

University of Nebraska - Lincoln

DigitalCommons@University of Nebraska - Lincoln

Philosophy Dissertations, Theses, & Student
Research

Philosophy, Department of

2011

Strategies for Defusing the Demandingness Objection

Justin J. Moss

University of Nebraska-Lincoln, moss@unlserve.unl.edu

Follow this and additional works at: <http://digitalcommons.unl.edu/philosophydiss>



Part of the [Ethics and Political Philosophy Commons](#)

Moss, Justin J., "Strategies for Defusing the Demandingness Objection" (2011). *Philosophy Dissertations, Theses, & Student Research*. 6.
<http://digitalcommons.unl.edu/philosophydiss/6>

This Article is brought to you for free and open access by the Philosophy, Department of at DigitalCommons@University of Nebraska - Lincoln. It has been accepted for inclusion in Philosophy Dissertations, Theses, & Student Research by an authorized administrator of DigitalCommons@University of Nebraska - Lincoln.

STRATEGIES FOR DEFUSING THE DEMANDINGNESS OBJECTION

by

Justin James Moss

A DISSERTATION

Presented to the Faculty of

The Graduate College at the University of Nebraska

In Partial Fulfillment of Requirements

For the Degree of Doctor of Philosophy

Major: Philosophy

Under the Supervision of Professor David Sobel

Lincoln, Nebraska

December, 2011

STRATEGIES FOR DEFUSING THE DEMANDINGNESS OBJECTION

Justin James Moss, Ph.D.

University of Nebraska, 2011

Adviser: David Sobel

Bernard Williams's formulation of the Demandingness Objection holds that living a moral life, as the consequentialist understands it, is incompatible with living a life that is good for human beings. This is because the demands of consequentialist morality threaten to overwhelm the life of the person who cares about being moral, thus leaving no time for their own projects and interests. Several prominent consequentialists have responded to the Demandingness Objection by seeking a more moderate and indirect form of consequentialism that does not require as strong a duty of beneficence as classical utilitarianism. I review and criticize three prominent moderate forms of consequentialism: Brad Hooker's rule consequentialism; the theories of Samuel Scheffler and Tim Mulgan, which share an agent centered prerogative; and Liam Murphy's collective principle of beneficence. As the primary method of criticism, I develop a type of collective action problem, which I refer to as the Polluter's Dilemma. This dilemma occurs when a moral theory permits agents to favor their own interests and in doing so create a very small harm that affects all other agents. These small harms accumulate, and the result is that the long-term interests of all agents are greatly harmed. I provide reasons to think that acceptable forms of consequentialism must avoid the Polluter's Dilemma, and I argue that the three mentioned forms of moderate consequentialism do not avoid the Polluter's Dilemma. In concluding, I review a form of consequentialism that, I argue,

avoids both the Polluter's Dilemma and the Demandingness Objection. Based on this result, I make recommendations about how future consequentialist moral theories should develop. Consequentialists should seek a moral theory that leaves agents room for their own projects, but that theory should be flexible enough to recognize which stringent demands are appropriate and which stringent demands are not, and the theory should not support the aims of agents that leave everyone worse off in the long term.

To my parents

Gary and Nancy Moss

who always demanded excellence from me.

ACKNOWLEDGEMENTS

I have been the recipient of a great deal of support from numerous sources in the writing of this dissertation. Above all, my thanks must go to my adviser, David Sobel, for his extensive support and advice. He guided me through every stage of the dissertation-writing process, offering valuable feedback on every draft, and helped me make this work better by making important suggestions for revision. I am very grateful to him for his guidance, mentorship, and friendship.

My other committee members, Joseph Mendola and Mark van Roojen, assisted me by providing valuable feedback on my dissertation and many suggestions for how to make improvements. Quite apart from their duties as my readers, their influence, leadership, and support of all graduate students in the UNL Philosophy Department has helped to make the positive intellectual and professional climate of the department possible.

This dissertation has benefitted from the input and influence of many others in the Philosophy Department, faculty and graduate students alike. I especially wish to thank Aaron Bronfman, who pushed me to think harder about a crucial point in the central argument of the dissertation; David Henderson, whose friendly support kept me motivated and on track; and the graduate students of the UNL Philosophy Department, who attended three sessions of the Graduate Student Colloquia series in which I presented pieces of this dissertation. For their feedback, encouragement, and questioning, I wish to thank Patrick Arnold, David Chavez, Matthew Dee, Allison Fritz, Cullen Gatten, Shane George, Christopher Gibilisco, Landon Hedrick, Cliff Hill, Andrew Jaeger,

Clare LaFrance, Timothy Loughlin, Christopher Richards, Steven Swartzler, Adam Thompson, and Preston Werner.

My wife, Vanessa Steinroetter, supported me and encouraged me throughout the process of my graduate education and especially through the intense months of drafting, revising, and completing this dissertation. I am deeply grateful to her for all her support, and could never have done it without her.

Table of Contents

| | |
|--|------------|
| Introduction..... | 1 |
| Chapter 1: Permissiveness and the Agent-Centered Prerogative..... | 14 |
| 1. Introduction | 14 |
| 2. Scheffler’s Original Agent-Centered Prerogative | 15 |
| 3. An Improved Prerogative: Mulgan’s Combined Consequentialism | 26 |
| 4. Polluter’s Dilemmas and the Demandingness of Combined Consequentialism | 30 |
| 5. Why the Agent-Centered Prerogative is Generally Unacceptable as a Component of a Consequentialist Moral Theory | 39 |
| Chapter 2: Rule Consequentialism’s Dilemma: Demandingness or Self-Defeat | 48 |
| 1. What is Rule Consequentialism?..... | 48 |
| 2. How Hooker Hopes to Avoid the Demandingness Objection..... | 53 |
| 3. Being Careful about Hypothetical Behavior and Belief..... | 57 |
| 4. Rule Consequentialism and Polluter’s Dilemmas | 66 |
| Chapter 3: The Inconsistency of the Collective Principle | 74 |
| 1. The Collective Principle and the Compliance Condition..... | 74 |
| 2. Does the Collective Principle Fairly Evaluate Rescue Cases?..... | 79 |
| 3. Mulgan’s Wrong Facts Objection and the Problems of Global Climate Change | 89 |
| 4. The Modified Compliance Condition | 97 |
| Chapter 4: Multiple-Act Consequentialism and Constraints on Acceptable Forms of Consequentialism..... | 104 |
| 1. Introduction | 104 |

| | |
|---|-----|
| 2. Consequentialism, Its Rivals, and Climate Change | 105 |
| 3. The Polluter's Dilemma as a Constraint on Acceptable Forms of Consequentialism | 108 |
| 4. Mendola's Multiple-Act Consequentialism | 118 |
| 5. Atomic Agency Within a Life? | 123 |
| 6. How Overlapping Group Agency Plausibly Solves the Polluter's Dilemma..... | 127 |
| References..... | 132 |

Introduction

Consequentialism – the view that an action’s moral permissibility depends solely on the rankings of states of affairs on some evaluative scale – is widely viewed as a candidate for the best general approach to doing ethics. Its most famous version – classical act utilitarianism, the view that the rightness of an action depends solely upon the propensity of that action to maximize happiness - has many attractive features. It is an impartial theory; it asks that we treat the interests of all relevant parties equally and impartially, all other things being equal. It is an inclusive theory; the moral community, or the class of beings whose interests we are morally obligated to take into account, includes not only all human beings within the bounds of the moral community, but all sentient beings. It is a flexible theory; the central goal of act utilitarianism is to produce the best result possible, but it may be paired with many different accounts of what this best possible result consists in. Finally, it is a reformist theory of morality which recommends that we change our ways, both at the collective and at the personal level, if a careful assessment of the facts leads us to conclude that things could be done better. Historically, it has served as the theoretical groundwork for many excellent causes. Women’s liberation and political liberalism (John Stuart Mill), reform of the legal and penal code (Jeremy Bentham), and animal liberation, environmentalism, and assistance to those living in extreme poverty (Peter Singer) are but a few of the famous and worthy causes that adherents of act utilitarianism have historically advocated. For all these reasons, and others, act utilitarianism – and consequentialism more generally, even though the specific version of

it being talked about may or may not share these features with act utilitarianism – has many proponents. Indeed, several contemporary philosophers refer to the following thought as the “Compelling Idea” that makes consequentialism such a perennially attractive approach: it is always permissible for an agent to bring about the best possible state of affairs.¹

What many see as beneficial, however, others see as problematic. Both critics of consequentialism and thoughtful consequentialists who are not utilitarians can and have taken issue with all the items that I have listed as attractive features of consequentialism.² I do not aim to focus on all of their criticisms in this dissertation. The focus of the present work is the fourth item on my list; namely, that act utilitarianism is a moral theory with reformist impulses. Although this feature of act utilitarianism lends itself to excellent uses, it also lends itself to what has become a traditional objection to act utilitarianism. Many philosophers have believed that act utilitarianism asks for far too much sacrifice from us to be a plausible ethical theory. If the main goal of morality is to produce as much good in the world as possible, then it appears that any action that is not somehow maximizing the amount of good in the world would be judged as wrong.

