as a whole.

How much weight is an agent allowed to put on her own interests in proportion to the interests of others? In order to answer this question, Scheffler introduces two criteria of moral permissibility of sub-optimal acts. According to Scheffler, someone is morally permitted to perform her preferred act P if and only if there is no alternative act A such that A would produce a better overall outcome than P, as judged from an impartial standpoint, and the total net loss to others from doing P rather than A was more than M times as great as the net loss to her of doing A rather than P, where M is a morally allowed ratio that each agent can give to her own interests rather than to the interests of others (Scheffler 1992: 378).

In order to see how the hybrid view works, imagine the following example. An affluent person has three options of how to spend her income. Option (1) is to donate all of her income to Oxfam. Option (2) is to donate most of her income to Oxfam and buy one ticket to the ballet. Option (3) is to donate nothing to Oxfam and to buy several tickets to the ballet. In order to apply Scheffler's principle we need to attribute impersonal value as well as the personal costs to every option. Thus, the impersonal value of the first option would be, let us say, 10,000 while the personal cost of the agent is 10. Giving almost all of the income would produce 9,900 of impersonal value and inflict the cost to the agent equal to 1. Finally, if she decides to give nothing, the impersonal value of such an option would be 5,000 with no personal costs whatsoever. We also suppose that an agent gives 600 times more weight to her own interests than to the interests of strangers.

Now, in order to understand which of the options is the one that is permitted, we need to calculate the extra impersonal value brought by extra donation made by an agent and compare

it to the cost brought to her of the same extra donation. For example, donating most of her income rather than not donating at all would bring 4,900 extra of impersonal value and inflict extra cost to our affluent person equal to 1. Since the agent makes a preference to donation over her personal projects, the extra personal costs inflicted should be multiplied by the ratio that each agent can give to her own interests rather than to the interests of others. This means that the loss to others from an agent choosing not to donate anything would be greater than the loss of an agent from choosing to donate most of her income.⁸ As a result, the affluent person is not permitted not to donate anything.

Let us not see whether our affluent is required to donate all of her income. The extra impersonal value from her donation everything rather than the most of it is only 100 while the extra personal cost would become 5,400.⁹ This means that the agent is not required to donate everything, hence, is allowed to buy one ticket to the ballet (Mulgan 2001: 146-7).

However, there is an internal problem with Scheffler's view. Scheffler does not provide us with any specific information about his variable M . The only guidance he gives us is that M should be constant and that it can be assigned by an agent depending on how much more weight she gives to her personal interests compared to the interests of others (Scheffler 1994: 20; also Mulgan 2001: 151). One might ask two questions:

- (1) How can M be constant, hence, completely independent from the nature of a project?
- (2) How can the approximation of M , which plays a crucial role in determining whether the agent is allowed for agent-centered prerogatives, be left to the mercy of the agent herself?

⁸ The loss to others is what they do not get if an agent chooses her personal project which is 4,900, whereas the personal loss to an agent is M times as great is 600.

⁹ The extra cost is calculated by multiplying the difference of costs by M . $10 - 1 = 9 \times 600 = 5,400$.

To see the implausibility of the issue stressed by the first question consider two examples. Suppose first the following situation A. Anna has donated 80 percents of her monthly income to Oxfam and is trying to determine whether she has the agent-centered prerogative to buy a couch. Suppose also that the amount of money she would spend on it would save one child in Africa from starving to death but Scheffler's hybrid view grants her an agent-centered prerogative for buying this couch because the marginal utility of the money spent for saving one more child in Africa is smaller than the marginal utility of the money spent on the couch, given the amount of Anna's self-assigned M .

Suppose now situation B. Anna has donated 80 percents of her monthly income to Oxfam and is trying to determine whether she has the agent-centered prerogative to buy life saving medicine for her sick child? By the same token, as in situation A, the agent-centered prerogative is granted and Anna may buy the medicine.

Now both projects, buying the couch and buying the medicine for the sick child, are personal and their importance to Anna is accordingly M times greater than the importance of interests of distant strangers. However, what strikes us as inappropriate is that, due to the fact that M is a constant variable, the importance of buying a couch for Anna is supposedly exactly the same as the importance of buying the medicine for her sick child.¹⁰

It is possible that Scheffler's reply to such intuitively repulsive implications would be that, if an agent would be able to assign different importance to different personal projects, his hybrid view would probably be as demanding as simple act-consequentialism. Given the fact that every extra dollar donated to Oxfam is a potential life saver, its marginal utility will almost always be higher with the exception of the cases in which an alternative personal project would be as important as saving human life. How can the approximation of M , which plays a crucial role in determining whether an agent is allowed for agent-centered

¹⁰ For mathematical precision we assume that the amount of money Anna is to spend for a couch is exactly the same as the amount which she needs to buy the medicine for her sick child.

prerogatives, be left to the mercy of the agent herself, is a question that presents a serious problem for Scheffler's hybrid view.