This implication could seem intuitive in some cases. For example, a critic of a particular government policy, such as the War on Drugs in the United States or the practice of capital punishment, might argue on consequentialist grounds that her target policy should be discontinued because it fails to maximize the good. But the traditional

¹ Douglas Portmore (2005, 98) puts the thought this way: “what about [consequentialism] is so compelling? Well, it seems to be the very simple and seductive idea that it can never be wrong to produce the best available state of affairs.” See Schroeder 2006, 2007 and Dreier 2011 for more uses of this terminology.

² There is, to take one example into consideration, a lively debate over whether consequentialism must be thought of as impartial and agent-neutral. What I think of as the standard view is that consequentialism must be agent-neutral (see McNaughton and Rawling 1991 and Pettit 1997). Portmore 2001 and Broome 1991, by contrast, argue that consequentialism can incorporate agent relativity.

form of consequentialism is not a theory of morality at the collective level only; it is also meant to be a theory of morality at the individual level. Our own actions, or lack thereof, also come in for judgment under this theory. If it turns out that our own actions are failing to maximize the good, then we will turn out to be much less morally good than we might have thought we were. If we, for example, have a choice between spending some amount of our own money on personal entertainment, or spending that same amount of money on an effective charity which would provide aid to someone living in dire poverty, it seems intuitive to suppose that the greater amount of happiness would be generated by giving this amount of money to charity. Though our own entertainment matters from the point of view of act utilitarianism, we are not suffering much if we lose out on, for example, a single trip to the movie theater. We cannot say the same of someone, to give two examples, whose supply of food and fresh water is insecure, or who lives without access to competent medical services. Our trip to the movies is the lesser good. We would not be giving up much of moral importance if we donated our money to charity.

That might be enough to persuade a considerate and thoughtful person to give up one such trip to the movie theater. But the same considerations apply for the next trip to the movie theater, or the next impulse purchase, or the way we spend our free time, or the careers we choose, and so forth. Indeed, the traditional objection to act utilitarianism to which I alluded earlier has come to be known as the Demandingness Objection.

This objection, as I shall understand it for the purposes of this dissertation, holds that any system worth calling a system of morality should be compatible with living a life that is an enjoyable life for a human being to lead. The way a moral system achieves this compatibility is to not overwhelm the life of the person who wants to be moral. The

friend of the Demandingness Objection thinks that a good life for a human being includes the freedom to choose interests and projects that matter to us and that make our lives meaningful and enjoyable. Naturally, any system of morality worthy of its name will sometimes impose demands on us. But the friend of the Demandingness Objection believes that there is a significant difference between the number and scope of the demands an acceptable moral theory may place on us and the number and scope of the demands that a moral theory may place on us. Because there is so much good to be done in the world, act utilitarianism requires us to scrutinize each of our actions and consider how we might best work to generate the greatest good for all. The friend of the Demandingness Objection believes that, in weighing what we ought to do, it will almost always be the case that a person ought to work for the greater good, which she may not identify with or care about, and forego working on his own projects and interests, which are things that give her life meaning and enjoyment. This will hold for every decision a person must make.

To drive the point home, Tim Mulgan describes a case in which a person who has already contributed a significant sum of money to charity is once again faced with the choice of buying expensive tickets to the theater or giving the money to an effective charity. Suppose that this person chooses to buy the tickets. As Mulgan says, “the Demandingness Objection says that Consequentialism must condemn Affluent’s behavior, and that this [the condemnation] is unreasonable.”³ Thus, it appears, act utilitarianism – and perhaps all theories which go under the heading of ‘consequentialism’ as well – is a highly demanding moral theory. It is demanding

³ Mulgan 2001, 4. The capitalization of “Consequentialism” is his, and Affluent is the name of the person in the thought experiment.

because it alienates agents from the projects and interests that make their lives meaningful and enjoyable.

Consequentialists and non-consequentialists alike feel the pull of this Objection. For non-consequentialists, one way of prodding the consequentialist on this matter is to flesh out an unintuitive implication of the Objection. Non-consequentialists will sometimes point out consequentialism does away with the possibility of supererogation. No longer is there a certain class of actions which are morally praiseworthy for an agent to perform but not morally required of that agent; no longer could we say that an agent goes above and beyond the call of duty in doing morally praiseworthy actions when that agent donates to charity, or so it seems. It would appear that something that we intuitively felt was true - that there really are such things as supererogatory actions - cannot be the case under consequentialism.⁴ Such an observation is what Peter Unger might label as a “Preservationist” response. According to Unger, the Preservationist believes that our reaction to particular cases accurately reflects our moral commitments, and this appearance of an accurate reflection is trustworthy and ought to be preserved.⁵ In this case, we commonly have the intuition that giving to charity is not morally required, and one ought not to be criticized if one gives nothing to charity. Moreover, philosophers and non-philosophers have the intuition that there are many kinds of actions which might produce the most good, but which nevertheless are not morally required. That is something the consequentialist may have to give up on, the non-consequentialist might claim.

What are some responses to the Demandingness Objection that utilitarians, and

⁴ This is a standard enough way of fleshing out the Demandingness Objection that it is commonly taught at the undergraduate level. See, for example, Shafer-Landau 2010, 132.

⁵ See Unger 1996, 11.

consequentialists more broadly, have offered?

One possible response is to deny that there is a problem for consequentialism here. Indeed, according to this response, we ought to think of this as a reason that consequentialism might actually be the correct moral theory. Morality is, in fact, very demanding. This is the response preferred by philosophers like Peter Singer, Shelly Kagan, and Peter Unger.⁶

Another sort of response is to think of such demandingness as a real problem for any form of consequentialism. Indeed, it is sometimes claimed that if there is anything wrong with consequentialism, it is its propensity to demand so much sacrifice.⁷ What form might a response from the consequentialist to this problem take?

One response of this type is to treat the phenomenon of demandingness as stemming from a misunderstanding of how consequentialists ought to go about deciding what to do. On this view, consequentialists distinguish between understanding consequentialism as an objective standard of rightness and truth-maker for moral claims, and consequentialism as a decision procedure that each agent ought to subjectively follow when wondering about what to do. It would be a mistake if agents accepted a commitment to regularly and directly evaluate acts in purely consequentialist terms. Instead, one should be committed to leading the sort of life and doing the kinds of actions which, if evaluated in consequentialist terms, would in fact make things go best for everyone concerned. Peter Railton has proposed just such a response to the Demandingness Objection.⁸ Railton argues that the fundamental notion that the

⁶ Singer 1972; Kagan 1989; Unger 1996.

⁷ Kagan 1989, xiii., describes such a claim.

⁸ Railton 1984, especially 152-153. To be more specific, Railton views the Demandingness Objection as involving a phenomenon that he dubs *alienation* and he characterizes as a separation between the rational

consequentialist accepts is that the maximal amount of human value should be promoted. The fundamental notion is *not* that the maximal amount of human value must be promoted *in a specific way*. Thus, a sincere consequentialist may behave in such a way that promotes the greatest good over the long run without being strictly committed to a difficult kind of calculating and alienating decision procedure. It may even be practically necessary for such a person to behave wrongly from an objective consequentialist point of view, such as when a person chooses to spend the day relaxing with his or her partner, rather than spending his or her day constructing homes for the homeless. However, the wrongdoing may be a blameless kind of wrongdoing. The blamelessness of this act lies in how the wrongdoing affects the agent. If the agent did not spend the day relaxing (on occasion) with his or her partner, then the agent might ultimately be a more cynical person. His or her capacity to bring about good in the world would be less-developed and his or her overall contribution to well-being would be less in the end because of his or her cynicism. So in order to ward off cynicism, consequentialism might recommend that the committed consequentialist not behave subjectively as a consequentialist.

Plausible as Railton's proposal may seem, it has not been universally adopted. Several philosophers who self-identify as consequentialists or who hold views sharing much in common with consequentialism treat the Demandingness Objection as indicative of a deeply-rooted problem with consequentialism that only much modification can solve. These philosophers think, with Railton, that morality ought to be compatible with allowing us time and resources to pursue a wide variety of relationship commitments and

self that decides what to do based on the available reasons for acting and the affective self that is more concerned with other-regarding relationships and an agent's own sentiments. In what follows, I shall focus less on the broader phenomenon that Railton describes, and instead limit my attention to the kind of alienation that is specifically about being estranged from one's projects and goals.

non-moral projects that might matter to us. The project of being a good person ought not to take over our lives (though it should certainly be an important project in our lives).⁹ But Railton and these philosophers part ways over the issue of whether consequentialism has the theoretical resources to preserve our projects, relationships, and personal integrity. Railton believes that consequentialism, properly understood, has such resources. The philosophers I allude to here – Samuel Scheffler, Tim Mulgan, Brad Hooker, and Liam Murphy – feel that consequentialism requires extensive modification in order to be a moral theory that can be plausibly thought to safeguard our personal integrity, and they have attempted to amend consequentialism so that it is less demanding, and therefore less vulnerable to the Objection.

Before I proceed to outlining the subsequent plan of the work, I should address a methodological concern. My use of the term “consequentialism” to describe the views of some of the philosophers mentioned above may strike some readers as problematic. The problem does not lie in identifying philosophers like Railton, Singer, and Unger as consequentialist. If anyone is a consequentialist, it is someone who believes that the rightness of action is solely determined by where an outcome ranks on some evaluative scale, and that outcomes must not be ranked an agent-relative (and perhaps also, time-relative) way. Philosophers like Railton satisfy this definition. But several of the other philosophers mentioned – such as Scheffler and Mulgan – use agent-relativity to rank states of affairs. This is a departure from the paradigm understanding of consequentialism. And the departure from the paradigm understanding of consequentialism may be problematic because of recent work arguing that any moral theory may be modeled using an agent-relative and time-relative ordering. I refer here to

⁹ Or so Bernard Williams argues in Williams 1973.

what is known in the literature, variously, as Dreier's Conjecture,¹⁰ or the Deontic Equivalence Thesis.¹¹ This thesis holds that for any non-consequentialist theory, a consequentialist equivalent may be devised. Because any theory may be 'consequentialized', some philosophers have worried that consequentialism as such may have no substantive content. Thus, if the term 'consequentialism' is to have any meaning at all, it will be necessary to draw more carefully the distinction between consequentialism and non-consequentialism. More to the point, theories that give different aims to different agents, such as Scheffler's and Mulgan's theories, may count as forms of non-consequentialism.¹² In addition, theories which may sometimes have us do actions that will not maximize (or satisfy) the good may also not count as consequentialism on some views. So why count these theories as forms of consequentialism in this dissertation?¹³

I count them as such for a couple of reasons. First, the purpose of my project is to examine modifications of consequentialism that are motivated by the desire to avoid the demandingness objection. For my purposes, nothing much shall hang on the specific issue of whether something is, when all is said and done, rightfully called a form of consequentialism. I will, of course, be concerned to draw lessons for consequentialism from examining the successes and failings of these theories. But those lessons that I draw shall not ultimately be affected by whether, for example, Scheffler's view is a form of consequentialism or not.

Second, there is some precedent in the literature for taking a broad view of what

¹⁰ See Dreier 1993 for the original conjecture, and Dreier 2011 for further discussion, plus Portmore 2011, 87 for other uses of the term.

¹¹ This is Douglas Portmore's term for Dreier's Conjecture (Portmore 2011).

¹² Campbell Brown has recently argued that agent-relative theories violate a necessary component of consequentialism; see Brown 2011 for this argument.

¹³ Thanks to Mark van Roojen for pressing this point.

counts as consequentialism. The one criterion that all sides in this debate seem to agree on is that if a theory does not make moral rightness or wrongness solely a function of the value of outcomes, then that theory is not consequentialist. Beyond this, there is little agreement. It is not clear that a focus on outcomes alone is sufficient to call a theory consequentialist. Some philosophers think that this is so, and argue that consequentialism can incorporate agent-relativity.¹⁴ That position is, of course, controversial; many philosophers do not accept the thesis that agent-relativity is compatible with consequentialism.¹⁵ And some philosophers take the view that ‘consequentialism’ is a term of art, and that there is no correct way to define the term (beyond the criterion I mention above), as James Dreier does.¹⁶ Although this is an important controversy, it is nevertheless a controversy that I see as beside the point of the present project. As mentioned before, what I am really interested in pursuing in this dissertation is a deeper understanding of the shortcomings of attempts to make consequentialism less demanding, and what that might mean for future consequentialist theorizing. Since that is my aim, and because I desire to respect the terminology that other authors use to describe themselves (particularly in Mulgan’s and Hooker’s cases), I shall, for simplicity’s sake, refer to them as forms of consequentialism.

In the subsequent chapters of this dissertation, I shall examine several of these amended, less demanding forms of consequentialism. I believe that the amended versions of consequentialism that I shall examine shall either fail to avoid the Demandingness

¹⁴ For example, see Broome 1991, Skorupski 1995, and Portmore 2001. Tim Mulgan also argues strenuously that consequentialists must give up on the idea that the right act on any given occasion is the act that produces the best consequences, because all the options available for an act consequentialist concerned to avoid the Demandingness Objection are unpalatable. For that claim, see Mulgan 2001, 25-49, esp. 49.

¹⁵ See McNaughton and Rawling 1991, Pettit 1997, and Brown 2011 for examples of philosophers who reject agent-relativity as compatible with consequentialism.

¹⁶ Dreier 2011, 97.

Objection, or - because of the way in which they are less demanding forms of consequentialism - they allow the creation of large-scale collective action problems. I claim that neither result is acceptable.

In chapter 1, I examine two versions of consequentialism that are related by a shared theoretical mechanism. This mechanism is the agent-centered prerogative, which allows an agent to give extra weight to his or her own non-optimal projects in order to preserve those projects (and thus, the possibility of a meaningful and enjoyable life for that agent) from being overwhelmed by the demands of morality. In considering the agent-centered prerogative, I develop the concept of the Polluter's Dilemma, which is the sort of large-scale collective action problem alluded to in the previous paragraph. I argue for the claim that the agent-centered prerogative generally permits the creation of these problems, and is thus unacceptable as a moral theory, especially in light of the real-world dangers presented by global climate change.

In chapters 2 and 3, I extend the line of criticism from chapter 1 to two other well-known versions of consequentialism. In chapter 2, I turn my attention to a well-known version of indirect consequentialism – that is, a type of consequentialism which does not directly assess whether an act is right or wrong based on the value of some individual act, but rather assesses rightness and wrongness based on whether the action satisfies some other criteria which, if chosen, would maximize the good. In particular, I consider how the most well known version of rule consequentialism would deal with the problem of the Polluter's Dilemma and argue that rule consequentialism faces a serious dilemma; that it either permits the generation of Polluter's Dilemmas, or fails to avoid the Demandingness Objection.

One other prominent tactic that consequentialists who attempt to avoid the Demandingness Objection use is to appeal to the notion of having a fair share of good that one is required to generate. They accept the consequentialist position that we have a general reason to promote the good, but they also accept the thesis that we are only required to generate a certain amount of good. Perhaps the most prominent example of this type of consequentialist theory is Liam Murphy's. Murphy claims that we are only required to sacrifice our time and resources up to the point where everyone would be required to in a situation where everyone complied with the demands of morality. I argue, in chapter 3, that Murphy's theory has been moderated so much that it regularly fails to make intuitively appropriate demands. However, I obtain an important result; Murphy's theory does not permit the generation of Polluter's Dilemmas.

The ultimate point of this dissertation, however, is not a negative one. My ultimate goal is to learn about what the best, most plausible version of consequentialism would look like. By the end of chapter 3, having seen the shortcomings of several versions of moderated consequentialism, we have learned several things about how the best version of consequentialism would work, and in the opening section of chapter 4 I summarize and defend these findings. A key component of this defense is a justification of my contention that generation of Polluter's Dilemmas is both unacceptable and constitutes a genuine constraint on acceptable versions of consequentialism, and this justification comes in chapter 4 as well.

Chapter 4 does not end there, however; nor does the positive project of this dissertation. I make a case study out of an interesting and novel version of consequentialism that I believe avoids both the Demandingness Objection and the

Polluter's Dilemma. In addition, in the final part of chapter 3, I consider what can be done to improve Liam Murphy's theory of beneficence, and I propose a modified version of Murphy's theory that, I believe, avoids the problem Murphy's theory faces with appropriate demands.