The second question reflects another problem with Scheffler's M. If this variable is completely dependent upon agent's subjective preferences, then the fact about the agent having or not having a moral obligation to do certain thing is partially dependent on those preferences. For M is a variable that determines whether someone has an agent-centered prerogative or not because the personal costs to agent of not doing his personal project (Scheffler's P) are changed with the value of M which, in turn, changes the relation between the personal cost and the impersonal value of doing some other available action (Scheffler's A). So whether an agent is required to do A or she is permitted to do P, ultimately depends on M. And if M is purely subjective, then we can imagine cases where this would fetter certain actions that are obviously morally required.

For example, suppose that I believe that my child should have the new PlayStation and that this project is billion times greater than any other one which can be done with the same amount of money. Suppose, for the sake of the example, that the new Play Station costs 1000 euro and that, for that money, the lives of 10 children in Africa could be saved. The hybrid view implies that I should not spend a cent to help those children. Of course, whether I have an obligation to give all that money to charity is debatable but I certainly have some moral obligations to the dying children. In other words, not only Scheffler's hybrid view is not overdemanding but it can be underdemanding.

To both questions concerning the variable M, Scheffler can respond in the following way. He can agree that holding M constant and grounding it on purely subjective preferences can generate unacceptable implications. However, Scheffler can say that, when determining the permissibility of an agent-centered prerogative, M should be specified with regard to certain objective standards, such as the standard of reasonableness. It is unreasonable to put a

billion times greater weight to your child's fun but it is reasonable to put some importance on your child's life even if the resources you will in that way spend could save several lives of distant strangers.

Consider the following example. Anna is in a position to spend her monthly salary either by buying ballet tickets or by donating to Oxfam. Suppose that the impersonal value of outcome if Anna chooses to donate to Oxfam is 100, and if she chooses to buy the tickets is 2. At the same time, the personal cost for Anna of donating to Oxfam is 1 whereas there is no cost at all from her buying the ticket. Now whether Anna is morally permitted to buy the tickets depends on the weight she is permitted to give to her own interests over the interests of others. It is plausible to assume that Anna can give some weight to her own interests but not 50 times more than to the interests of people in need, even if they are distant strangers. Anna, therefore, cannot have the agent-centered prerogative to buy the tickets.¹¹ However, suppose that, rather than buying ballet tickets, Anna needs to buy the expensive medicine for her sick daughter. Under the assumption that the impersonal value of outcome if she buys the medicine will remain the same as in the previous case, Anna can put more than 50 times more weight to the interests of her child than to the interests of others. In this case, thus, Anna has an agent-centered prerogative to buy the medicine for her child, even though, from the impersonal perspective, she would do more good by donating to Oxfam.

Introducing the notion of reasonableness can distinguish in a very broad way between weights that agents can put on their personal projects in comparison to the interests of others. However, any notion of this kind is vague and cannot give a determinate answer to the question how much weight should we put on our personal projects. If I decide to buy PlayStation for my child and if this is the only gift that I will have given to her in a long time, is it reasonable or unreasonable to put even hundreds of times greater importance to this than

¹¹ The loss to others from Anna's buying the ticket is 98. The loss to Anna from donating to Oxfam is 50 $((1-0) \times 50)$. Since the loss to others is greater than M times loss to Anna, the latter is not permitted not to donate money to Oxfam.

to giving the same amount of money to Oxfam? Our intuitions diverge here. It is debatable whether labeling the weighting of our personal projects as reasonable or unreasonable can be done without a help of an independent criterion of reasonableness. Since Scheffler provides none, moral permissibility of actions, according to the hybrid view, remains determinable by, among other things, subjective considerations.

The hybrid view requires some sacrifices but avoids unreasonable demands of simple act-consequentialism. However, besides the problems stated above, the hybrid view also suffers from the problem of permitting harm to others for the sake of achieving personal projects (see Kagan 1984: 251). Scheffler specifies the agent-centered prerogative as sensitive only to the size of the loss to others, but not to whether the loss is caused by the agent's act. The agent-centered prerogative will then not only permit agents to allow harm to others, it will also permit them to do harm in the pursuit of their non-optimal projects.