Chapter 1: Permissiveness and the Agent-Centered Prerogative

1. Introduction

One general strategy of attempting to defuse the Demandingness Objection is to moderate the demands that a consequentialist moral theory makes. There are several tactics one might make use of to achieve this end, but in this chapter I will focus on one; the agent-centered prerogative. The prerogative's most notable proponent is Samuel Scheffler, who has made it the centerpiece of the Hybrid Theory of Morality that he has proposed in his book, *The Rejection of Consequentialism*,¹⁷ and defended it against objections in later writings.¹⁸ The prerogative has also been adopted in modified form by Tim Mulgan, who has used it as a crucial ingredient in a novel and complex form of consequentialism which, Mulgan believes, avoids the Demandingness Objection.¹⁹

In what follows, I shall argue that these two prominent versions of the agent-centered prerogative are not acceptable ways of solving the Demandingness Objection. I argue that both versions of the prerogative are too permissive, in that they permit agents to cause harms in the pursuit of their own goals. In arguing this, I draw upon and develop an objection first raised by Shelly Kagan. In section 2, I first discuss Scheffler's version of the prerogative and argue that, despite Scheffler's defense of it, it cannot serve as a suitable way of dissolving the Demandingness Objection. In sections 3 and 4, I then

¹⁷ Scheffler 1994.

¹⁸ Scheffler 1992.

¹⁹ Mulgan 2001.

extend this critique to Mulgan's version of the prerogative and argue that the prerogative causes Mulgan's theory to generate a paradoxical result. His theory seems to recommend drastic action, but at the same time permits agents to defect from the course of drastic action far too easily. In section 5, I conclude by arguing that the discussion of Mulgan and Scheffler's prerogatives reveals a general reason why the agent-centered prerogative, as a tactic by which the consequentialist may avoid the Demandingness Objection, is unacceptable. This general reason, I argue, is that by allowing agents the freedom to privilege their own projects and opt out of doing the actions required to maximize the general welfare, agent-centered prerogatives will generate situations in which each person's position is less good overall than it would be in a situation where agents are not permitted to privilege their own projects and interests.

2. Scheffler's Original Agent-Centered Prerogative

It will first be helpful to review Scheffler's version of the agent-centered prerogative, his motivation for introducing it, and the criticisms that have been leveled against it. The motivation for Scheffler's introduction of the prerogative is to respond to concerns eloquently raised by Bernard Williams.²⁰ Williams's concern is that consequentialism threatens to make morality into a pursuit of such overriding importance that it alienates us from the pursuit of our own projects and interests. Scheffler's Hybrid Theory is an attempt to combine consequentialism with a special status granted to the agent's own personal point of view, in effect giving the agent some control over whether she ought to respond to the demands of morality or pursue some project of her own

²⁰ Williams 1973.

instead. What Scheffler's version of the agent-centered prerogative consists in, in situations when we are weighing the costs to ourselves of doing an action against the cost to others of not doing that action, is allowing us to give weight to our own interests such that the weighted value of our own interests will sometimes outweigh the unweighted impersonal value of our actions. Scheffler says that this will enable us to save the agent's personal point of view in the following manner:

... It would then allow the agent to promote the non-optimal outcome of his choosing, provided only that the degree of its inferiority to each of the superior outcomes he could instead promote in no case exceeded, by more than the specified proportion, the degree of sacrifice necessary for him to promote the superior outcome. If all of the non-optimal outcomes available to the agent were ruled out on these grounds, then and only then would he be required to promote the best overall outcome.²¹

In general, the factor by which we give weight to the costs to ourselves will be quite high. This is because if the factor is set too low, then the Hybrid Theory will be too weak to deflect the Demandingness Objection, especially when relatively small donations have the potential to do a lot of good in the hands of the appropriate aid organization. So we must set the weighting factor high. Scheffler is not specific about just how high the weighting factor is. It must, however, be enough to give us a prerogative that will be of some use to us. Tim Mulgan, in constructing an illustration of how Scheffler's prerogative is supposed to work, suggests that agents might be allowed to give 600 times more weight to their own interests than to the interests of others. This is bound to seem unusually high to some readers. Nevertheless, nothing of importance seems to hang on the exact number; it is simply that the weighting must be high if Scheffler's prerogative is

²¹ Scheffler 1994, 20.

to be of any use in helping us to avoid being overwhelmed by an obligation to aid the desperately needy in far-flung places.²²

Shelly Kagan's challenge to Scheffler's view is that the agent-centered prerogative permits agents to not only allow but actually cause harms in the pursuit of their own non-optimal projects.²³ The challenge may be illustrated by an example. Consider two cases in which I require a large sum of money to pursue my own projects successfully. Let us stipulate that this sum of money is \$10,000, and I require this amount because I have an antique car that I want to transport to a car show. In one version of this case, I already have \$10,000, and I chose to spend it on transporting my car, instead of donating the \$10,000 to charity organizations that would have saved the life of a stranger. In a second version of this case, I do not have \$10,000, so I secretly kill my uncle in order to inherit \$10,000, which I then use to transport my car to the car show.²⁴ Kagan argues that there is a morally significant difference between these two cases that Scheffler's Hybrid Theory is unable to take into account. According to the Hybrid Theory, I am permitted to value my own project that costs \$10,000 over the life of another person, thereby allowing that person to die. But because the Hybrid Theory contains no restriction on the types of action that an agent might pursue, I am equally permitted to value my own project that costs \$10,000 over the life of another person, and killing that person in doing so. So, according to Kagan, it is the lack of a deontological restriction - an agent-centered restriction - that generates such troublesome cases for Scheffler.

²² Mulgan 2001, 147.

²³ Kagan 1984.

²⁴ Kagan 1984, 251.

One might wonder exactly what it is that differentiates Scheffler's Hybrid Theory from ordinary consequentialism on this point. After all, ordinary consequentialism also permits us to cause harms in some cases. For example, if I am a commissioned officer on board a naval ship and my ship will sink unless I order a subordinate to perform an extremely hazardous repair in a burning room that will surely result in his death, consequentialism would say that I ought to order the subordinate to his death in order to prevent a greater catastrophe, which would be the loss of the ship with all hands. Other instances of this general type of case can be constructed, and one might question the consequentialist about whether such cases are generally permissible. I will not pursue that line of thought. What I am interested in pointing out is that there is a general difference in what ordinary consequentialism permits and what Scheffler's Hybrid View permits. This general difference is that ordinary consequentialism permits agents to cause harm in pursuit of an optimal outcome,²⁵ whereas Scheffler's Hybrid Theory, as a result of his inclusion of an agent-centered prerogative, must not only sometimes permit harms caused in pursuit of an optimal outcome, but also sometimes permit agents to cause harm in pursuit of a non-optimal outcome.²⁶ The case of the car show mentioned earlier is an example of this general type of case. Because the \$10,000 I spend on taking my car to the car show would do far more good if spent elsewhere, and because ordinary consequentialism does not permit me to weigh my interests more heavily than those of other people, ordinary consequentialism would not permit me to take my car to the car show. And because this action does not maximize the good to begin with, ordinary

²⁵ The thought here is that the status of that outcome as the optimal one, even when combined with the harm caused in pursuit of it, justifies the causing of the harm.

²⁶ On this point, Mulgan agrees and thinks that this makes Scheffler's position look bad as compared to what Mulgan calls Simple Consequentialism and what I have just been calling ordinary Consequentialism. See Mulgan 2001, 153.

consequentialism would furthermore not permit me to cause extra harm in pursuit of this aim. So Kagan's case would not pose any problem for ordinary consequentialism. It is, however, a problem for Scheffler's Hybrid Theory precisely because the Hybrid Theory allows us to weigh the costs to ourselves differently than the costs to others.

In responding to Kagan, Scheffler says that if his theory permits us to allow harms of a certain size to befall others for us to pursue our own projects, then it must also allow us to directly cause harms of the same size in order for us to pursue our own projects.²⁷ But in practice, Scheffler says, agents will not find it an equal choice between allowing a harm to happen in order to pursue their own projects, and committing a harm in order to pursue their own projects. Committing harm is costlier for an agent, leaving themselves less time and resources for the projects that they would rather be working on. In particular, Scheffler says that if the benefit an agent is looking for is only obtainable by an act of killing, then the agent can obtain the benefits only along with increased mental suffering.²⁸ An agent might be filled with any combination of items on this (not necessarily exhaustive) list: self-loathing, disgust, guilt, shame, humiliation, fear of being caught, and horror at their own monstrous nature. In addition, they might suffer from "profound distortions of personality and of the capacity to lead a fulfilling life."²⁹

Such a response is plausible; it is not, after all, unheard of for criminals to turn themselves in for their crimes, and it certainly seems possible that a criminal might be consumed with guilt and self-loathing after the crime has been committed. It is strange, though, that Scheffler would rest his argument on psychological generalizations that may,

²⁷ See Scheffler 1992, 380.