For example, the hybrid view does not distinguish between the moral permissibility of acts in the following two cases. In the first case, an agent kills her rich cousin in order to inherit \$10,000 and spend it on her personal projects. In the second case, an agent, who already has \$10,000, decides to spend his money on personal projects rather than to donate it to Oxfam. Suppose that the weight that the agent gives to her personal projects in the second case is large enough so that it grants her the desired agent-centered prerogative. However, the weight that the agent can give to his personal projects can be the same in the first case too. The hybrid view then implies that the agent is morally permitted to kill her rich cousin.¹²

It might be replied that this is not a new objection to act-consequentialism. After all, this theory is outrageously permissive, allowing us to kill an innocent person in order to save a few others because this is the best outcome overall. However, while simple act-

¹² One might object that the hybrid view contains an independent constraint against killing. However, even then we could imagine equivalent amount of harm that an agent can inflict on others in order to achieve her personal prerogatives that would be found permissible by the hybrid view.

SUB-MAXIMIZING VIEWS

consequentialists can defend themselves by saying that their theory allows for killing only if this would bring the best possible outcome, Scheffler is not allowed to make the same move. The hybrid view implies that it is morally permissible to kill someone even if that would bring the sub-optimal outcome, if by doing so, the agent would satisfy her personal projects. This is a clearly counterintuitive result that shows deep a inner problem in the hybrid view.

Chapter 4: Collective Act-Consequentialism

In comparison with the previously discussed simple act-consequentialism as well as sub-maximizing views, Liam Murphy's collective principle of beneficence, on the face of it, seems to deal with the demandingness objection successfully (see Murphy 2000). Murphy draws our attention to the obvious fact that the demands of simple act-consequentialism, when applied to redistribution of resources, become unreasonably high due to the non-full compliance with the principle (Murphy 2000: 76). That is why, Murphy believes, our moral obligations towards the world's poor should not be bigger than the ones we would have if every affluent person fulfilled his or her obligation. In this chapter, I will examine Murphy's view on our obligations to the world's poor, known as collective act-consequentialism, which in recent years gain on popularity. I will demonstrate that while it gains plausibility in one aspect, it loses out in another one.

4.1 Murphy's Collective Principle of Beneficence

In creating a principle of morality the main stress is usually placed on the recipients of the beneficence. This is most obviously the fact with simple act-consequentialism. Act-consequentialists require people to stop investing time and money on their personal projects, forbid them from buying toys for their kids, spending money on hobbies, and expect them to go as far as cutting off their own limbs - all for alleviating third party's sufferings. However, the sacrifice would be not so high if everyone would do their share.

Even if the principle of morality on the normative level is sound, it often has less than desirable implications on the level of practical application. On a practical level, if one would

be told that, unless she abandons all her life projects and pleasurable activities in favor of giving all the time and resources to help the poor, she is living an immoral life, what are the chances that she would still do something disregarding the stigma of her own amorality? The aim of Murphy's account is to reconcile normative standards and practical implications of the principle of morality, more precisely, to maximize the overall good on a practical level by decreasing the demands on the normative one.

Murphy examines act-consequentialists' argument as applied to the sphere of famine relief through the notion of fairness, intuitively understood as an equal share. Simple act-consequentialism requires us to do the best we can independently of what others do. This makes the aim of famine relief to be an individual project of each and every person, which is clearly more than one person can surmount. Murphy's solution hence is to make beneficence a cooperative venture of all people and to ground it on the notion of fairness. He argues that the principle of beneficence is agent-neutral principle of morality, meaning that each of us has one common moral aim and that is the aim of acting as if being a member of one cooperative enterprise (Murphy 2000: 75; see also Tanyi 2011: 16).

Murphy argues that beneficence understood as a common project shows non-compliance of one's share in another light. It makes it to become unfair and morally significant. In the world of A, B, and C, if A and B do not comply with their share of moral obligations, simple act-consequentialism would require C to fulfill their share, which is, according to Murphy, unfair. He, thus, comes up with his own principle of beneficence based on common aim and his notion fairness.

As act-consequentialism does, Murphy's collective principle of beneficence also requires a person to maximize the good, but with two constrains combined together: (1) maximizations of the good should not result in one's life becoming worse off under the conditions of (2) full compliance with the requirements of the principle. In other words, under

full compliance with the principle, we expect the aggregated well-being to be increased until the optimum level. However, in order to achieve this optimum level of aggregated well-being, people are not required to sacrifice more than they would have to sacrifice under the full compliance. Partial compliance with this principle does not require us to sacrifice more so as to optimize the aggregated well-being but rather to make a sacrifice equal to the one which is necessary for bringing the expected aggregated well-being under the full compliance. A person is welcome to make a bigger sacrifice but she is not morally required to do so (Murphy 2000: 117-8; see also Mulgan 2001: 105-6). Thus, Murphy's collective principle of beneficence allows for supererogatory acts. That is to say, a person may do more than it is morally required from her, which would be praiseworthy, but not morally wrong not to do (Murphy 2000: 121).