²⁸ Ibid. 381

²⁹ Ibid.

after all, turn out to not apply to all agents. As Mulgan notes, an agent without the relevant psychological capacities just detailed would be permitted by the agent-centered prerogative to kill her own uncle.³⁰

But Scheffler's critics need not appeal to the existence of agents for whom killing would not be a psychological cost to make trouble for the agent-centered prerogative. Highly plausible Kagan-style objections can be developed without resorting to examples involving killing. Ramon Das has recently developed just such a case.³¹ Das asks us to imagine an accountant on the low end of the pay scale working at a highly profitable law firm. She is aware that her work is worth more to the firm than she is paid for it; still, she is able to lead a reasonably comfortable life. One day, she is able to embezzle \$10,000 without getting caught, and does so without hesitation. She reassures herself of the permissibility of her actions by this reasoning: "Scheffler's agent-centered prerogative permits me to forego donating \$10,000 to save the life of a person in desperate need, since I am allowed to give weight to my own interests and prefer a non-optimal outcome if the weighted value of making a huge sacrifice on my part is greater than the impersonal value that would be produced if I made the sacrifice. Since that is the case, there is no morally significant difference between that case and my taking \$10,000 from the law firm, because the personal value of my having this money is greater than the impersonal value of having it in the possession of the firm; my multi-millionaire bosses have more money than they know what to do with anyway. And since there is no deontological restriction in Scheffler's theory on me to avoid doing harm and no distinction between doing a harm and allowing a harm to occur, what I have done must be permissible."

³⁰ Mulgan 2001, 162.

³¹ Das 2000.

Making use of the same tactic against Das that he used against Kagan will not help Scheffler much. It is plausible to claim that killing involves a high psychological cost to (some) killers. It is, however, far from clear that embezzling from a workplace and boss that one resents involves as heavy a cost, if it involves a psychological cost at all.

If we move away from cases that involve committing harms directly against other agents and consider instead cases which involve indirect harms to other agents, we may find that the Hybrid Theory is, again, unacceptably permissive. Here I will introduce a variation on Kagan's objection to Scheffler, the intent being that Kagan-style cases are actually quite widespread and need not involve cases involving obvious crimes, such as murder or embezzlement. Consider a case that involves me working on some non-optimal project, and in doing so I indirectly contribute to a harm that befalls other agents. Let us suppose again that I am an

Antique Car Aficionado. I desire to take my car to a classic car show. This car would no doubt be the talk of the show, for it is an unusual 1937 Aston Martin of a limited production run, and being the talk of the show would benefit me greatly; both professionally, because I want to be invited to more car shows, and personally, because I have recently been divorced because I spent too much time on the car and I need to show my former partner that my investment was worthwhile. Transporting my car would involve hauling it to the car show on another, heavier car, and this hauling would involve the release of automobile emissions into the atmosphere, which pollute the environment and contribute in a small way to climate change.³²

So then, I have two options:

³² Naturally, I am following scientific consensus in the assumption that climate change is a well-established scientific fact.

1) *Stay Home*. I bear a high cost of owning and maintaining the car while gaining no benefit from it. In addition, I am saddled with the cost of knowing that I lost my relationship for a car. However, because I am not hauling a car around needlessly on another car, the maximum impersonal value for everyone else is generated, as compared to option 2.

2) *Show The Car*. I bear no cost from making any sacrifice. Somewhat less impersonal value overall is generated, because of the resulting air and noise pollution.

Now, I am a well-informed citizen who is well aware of the negative effects of my car show projects. But, because I am a moral philosophy aficionado as well and have read Scheffler in my spare time, I reason that since I am allowed to multiply the cost to myself by a large proportion and weigh that result against the difference in impersonal value between *Stay Home* and *Show The Car*, and since I find that the weighted cost outweighs the difference between the options, I am allowed to pursue *Show The Car*.

In the case just described, I rely on the assumption that my share of harm that might result from contributing to climate change is rather small -- small enough, that is, to be outweighed by the weighted value an agent is allowed to give to her own interests by the agent-centered prerogative. One might challenge this assumption with an analogy:

Shooting Range. I am standing on the Nebraska prairie with a deer-hunting rifle. I can see some deer out in front of me on the prairie, but in addition to the deer, there are several people standing on the prairie. They are widely spaced out and not standing near any of the deer, so it is unlikely that any of my shots will hit them.

There is, of course, a reasonable expectation that if I shoot, I will misjudge the wind

or my hand will not be steady, and one of my shots will go off course, thus seriously wounding or killing a person standing on the prairie. But even if I shoot and I miss a person, whether I hit a deer or not, I did something wrong. The harm lies in the negative expected value of my actions.³³

Perhaps it would be proper to think of the harms resulting from my polluting as more akin to shooting a gun, at random or in a directed manner, into a field where people are randomly and widely spaced out. We might say that the potential harms resulting from my polluting actions and the potential harms resulting from my pulling the rifle's trigger are analogous in one way because they are both diffuse harms. The harm may not be solely borne by any one person; by shooting into the field, I have harmed everyone by putting them in a dangerous situation; by driving my car, I have made the climate a little more inhospitable for everyone. The harms are analogous in another way; perhaps the harms are not diffuse. After all, it may be the case that the pollution from my car is just enough to cause a nearby asthmatic person to have a severe asthma attack, or just enough to cause a crucial bit of an ice sheet to melt, thus causing a chain reaction in which the houses of a low-lying island nation's residents are submerged. Just as with *Shooting Range*, the expectation is that people will be injured, and impersonal value will be worse off overall than if I had not exercised my prerogative to *Show The Car*. This suggestion is plausible. If the analogy holds, then I will have a much harder time demonstrating the permissiveness of the agent-centered prerogative.

It is not clear that the analogy completely holds, however. While the harms themselves may be analogous in that they may both either be diffuse or concrete, the causes of the harms are not analogous. In *Shooting Range*, I am the direct cause of the

³³ I owe this analogy to Aaron Bronfman.

harm and I can reasonably foresee that my actions will result in concrete harm to someone. In *Antique Car Aficionado*, it is both unclear that I am the direct cause of the harm, and it is not the case that I can reasonably foresee that my actions will result in concrete harm to someone. At best, I am a small part in the causal network that leads to a harm from climate change, or even a harm resulting from the interaction of my car's emissions and someone else's severe asthma. The empirical facts bear this out. Climate change is not due to the actions of any one person; rather, it is due, at least in large part, to the aggregate industrial activity of billions of human beings. Furthermore, driving a heavy truck is not likely to elicit gasps of condemnation from bystanders, whereas I expect that firing a rifle in the general direction of people would elicit such gasps. Therefore, I conclude that the analogy does not hold.

If I am correct, then the harms permitted by Scheffler's agent-centered prerogative are widespread. Perhaps I desire to regularly eat filet mignon at my favorite steakhouse. Or perhaps, to make the harm even less, we might suppose that I am a vegetarian who nevertheless loves to eat gouda. The production of the filet mignon, or the gouda, inevitably involves pollution of various kinds as well, to say nothing of the harms that befall cows directly in the process of producing the filet or the gouda. Nevertheless, my share of any harm that befalls other people as a result of pollution or climate change, or as a result of my participation in the production and consumption of meat and/or dairy, is quite low. Because I am not directly causing huge harms to befall other people, and because the good which I derive from pursuing my own projects can be quite intense for me, the extra weighting that Scheffler's agent-centered prerogative permits me to give

my projects will have the result that a very wide range of sub-optimal behaviors with collectively terrible long-term consequences for all agents will be permitted.

Perhaps Scheffler might say that I have made a mistake in thinking that my share of contributing to global warming is small; in reality, my impact is larger than I think, and I ought to be weighing the personal value of my actions against the impersonal value of the foreseeable long-term effects of climate change. Perhaps the best version of his agent-centered prerogative requires me to think about weighing my non-optimal interests against the impersonal value that might foreseeably result from me (and everyone else) doing optimal actions. If that is Scheffler's response, then I would find it hard to see what the principled rule is by which I decide whether to weigh my non-optimal interests against the impersonal value that can be directly derived from my optimal actions or against the impersonal value that might be indirectly derived from many people doing optimal actions. But perhaps more importantly, making this shift would seem to me to make the impersonal value of my optimal actions so great that the agent-centered prerogative ineffective at allowing me to sometimes prefer my own non-optimal projects. So this response is not available to Scheffler.

Scheffler's version of the agent-centered prerogative, as we have seen, is subject to the charge that, in the absence of deontological restrictions, it allows agents to commit harms in pursuit of their non-optimal projects.³⁴ Although it is possible that Scheffler is correct and agents generally find it costlier to commit harms than to make sacrifices, my project in this section was to demonstrate that it is possible to generate Kagan-style counterexamples to Scheffler's view that do not rely on harms that consist in obvious crimes; indeed, it is possible to describe harms that Scheffler's prerogative allows us to

³⁴ This is Kagan's objection.

commit that not only do not rely on the harms in question being obvious crimes at all, but in fact work in the long run to produce a sub-optimal outcome. If the agent-centered prerogative is to remain a live option for the consequentialist seeking to avoid the Demandingness Objection, we will need to find a better version of the prerogative. The question remains, however; is Tim Mulgan's version of the agent-centered prerogative any better than Scheffler's?