Moreover, the argument contains what Murphy calls a third person rider. While making a decision among available options, an agent ought not to choose such option, among the consequences of which would be the decrease of another person's expected well-being under full compliance. The moral requirement to make a sacrifice disappears if the loss caused by it would bring a person to the level of well-being which is lower than if it were under full compliance. In fact, the sacrifice is not morally required at all if a person's well-being is already lower than her expected well-being under full compliance.

A fair question to ask now is how would this collective principle operate? In order to answer this question and illustrate the principle let us appeal to a small two-steps thought experiment.¹³ The first step is to determine the ideal situation. By doing so we help an agent to find a criterion for practical standards necessary for guidance in the decision-making process. This ideal situation presupposes that everyone who is able to comply does so and thus acts optimally. Suppose that on the planet X there are 100 affluent people, 20 poor ones

¹³ The method is mentioned in Mulgan 2001: 106. I incorporated it into a thought experiment.

and 10 extremely poor ones. The ideal situation presupposes that all 100 affluent people make sacrifices in order to help the 10 extremely poor people, thus aggregating the expected well-being. Being very affluent, these 100 would not make their own level of expected well-being worse. Being very poor, the 20 people are not required to make sacrifice because it would push them below their level of expected well-being. Suppose now that in order for 100 affluent people to help 10 extremely poor people, the former should donate 5% of their income monthly for a period of one year.

The second step is to estimate your own share of sacrifice as if you are in the ideal situation that is to say, if everyone complies with the principle. This means that you, being one of 100 affluent citizens on the planet X, are morally required to donate 5% of your income monthly for a period of one year independently of whether all 100 are complying with the principle. Even if only 50 affluent people are complying with the principle you are not required to give more.

Plausible as it seems, Murphy's collective principle of beneficence has certain problematic points. One of them is demonstrated by the following example brought up by Murphy himself (Murphy 2000: 127-33). Suppose you and another stranger are walking on a bridge over a shallow pond where two children are drowning. The stranger is standing on the bridge but doing nothing.¹⁴ Being a compliant citizen, you estimate that in idealized situation your share of moral duty is to save one child and everything above it, is supererogatory. Murphy says that you may save the second child but you are not required to do so.

Even though Murphy understands the counterintuitiveness of his own conclusion, he explains it in the following way. Firstly, according to the collective principle of beneficence an agent is not required to save the second child but she is nevertheless required to follow the rule of thumb, which, in the case when saving the second child does not impose extreme costs

¹⁴ According to Murphy's thought experiment, the stranger passes by. I considered it necessary to change it slightly. If a stranger passes by and you find yourself all alone on the bridge, one might have said that Murphy's collective principle would require you to save both children.

on her, will tell her to save the child. Secondly, the fact that an agent fails to save the second child would display her as someone who has a bad character. That is why, Murphy concludes, it seems wrong not to save the second child while according to his principle it is morally permissible not to do so. Finally, he suggests that the cases of rescuing might be regarded as special and treated accordingly through a special moral principle devoted to it (Murphy 2000: 127).

Murphy's defense of his collective principle in the face of two drowning children objection is infinitesimally weak. It is simply implausible to say that, by failing to save the second child in the situation where the stranger left the bridge without doing her share, the agent is violating the rule of thumb or demonstrating her bad character. Few people would disagree that in this example an agent in fact fails to do something which she is morally required to do (Streumer 2004: 357-62; see also Philips 2007: chapter 5).

The question which should be asked now is analogous to the one that Peter Singer once asked.¹⁵ Is it possible that fairness leads us as far as not saving the second child because the stranger failed to do her share, even if we know that the child will die if we do not do the stranger's share? I believe that the answer might lie in the correct understanding of the notion of fairness. Our intuitive understanding of fairness when applied to the question of moral obligations to the world's poor is the one which require us to have equal share of obligations. However, when the whole conception of morality relies on our understanding of fairness and, as a result, has such implications as not saving the child from being drowned, it seems appropriate to include some other more substantive considerations into our account.