3. An Improved Prerogative: Mulgan's Combined Consequentialism

One of Mulgan's most significant contributions to the literature on the Demandingness Objection comes in his book, *The Demands of Consequentialism*. Here, after extensive discussion and criticism of many previous attempts to avoid the Demandingness Objection, he proposes a theory that he terms Combined Consequentialism, which is so named because it is a combination of several different consequentialist theoretical mechanisms. Combined Consequentialism's basis is Mulgan's distinction between two "realms of moral choice," the Realm of Necessity and the Realm of Reciprocity.³⁵

Mulgan's thesis is that different consequentialist theories apply in the different realms. In the Realm of Necessity, current members of the moral community are faced with the choice of whether to enable those excluded from the moral community to participate in it, the crucial interaction being between those who are currently active

³⁵ Mulgan 2001, 172. Mulgan actually claims that there are three realms of moral choice. He terms the third realm, which I do not discuss here, the Realm of Creation and says that this realm involves our decisions about bringing lives into existence. For the purposes of his book, however, he ignores the Realm of Creation and focuses instead on the Realms of Necessity and Reciprocity.

members and those who cannot even participate.³⁶ But in the Realm of Reciprocity, equal members of the moral community decide how to interact with one another. According to Mulgan, when we are making decisions in the Realm of Necessity, maximizing the overall good is the right approach. This requires us to make sacrifices in order to turn potential members of the moral community into active members of the moral community. When we are making decisions in the Realm of Reciprocity, however, we ought to follow a rule consequentialist code of conduct that hews more closely to common-sense morality, and avoids some of the various problems associated with simple act consequentialism.

These realms are generated by the distinction that Mulgan draws between different categories necessary to human well-being. Mulgan's account of well-being, on the whole, resembles what Derek Parfit has referred to as an Objective List Theory of Welfare.³⁷ On this sort of view, well-being consists in satisfying a number of different categories that work together to promote one's overall sense that one's life is going well, and whether they are good for us or not does not depend on what we think of them or how much happiness they produce for us.

Mulgan's view on well-being, and the basis for his distinction between the two moral realms, is that individual well-being is comprised of the satisfaction of needs and goals. Needs include the things that are necessary for an agent to take part in the moral community, among them food, shelter, freedom from unnecessary suffering, and the like. Needs are the foundation for goals. Once a person has their needs fulfilled, they are in a position to set goals for themselves. Goals, for Mulgan, are the things that the satisfaction

³⁶ By "members of the moral community," Mulgan means those who are relatively equally well-off agents who can interact in mutually advantageous ways.

³⁷ Parfit 1984, 4.

and setting of which determines, all things considered, how well a person's life goes. But goals only matter morally if they are achieved in the right way; they require autonomy, which can only be had if the agent's needs are satisfied. Needs are not like this; to Mulgan, it does not matter how needs are promoted, since the fulfillment of the need is what gives the need its value.

The third major component of Mulgan's Combined Consequentialism is a modification of Scheffler's agent-centered prerogative. As with Scheffler, the motivation for including the prerogative in Combined Consequentialism is to ensure the agent's autonomy to pursue her own goals. Mulgan, having extensively reviewed the variety of objections to Scheffler's Hybrid View and having stated his own, modifies the prerogative by replacing the constant with a variable; thus, "As the cost the agent must bear to produce a given amount of good increases, the weight she is allowed to give to her own interests also increases."³⁸

We are now in a position to see how Mulgan's Combined Consequentialism works.

Here is a case that Mulgan frequently refers to:

Affluent's Tale. Affluent is, as her name implies, an affluent citizen of a developed country, sitting at her desk with a checkbook in one hand and two pamphlets in the other. One of the pamphlets is from a reputable international aid organization. The other is from a local theater company, advertising their latest production. Affluent has enough money to make a donation to the charity, or to buy tickets to the theater, but not enough money to do both. Because of her love for the theater, she buys the tickets, even though she knows the money would have done much more good in the hands of the charity.³⁹

What Affluent is permitted to do, under Combined Consequentialism, depends on what sort of sacrifices she has made already. If she has not previously made any

³⁸ Mulgan 2001, 268. For Mulgan's review of the problems with Scheffler's theory, see Mulgan 2001, chapter 6.

³⁹ Ibid. 4. Adapted with minor modifications.

contributions to charity, Combined Consequentialism would intuitively require her to donate a large portion of her money to charity - most probably, the factor by which she is allowed to give weight to her own interests is not yet large enough to outweigh the needs of others. If she has previously made donations, then the increasing cost to Affluent warrants her giving more weight to her own projects, until at some point the weighted cost to Affluent outweighs the difference between the impersonal value produced by donating and the impersonal value produced by not donating.

It seems Mulgan's Combined Consequentialism would deal adequately with Kagan's original objection and Ramon Das's embezzlement case. Das's accountant, presumably having sacrificed nothing prior to her decision to embezzle \$10,000, would not be permitted under Mulgan's prerogative to favor her own interests and embezzle \$10,000. If we stipulated that she is a generous person who has already given much of her money to charity, Mulgan's prerogative might permit her to embezzle \$10,000; but another component of Combined Consequentialism - that of Rule Consequentialism, which governs the Realm of Reciprocity - would not allow it. To see why, note that Das's accountant is not desperately needy. In Mulgan's view, she is a full member of the moral community. Her assertion that she needs the money more than her multimillionaire bosses does not matter like it would for someone whose needs have gone unmet, for she does not "need" the money. It would, rather, satisfy a goal of hers to have more money. But according to Mulgan, goals only matter morally if they are satisfied in the right way - the way we have reciprocally agreed on to be the best way, which is why rule consequentialism governs this realm. Since the optimal set of rules will forbid stealing,

Mulgan's theory is not subject to Das's modified Kagan-style objection. Mulgan's theory thus is an improvement over Scheffler's.

We need to ask, however, whether Mulgan's theory is adequate to deal with all Kagan-style objections. I will suggest that this is not the case. Specifically, I argue that Mulgan's view is subject to my modified Kagan-style objection, which I introduced in the previous section and which I develop more in the next section.

4. Polluter's Dilemmas and the Demandingness of Combined Consequentialism

We are living in an age of unprecedented climate change. The scientific consensus has it that aggregate human activity since the Industrial Revolution is a major contributing factor to the warming of the planet. The warming of the planet has catastrophic implications, among them: increasing desertification of regions near the Earth's equator; increasing numbers of catastrophically powerful storms because of the warming oceans, especially in tropical regions; increasing levels of famine in the developing world; increasing levels of migration from regions that become intolerably hot; increasing competition for increasingly scarce resources; the rise in sea levels because of the disappearance of the polar ice caps; and the alteration of the food chain of the oceans, which will likely result in the disruption of the food supply derived from the ocean, to name a few possible outcomes of the current state of affairs. Though it is clear that the Earth experiences climate fluctuations naturally, it is hardly deniable that technologically-enabled human activity has not played some role in the current precarious state of affairs.

Every day, we make choices which can contribute to, lessen the impact of, or completely alleviate various moral problems which confront us. Some of these daily choices to make include: what we spend our money on, how to get ourselves from one place to another, which foods to eat and which drinks to drink, how much food to eat, and so forth. Take the example of choosing one's form of transportation. Common options here include walking, bicycling, taking public transportation, or driving one's own car. Such a choice may seem innocuous, and indeed seems so to many people; the same is true of the other problems mentioned. The downside of such actions is often obscure. It is not obvious what sort of harm there could be from making the choice to drive one's own car. Yet the moral problems that these choices can contribute to are anything but innocuous.

I am attempting to motivate and systematically describe a certain kind of problem case. Derek Parfit, in *Reasons and Persons*, describes this kind of case as a "Many-Person Dilemma."⁴⁰ There are many kinds of Many-Person Dilemmas, but the basic template of the dilemma is the same: If most or all agents choose to benefit themselves in a certain way, the overall results will be worse for all. Consider, for example, the result of all the individual tokens of the decision to use one's own automobile to get from one place to another. One instance of such a decision imposes no particular burden on anyone else.⁴¹ However, the same decision, repeated many times for decades by the overwhelming majority of all citizens of the developed world and a non-trivial number of citizens in the developing world, has contributed to a situation which threatens the general well-being. One instance of a decision to get one's food from an unsustainable

⁴⁰ Parfit 1984, 59-62.