Suppose the following scenario. A and B are two children. They are sick and in need of medicine. The state of health of A is graver the B's. She needs a double dose of medicine of what B does. In case she has the treatment she will live 60 more years. B needs only one

¹⁵ "While the idea that no one need do more than his or her fair share is a powerful one, should it prevail if we know that others are not doing their fair share and that children will die preventable deaths unless we do more than our fair share?" (Singer 1999b: 63)

dose to be completely cured and live a long life of 60 more years. Meanwhile, if B does not have any treatment, she will not die but live 5 years less than otherwise. Suppose also that you, a benefactor, have only two doses of the medicine they need to distribute. Fairness, understood as equal share, would require us to give one dose to A and another one to B. The result of such distribution would be that A dies and B is cured and lives a long life. This result seems dramatic. Few people would agree that it is fair to let A die in this case. Fairness requires that we give both doses of medicine to A¹⁶.

Defenders of Murphy could say at this point that the collective principle of beneficence is in general not problematic but only limited to the cases where the resources are divisible and not lumpy, as it is in the rescue cases. What distinguishes human lives from material resources is that lives can be either saved or not saved while material resources allow for degrees. Under this assumption, defenders of Murphy can plausibly conclude that the collective principle of beneficence works for the exception of rescue cases. To be sure, this would be a significant complication for Murphy's view since we often expect that a principle of morality functions equally well in any case where the question of what is a morally required thing to do appears. As Murphy himself notes, there are no morally relevant factors that distinguish rescue cases from other cases of saving human lives so we should not treat the two differently (Murphy 2000: 130). However, let us assume that the collective principle of beneficence does not apply to the rescue cases. Is it any more plausible in other cases where beneficence towards the people in need is in question? In order to see that it is not, consider two other objections to it presented by Mulgan (see Mulgan 2001: 107-10; 118-20).

The first objection, the so-called 'wrong facts objection', roughly says that direct and stringent dependence on empirical facts such as quantity of the poor and affluent people does not serve as a sufficient justification in the face of sacrifice which these numbers require us to

¹⁶This example I owe to Anton Markoc.

make. Mulgan illustrates this objection with the following case. Suppose there are 50 million of poor people in the world. Its ratio with the amount of affluent people triggers moral obligations equal to 10 percents of each affluent's income, which seems reasonable. However, if we change the number of the poor people from 50 million to 1 million, the amount of sacrifice decreases accordingly from 10 percents to one-fifth of 1 percent of the affluent's income. Finally, if we change the amount of the poor people to 2500 million, the collective principle would require from the affluent to sacrifice everything.

By the same token, these fluctuations can also be demonstrated if we change the amount of the affluent people in the world. Suppose that in the ideal situation there are very few affluent people. In such a situation, even the condition of full compliance would not save an affluent person from unreasonable demands. Contrary to this, in the situation with very many affluent people, why would an affluent be morally required to sacrifice so little when she could do much more, especially knowing that there are still many people suffering from other people's non-compliance? The claim Mulgan is trying to make is not that the empirical data is irrelevant in the face of our moral obligations to the world's poor but that it should not constitute such direct and stringent dependence (Mulgan 2001: 107-9).

One can respond that this kind of dependence is necessary given the fact that duties related to distributive shares inherently depend on facts about how other persons fare. So even if, by imagining possible circumstances, collective principle of beneficence comes to be overdemanding or underdemanding in these circumstances, this would not be an objection to it. This argument is, however, too quick. It is one thing to say that moral duties generating distributive shares must depend on empirical facts about persons bearing those shares but we cannot conclude that, because of this inherent dependence, no possible situation in which those moral duties become counterintuitive serves against the principle proposing those duties. When we test the soundness of a certain moral principle, it is enough to consider possible

situations with empirical assumptions relevantly similar to the ones found in the actual world and see the implications of the given principle in those situations. If the implications are counterintuitive, we have enough ground to assume that something is wrong with the principle. This is exactly what happens with Murphy's principle of beneficence when applied to possible worlds with many affluents and few destitutes or vice versa. In the former, it generates unreasonable demands while in the latter, it proposes the demands which we intuitively consider to be too small (given the fact that there are still many people in need and you actually know that there are many affluents who do not comply with their moral duties).

Another objection Mulgan raises against Murphy's collective principle of beneficence consists in the undesirable outcome which may be produced by the first step of this principle, that is, by determining the ideal situation, together with the causes of global poverty. By complying with the collective principle, according to Murphy, people act optimally. In order to estimate the amount of sacrifice, an agent should resolve her ignorance by determining the ideal situation. Mulgan claims that, while resolving the ignorance, an agent cannot avoid paying attention to the causes of the poverty which the collective principle is sensitive to.