⁴¹ I defended this claim in section 2 of this chapter against the suggestion that our individual share of the harms of climate change are quite large.

food source does not impose any great burden on anyone, though it is perhaps harmful in some small way.⁴² However, the same decision, repeated countless times for decades by the overwhelming majority of all citizens of the developed world and a non-trivial number of citizens in the developing world, has contributed to a situation which threatens the overall well-being of all. For brevity's sake, I shall refer to these types of Dilemmas as *Polluter's Dilemmas* in what follows.

How are Polluter's Dilemmas like the objections that Kagan raises and that Das has further developed, and how are they different? They are similar because they propose that the agent-centered prerogative allows an agent to create harm in pursuit of her own goals, because there is no deontological restriction on doing harm. This is one of Kagan's goals in objecting to Scheffler's Hybrid Theory; the other, naturally, is to point out that the difficulty arises in part because the Hybrid Theory lacks the distinction between doing and allowing. Here I am focused on the first aspect of Kagan's objection that I have listed. My argument shall be that Mulgan's version of the agent-centered prerogative permits Polluter's Dilemmas to occur, because (as argued in section 2 of this chapter) each individual's share of the harm done by pollution is small enough to be easily outweighed by the weighted value accorded to an agent's interests by the agent-centered prerogative. Furthermore, I shall argue that Mulgan's theory, despite its ostensible goal of being a moderate consequentialist theory, turns out to be far too permissive in the case of Polluter's Dilemmas, even when the Polluter's Dilemma seems to fall into the Realm of Necessity.

⁴² I pursue this thought in section 5, where I discuss a distinction between harms experienced diffusely and harms experienced in concentration.

First, I reiterate a crucial argument of the second section of this chapter. We saw previously that Scheffler's agent-centered prerogative permits Polluter's Dilemmas to occur. It permits them to occur because it permits an agent to multiply the cost incurred by them of making a sacrifice by the weighted value of the agent-centered prerogative, and then compare the result against the difference in impersonal value between the superior and inferior alternatives. I argued that, frequently, the agent-centered prerogative will generate the result that agents are allowed to prefer their own interests and carry out an action that involves pollution. I also argued against the possible objection that the harms involved in one individual's polluting act are actually quite great. This was accomplished by showing that a proposed analogy fails to be analogous to the case of the harms involved in pollution.

It would not do to simply reiterate the points above and proceed to the second part of my argument against Mulgan, however; I must take into account the modification that Mulgan has made to his version of the agent-centered prerogative. According to Mulgan, an agent is not permitted to assign a constant weighted value to her own interests; rather, an agent is allowed to assign more weight to her own interests if she is asked to bear more sacrifice. This has the result that if an agent has either made many sacrifices in the past or is being asked to bear a very large sacrifice right now, that agent may assign a high weight to her own interests. If she has not previously made any sacrifices, Combined Consequentialism does not permit her to perform her preferred non-optimal action; she must, instead, make the sacrifice for the impersonal good and bear the cost.

Thus, in *Antique Car Aficionado*, if I have not previously made any sacrifice toward the greater good, there are two possibilities. The first is that it is possible that I

would be required to refrain from taking my car to the car show. It is possible that, since I would not be permitted to assign any extra weight to my own interests, the unmodified cost that I would be required to bear would not be high enough to outweigh the impersonal value that my sacrifice would generate. The other possibility is that, because the harms involved in my act of pollution are so small, the unmodified cost that I bear would, by itself, allow me to prefer my own interests.

I do not believe that it matters which possibility is actually the case. If the second possibility is the case, then it will be trivially easy for me - whenever I experience a moderate inconvenience - to truly claim that my interests trump the general good. If the first possibility is the case, then I will still be permitted to perform the polluting action if I make some relatively minor sacrifices elsewhere. Perhaps I constantly spend time and effort recycling what can be recycled out of my garbage. Perhaps, instead, I go to great effort to recycle electronic waste. Whatever my sacrifice, this sacrifice will enable to me put the agent-centered prerogative to work, pursuing my own interests (and thereby committing polluting actions) when it matters to me.

Now, what does Mulgan's Combined Consequentialism say about such cases? In trying to answer this question, a natural starting point is to ask the question of which realm of moral choice such cases are sorted into.⁴³ There are a couple of choices: decisions about such choices may be sorted into the Realm of Reciprocity, or into the Realm of Necessity. The Realm of Reciprocity, it will be recalled, is the realm of moral choice in which agents interact with one another on more or less equal grounds; it is where individual members of the moral community promote their own interests, work on

⁴³ Here, following Mulgan, I am ignoring his third realm, the Realm of Creation, which he ignores in Mulgan 2001.

joint projects together, and set and achieve goals for themselves under conditions of their own autonomy; according to Mulgan, this realm is to be governed under rule consequentialist terms. By contrast, the Realm of Necessity is where agents interact in an unequal way; in this realm, members of the moral community are confronted with people who are not yet members of the moral community, and are required under act consequentialist terms to sacrifice either up to the point where such members all have their needs fulfilled, or up to the point where Combined Consequentialism's version of the Agent-Centered Prerogative kicks in, allowing the ones who sacrifice to pursue their other projects.

Perhaps because Mulgan's theory was not designed with Polluter's Dilemmas in mind, it is possible to describe ways in which Polluter's Dilemmas may be addressed in both realms of moral choice. Some kinds of choices affecting the ways in which we interact with our environment are made at the collective level, and this suggests that such a Polluter's Dilemma falls into the Realm of Reciprocity. We freely and autonomously interact with one another, and on the basis of both the best available evidence of the risks involved and the extent of our own willingness to put up with restrictions on our autonomy and the types of goals we may pursue, we collectively decide (and constantly revisit our collective decisions) on what types of goals and projects are permissible with regard to our environment, and what types of goals and projects are incompatible with the protection of our environment. The environment is valuable to us, on this view, because to damage our environment is to damage our ability to achieve our goals. Our response to the Polluter's Dilemma would be essentially on rule consequentialist grounds.

This answer is insufficient. Given the myriad possible catastrophic effects of climate change canvassed above, it would be naïve to say that such changes our environment have the ability to affect us only at the level of our ability to pursue our own goals as free and equal members of the moral community. We should, instead, think of the Polluter's Dilemma as falling into the Realm of Necessity. If Polluter's Dilemmas fall into the Realm of Reciprocity, then the harm from Polluter's Dilemmas must be confined to the harms covered by the Realm of Reciprocity. These sorts of harms include damage to our own autonomy, damage to our ability to pursue the kinds of projects that hold significance for us and that we freely choose to associate ourselves with. But this is not what we find. The possible harms of climate change are not like harms to our political freedoms, such as our freedom to choose our own religious path and the freedom to associate with whom we please. Rather, they are harms that strike at our *needs*, in Mulgan's sense of that term. They threaten our food supply, our shelter, and our general safety. It is true to say that the effects of climate change harm our autonomy and our ability to pursue our own goals, but they do so not by striking directly at those things, but rather by causing our needs to go unfulfilled. Without those needs fulfilled, we can hardly be expected to interact with each other as free and roughly equally positioned members of the moral community.

There is a second reason to think that Polluter's Dilemmas fall into the Realm of Necessity. This second reason is that the burdens on human lives imposed by climate change will fall most disproportionately on the desperately needy. Those who are already desperately needy do not have the resources to respond to a changing environment as easily as those who do have the resources. People living in developing nations are,

according to the United Nations Development Programme, are up to 79 times more vulnerable to the effects of climate disasters than people living in developing nations.⁴⁴ Desertification and the disruption of water systems could mean that hundreds of millions of people in Africa and Asia will face water shortages.⁴⁵ Some low-lying countries with high population density, such as the Netherlands and Bangladesh, face the loss of much of their land area. Bangladesh, in particular, could lose up to 18 percent of its land area.⁴⁶ Likely, the Netherlands will be better able to adapt to climate change, given its position of relative wealth. Bangladesh, one of the poorest nations in the world, will have more difficulty in this regard. Low-lying coastal areas that have a disproportionately high share of a nation's GDP, such as in wealthier coastal cities in India, China, and Guyana, are at risk as well, and the disruption of such areas can have a devastating effect on the economic development and well-being of entire nations.⁴⁷

I conclude that we ought to think of Polluter's Dilemmas as primarily falling into Mulgan's Realm of Necessity. This means that our decision making about Polluter's Dilemmas is to be governed under an act consequentialist framework; that is, we are to act so as to maximize the good, at least until we reach the point where Mulgan's version of the agent-centered prerogative kicks in. This conclusion probably implies a number of changes in our own behavior at the individual level, including things like adopting a vegetarian diet, driving one's car far less, riding one's bicycle far more, purchasing carbon offsets to alleviate the costs of our polluting behavior, living in a smaller and more energy-efficient dwelling, and generally attempting to tread more lightly on the Earth's

⁴⁴ United Nations Development Programme 2009.