To illustrate the objection imagine the following. Anna is a complying affluent. She has a well paid job and she regularly makes her donations to Oxfam, preliminarily calculated with the help of Murphy's collective principle of beneficence. The economic crisis comes and she foresees the need to cut her expenses in the future. She decides that in order to be able to donate to Oxfam in future she would better start saving up for her little personal poverty relief fund. Mulgan claims that in attempts to apply Murphy's collective principle, Anna would start idealizing situations of future famine and, if this future famine is caused by people's failure to act optimally, then this famine would not exist in the future at all. Hence Anna would not have any moral obligation to save up for future starving people.

COLLECTIVE ACT-CONSEQUENTIALISM

In this chapter, I examined Murphy's collective consequentialism. This theory incorporates a great amount of our intuitive understanding about what is morally obligatory to do in order to help the people in need. It allows for supererogatory acts by suggesting that each person should sacrifice no more than if she or he would sacrifice if all the people acted as maximizers of overall well-being. Nevertheless, I rejected Murphy's view due to its counterintuitive implications, showing that, in some cases, it is no less demanding than the extremist position, while in the others, it is less demanding than it should be. I conclude that act-consequentialists cannot choose Murphy's way in order to get out from the counterintuitive implications of their theory.

Conclusion

That a great number of people live in extreme poverty while others enjoy the pleasures of affluence is a fact that hardly anyone can question today. Like in many other issues in applied ethics and political theory, empirical facts can raise questions of serious theoretical concern. As members of affluent countries, the reasonable among us know that they owe something to those living in dire straits around the globe. But how much they owe is an ambiguous question. Act-consequentialism is a moral theory that gives a straightforward answer. In general, act-consequentialists argue that we should bring about as much overall good in the world as possible even if, by doing so, we would make our own lives worse than they used to be. On the face of it, the theory seems to imply unreasonable moral demands. Is this a reason enough to reject it? This was the main issue of my thesis and I argued for an affirmative answer.

I have examined three groups of act-consequentialist arguments. The first group consists of simple act-consequentialists or extremists. The most famous extremist argument was offered and is still defended by Peter Singer. Singer argues that no individual with a decent sense of morality would question that we have a moral obligation to prevent something bad from happening if, by doing this thing, we would not sacrifice something of similar moral importance. For example, we all agree that you have a moral obligation to save the drowning child, even if you would thereby ruin your new pair of shoes. Singer then argues that each of us, members of affluent countries, can do a similar sacrifice, by, for example, donating to renowned aid agencies and save the people who starve in the third-world countries. We ought to help the starving people all the time, even if our lives become ascetic and even if no one else is doing the same.

CONCLUSION

Singer's argument strikes us as mistaken because we believe that morality cannot demand such extreme sacrifices. In simple terms, we believe it to be counterintuitive. The main extremist argument is an indirect one. Extremists want to show that relying on our intuitions in morality is unacceptable and that, hence, any accusations of extremists' conclusions concerning our moral duties towards the world's poor as overdemanding are unfounded. I discussed two arguments of the kind: Greene's and Singer's argument that moral intuitions are unreliable because they stem from parts of the brain responsible for our emotions (allegedly confirmed via fMRI research) and Unger's argument for the unreliability of moral intuitions because of their mutually conflicting nature. I argued that both of these strategies are unconvincing. While the argument from fMRI research rests on unargued normative claims about the unimportance of the factor of distance in morality, Unger's method of several options contains erroneous methodique of projective grouping and thus mistakenly considers mortal harm inflicted on an innocent person in order to save several others a morally irrelevant factor.

The second group of act-consequentialist arguments represents the sub-maximizing views of Slote and Scheffler. These authors admit that the implications of extremism concerning our moral obligations to those in need are too demanding and they try to incorporate this intuitive understanding of moral demands within a broadly consequentialist framework. Sub-maximizing views suggest that, in general, we ought to bring about as much overall good as possible but that we are also sometimes allowed not to follow this requirement for the sake of achieving certain personal prerogatives. There are several problems with Slote's satisficing consequentialism and Scheffler's hybrid view. Being ambiguous between two interpretations, Slote's account implies either no less unreasonable demands than extremism or demands which we intuitively consider to be less than required. Scheffler's view is, on the other hand, more precise about the amount of sacrifices that an

CONCLUSION

individual can evade for the sake of achieving personal interests. The hybrid view, however turns out to be not more plausible than the Slote's satisficing consequentialism. Allowing an agent to attribute more weight to her personal interests than to the interests of distant strangers, Schefflers commits a mistake by making *M* purely subjective and constant. The implications of such a mistake make the hybrid view unreasonably demanding in some cases and repulsive in the other ones. Sub-maximizing strategies, then, fare no better than the extremist ones.

Finally, the third group are collective consequentialist arguments. I examined the most renowned argument of the kind, Murphy's collective principle of beneficence. Collective consequentialism is a sub-maximizing view in a sense that it permits for the satisfaction of personal prerogatives, although only indirectly. What makes it recognizable is the idea that an agent should not strive to bring about more overall good than to the extent that she would be morally required to do if everyone acted the same, that is to say, if everyone followed the principles of simple act-consequentialism. This makes Murphy's view intuitively appealing and seemingly unsusceptible to the demandingness objection. However, as it was demonstrated, the intuitive appeal of the collective principle of beneficence is significantly undermined by the rescue cases to which, when applied, the principle generates repugnant results. In other, non-rescue cases, I argued that it is simply not the case that Murphy's collective consequentialism avoids the demandingness objection. What is more, in certain circumstances, Murphy's view implies that someone can go with less moral burden than that is intuitively considered as required. It, thus, cannot serve as a safe-ground for act-consequentialists efforts to avoid the overdemanding implications of their theory.

The demandingness objection is definitely not the only objection that one can invoke to challenge Singer's or any other act-consequentialist argument. However, in discussions about global poverty, demands of some moral theory are more than important issue to argue about

CONCLUSION

given the enormous practical implications that can be derived from different conclusions one can reach at. In this thesis, I argued that the most important act-consequentialist theories cannot avoid the demandingness objection and that, some of them, have other unacceptable implications.

My argument surely neither establishes the truth of common-sense morality nor it can be used as a support for the popular but certainly wrong ideas that we owe nothing to those starving in the third-world countries either because they are distant from us and we should give priority to our conationals or because we are not responsible for the horrendous state in which they live. One can only hope that some other well-argued moral theory, consequentialist or otherwise, will one day come to nicely fit into our intuitive picture of what we should do regarding global poverty. Until then, at least on the theoretical level, we can abandon act-consequentialist arguments for our obligations towards the world's poor as unsound. At the level of practice, progress in famine relief, however small, will be definite if only we comply with the ordinary intuitions telling us how to help the ones in grave needs.

Reference list

- Appiah, Kwame Anthony. *Experiments in Ethics*. Cambridge, MA: Harvard University Press.
- Berker, Selim. 2009. The Normative Insignificance of Neuroscience. *Philosophy and Public Affairs*, 37 (4): 330–45.
- Chapell, Timothy (ed.). 2009. *The Problem of Moral Demandingness: New Philosophical Essays*. Basingstoke: Palgrave.
- Foot, Philippa. 1967. The Problem of Abortion and the Doctrine of Double Effect. *Oxford Review*, (5): 5–15.
- Greene, Joshua et al. 2001. An fMRI Investigation of Emotional Engagement in Moral Judgment. *Science*, 293 (5537): 2105–8.
- Greene, Joshua. 2003. From Neural ‘Is’ to Moral ‘Ought’: What are the Moral Implications of Neuroscientific Moral Psychology? *Nature Reviews Neuroscience*, 4 (1): 847–50.
- Greene, Joshua. 2008. The Secret Joke of Kant’s Soul. In Walter Sinnott-Armstrong (ed.), *Moral Psychology*, Volume 3. Cambridge, MA: MIT Press.
- Greene, Joshua. 2009. Dual-Process Morality and the Personal/Impersonal Distinction: A Reply to McGuire, Langdon, Coltheart, and Mackenzie. *Journal of Experimental Social Psychology*, 45 (3): 581–4.
- Greene, Joshua. 2010. Notes on “The Normative Insignificance of Neuroscience” by Selim Berker. Ms. (presentation at Arizona State University). 34pp.
- Gruen, Lori. 1999. ‘Must Utilitarians Be Impartial?’ In Dale Jamieson (ed.), *Singer and His Critics*. Oxford: Blackwell Publishers.
- Hooker, Brad. 2009. The Demandingness Objection. In Timothy Chappell (ed.), *The Problem of Moral Demandingness: New Philosophical Essays*. Basingstoke: Palgrave.
- Kagan, Shelly. 1984. Does Consequentialism Demand too Much? Recent Work on the Limits of Obligation. *Philosophy and Public Affairs*, 13 (3): 239–54.
- Kagan, Shelly. 1989. *The Limits of Morality*. Oxford: Clarendon Press.
- Kamm, Frances M. 1998. Grouping and Imposition of Loss. *Utilitas*, 10 (3): 292–319.
- Kamm, Frances M. 1999. ‘Feminine Ethics: the Problem of the Distance in Morality and Singer’s Ethical Theory’. In Dale Jamieson (ed.), *Singer and His Critics*. Oxford: Blackwell Publishers.
- Kamm, Frances M. 2007. *Intricate Ethics: Rights, Responsibilities and Permissible Harm*. Oxford: Oxford University Press.
- Kamm, Frances M. 2009. Neuroscience and Moral Reasoning: A Note on Recent Research. *Philosophy and Public Affairs*, 37 (4): 330–45.
- McGinn, Colin. 1999. ‘Our Duties to Animals and the Poor’. In Dale Jamieson (ed.), *Singer and His Critics*. Oxford: Blackwell Publishers.
- Mulgan, Tim. 2000. Peter Unger, Living High and Letting Die (Review). *Mind*, 109 (434): 397–400.

- Mulgan, Tim. 2001. *The Demands of Consequentialism*. Oxford: Oxford University Press.
- Murphy, B. Liam. 1999. 'Institutions and the Demands of Justice'. *Philosophy and Public Affairs*, 27 (4): 251–91.
- Murphy, B. Liam. 2000. *Moral Demands in Nonideal Theory*. New York: Oxford University Press.
- Nagel, Thomas. 2010. What Peter Singer Wants of You. *The New York Review of Books*, 27.
- Pettit, Phillip. 1984. Satisficing Consequentialism. *Proceedings of the Aristotelian Society*, Supplementary Volumes, 58 (1): 165–76.
- Portmore, Douglas W. 2011. Agent-Neutral and Agent-Relative. In James Crimmins and Douglas Long (eds.), *Encyclopedia of Utilitarianism*, Bristol: Thoemmes/Continuum.
- Rawls, John. 1971. *A Theory of Justice*. Cambridge, MA: Harvard University Press.
- Scheffler, Samuel. 1992. Prerogatives without Restrictions. *Philosophical Perspectives*, 6 (1): 377-97.
- Scheffler, Samuel. 1994. *The Rejection of Consequentialism: A Philosophical Investigation of the Considerations Underlying Rival Moral Conceptions*. Revised edition. Oxford: Oxford University Press.
- Singer, Peter. 1972. Famine, Affluence, and Morality. *Philosophy and Public Affairs*, 1 (1): 229- 43.
- Singer, Peter. 1993. *Practical Ethics*. 2nd edition. Cambridge: Cambridge University Press.
- Singer, Peter. 1999a. 'A Response'. In Dale Jamieson (ed.), *Singer and His Critics*. Oxford: Blackwell Publishers.
- Singer, Peter 1999b. The Singer Solution to World Poverty, *New York Times Sunday Magazine*, available at url: <http://www.nytimes.com/library/magazine/home/19990905mag-poverty-singer.html>
- Singer, Peter. 2005. Ethics and Intuitions. *Journal of Ethics*, 9 (1): 331-52.
- Singer, Peter. 2009. *The Life You Can Save: Acting Now to End World Poverty*. New York: Random House.
- Sinnott-Armstrong, Walter, Young, Liane and Cushman, Fiery. 2010. Moral Intuitions. In John M. Doris and the Moral Psychology Research Group (eds.), *The Moral Psychology Handbook*. Oxford: Oxford University Press.
- Slote, Michael. 1984. Satisficing Consequentialism. *Proceedings of the Aristotelian Society*, Supplementary Volumes, 58 (1): 139–63.
- Slote, Michael. 1991. *Beyond Optimizing: A Study of Rational Choice*. Cambridge, MA: Harvard University Press.
- Streumer, Bart. 2004. Liam B. Murphy, *Moral Demands in Non-Ideal Theory* (Review). *Ratio*, 17 (1): 357–62.
- Tanyi, Attila. 2011. The Case for Authority. In Sebastian Schleidgen (ed.), *Ethics and Morals*. Vol. 1. Marburg: Tectum.
- Thomson, Judith Jarvis. 1985. The Trolley Problem. *Yale Law Journal*, 94 (6): 1395–415.
- UN Report on the World Social Situation 2010. 2009. *Rethinking Poverty*. New York: United Nations. Available at url: <http://www.un.org/esa/socdev/rwss/docs/2010/fullreport.pdf>
- Unger, Peter. 1996. *Living High and Letting Die: Our Illusion of Innocence*. Oxford: Oxford University Press.

Williams, Bernard. 1973. *A Critique of Utilitarianism*. In J.J.C. Smart and Bernard Williams, *Utilitarianism: For and Against*, Cambridge: Cambridge University Press.