⁴⁵ Ibid.

⁴⁶ Ibid.

⁴⁷ Nicholls et al, 2007.

surface than we have been doing by giving up some of the luxuries to which we have grown accustomed.

Whether one believes that climate change is better dealt with through penalizing polluting behaviors or through encouraging ecologically friendly technological development, the above recommendations in general seem like a plausible response to global warming, if widely adopted enough. Would Mulgan's view recommend such actions? Yes, it would; but the agent-centered prerogative, paradoxically, would allow agents to opt out of doing those actions far too easily. If Mulgan's view did not contain his version of the agent-centered prerogative, he could seemingly not escape the Demandingness Objection in the case of the Polluter's Dilemma. But I suggest -- because of my previous arguments that the harms leading to Polluter's Dilemmas are very small, and that agents will be allowed by Mulgan's agent-centered prerogative to prefer a non-optimal action very easily in the case of Polluter's Dilemmas -- that Mulgan's view allows us to far too easily exercise the agent-centered prerogative. Thus, Mulgan's view has the potential to generate Polluter's Dilemmas.

Mulgan's view thus succeeds in avoiding the Demandingness Objection. But this seems like a Pyrrhic victory. Indeed, the entire rationale for developing Combined Consequentialism was to find a plausible, moderate form of Consequentialism that made reasonable demands. Given the possible dangers posed by climate change, it seems like a plausible moderate form of Consequentialism should tell us to take a couple kinds of action; action to best position our society for the future, and action to mitigate the damage already done. But instead, the result was to find a theory that can permit agents to continue causing damage.

5. Why the Agent-Centered Prerogative is Generally Unacceptable as a Component of a Consequentialist Moral Theory

Where does this leave the agent-centered prerogative in general, then? The answer to this question, I believe, is that the very notion of an agent-centered prerogative is generally unhelpful for the Consequentialist seeking a way to dissolve the Demandingness Objection. The agent-centered prerogative will, in short, create more problems than it solves. The reason for this is that the agent-centered prerogative generates situations in which agents are worse off than they would be in situations where no agent-centered prerogative exists. In other words, Consequentialist theories augmented with an agent-centered prerogative are and will turn out to be directly collectively self-defeating. Derek Parfit has proposed this notion, in the first part of his book *Reasons and Persons*, as a formal constraint on moral theories, and suggests that an acceptable moral theory ought not to run afoul of this constraint.⁴⁸ According to Parfit, a moral theory is directly collectively self-defeating when, if universally and successfully followed by a community, it would make the overall state of affairs for each agent worse than the situation would be in which some other moral theory were universally and successfully followed.⁴⁹

⁴⁸ For an argument against this notion, see Mendola 1998.

⁴⁹ My formulation of the notion of direct collective self-defeat here is a paraphrase of Parfit's official definition. I give this paraphrase to encompass two versions of the notion of direct collective self-defeat that Parfit gives: one for moral theories that give the same goals to each agent, and one for moral theories which give different goals to each agent. Consequentialist moral theories with an agent-centered prerogative give the same goals to each agent, the goal being to maximize the good, but allow agents to occasionally pursue other goals. Thus, the former of Parfit's versions of direct collective self-defeat seems to apply naturally to these types of Consequentialist theories. The other goals that an agent may have are not directly given by the theory. See Parfit 1984, pp. 54-55.

Let us review, briefly, what we have seen. I have examined and argued against two versions of the agent-centered prerogative. Against Scheffler, we have seen that it is possible to generate a wide variety of situations in which agents will not find it costly to favor their own projects, and thereby cause harm in doing so. The problem is particularly acute in the case of pollution. Because the harms of pollution are small, Scheffler's agent-centered prerogative permits agents to cause a great deal of pollution in pursuit of their own projects. Against Mulgan's view, we can extend the same Kagan-style objections familiar from our discussion of Scheffler and develop them into the following critique of Mulgan: Combined Consequentialism, applied to the case of climate change and our response to it, generates seemingly demanding recommendations for dealing with the problem but, paradoxically, permits agents to defect from the actions that maximize the general welfare and prefer their own interests far too easily.

The general feature that both of these theories share is a commitment to the notion that agents are occasionally permitted to give the cost that they bear from making a sacrifice more weight than the cost the general population would bear if they did not make that sacrifice. This commitment, of course, stems from the desire to protect the agent's personal point of view from being overwhelmed by the demands of morality. Both of these theories, moreover, lack restrictions on the types of harms one might commit in the pursuit of one's own projects. As we have seen, Mulgan's view includes a type of restriction - namely, that an agent may not favor her own projects when she has not previously made any kind of sacrifice. But as I have argued, this is essentially a weak restriction that does not, in pollution cases, require a great deal of sacrifice to overcome.

It is weak in this way because it is not a restriction on the types of actions that the agent may do, but it is instead a restriction on when the agent may promote her own interests.

I argue that any moral theory that shares this commitment will be subject to the Polluter's Dilemma, and thus, to the result that it is a directly collectively self-defeating moral theory. To see why this is so, imagine us to be in a world where there is a consequentially optimal level of pollution, such that if we collectively adopt more stringent restrictions on polluting we will be worse off because industry, research, and commerce will be too restricted, and such that if we collectively adopt more relaxed environmental standards we will be worse off because climate change will outstrip our collective ability to adapt to it. In this situation, let us further stipulate that there is a specified amount of pollution that constitutes the maximum amount of pollution we are consequentially licensed to produce, such that producing more pollution beyond our personal maximum amounts to a departure from our previously-stipulated consequentially optimal level.

What does the agent-centered prerogative say about pollution beyond the consequentially optimal level? Here I am assuming, based on my discussion of the Shooting Range case and the discussion of the harms of pollution in sections 2 and 4, respectively, of this chapter, that the harms from each unit of pollution are small but equally harmful. In addition, the agent-centered prerogative must allow agents to value their own projects and interests many times more than the general welfare if it is to be effective in allowing the personal point of view some space. With these two assumptions in hand, plausibly, the agent-centered prerogative will permit agents to generate pollution above and beyond their personal maximum amounts. The ways in which this may occur

are not difficult to imagine. Perhaps I have been limiting myself in the amount of steaks I have been enjoying because, under my society's general pollution restrictions, I am not permitted to enjoy more steaks than I have already eaten for a specified time period. Perhaps I have been taking public transportation and riding my bicycle more often because I have already used up my allotment of personal car trips for a specified time period. The list can go on. But the agent-centered prerogative, in each case, will allow me to value my interest in eating steak, or getting to where I want to go more quickly, more highly than the harms generated from meat industry- and car-related pollution. By stipulation, though, in such a situation, every agent will be worse off because climate change will outstrip our ability to adapt to it.⁵⁰

One might object: "This does not necessarily constitute a case of being directly collectively self-defeating, for one of the given goals of the moral theory is to permit the agent some space for their personal point of view. Here, everything is fine because the agents are successfully following the moral theory that includes the prerogative, so the theory cannot be objected to on your grounds." I would reply that we are discussing consequentialist moral theories, whose general aim is that the good, whatever it is - overall utility, satisfaction of desires, etc. - be maximized. Protecting each agent's personal point of view is not a goal given to each agent by the theory. The theory merely permits agents to protect their own point of view; they are always, however, permitted to make sacrifices.

Still, might one not think that this result is perfectly fine? After all, protecting the agent's personal point of view does seem like a worthwhile goal, even if an unpalatable

⁵⁰ I thank David Sobel for suggesting this formulation of the general objection to the agent-centered prerogative given in this paragraph and the previous paragraph.

side-effect of this addition to a consequentialist moral theory is that we have the option to benefit ourselves a little bit at the cost of harming others a lot (measured from the point of view of overall value, of course). This, I claim, is not plausible. A central consequentialist insight, going back to Bentham and Mill, is that the promotion of the general good is most important, and thus doing what it takes to promote this end is what the good consequentialist ought to do. Working toward this end plausibly includes promulgating and observing regulations that protect our collective living space. Allowing those regulations to be subverted by a sub-optimal action stemming from some individual's idiosyncratic desire is antithetical to the consequentialist project. Put more generally, it is important that people get what actually will promote the general well-being. Allowing individuals the permission to benefit themselves at the expense of harming many others a little bit will not generally promote well-being. It is an empirical matter whether normal human beings would, in practice, avail themselves of the opportunity to benefit themselves a little at the expense of harming others a lot, naturally. But it does not seem like a point that stands in need of extensive defending, as I suspect we are all familiar with cases where we ourselves have done such actions, to say nothing of our familiarity with other people who regularly do such actions.

A further objection that one might give could run this: "Your general objection to agent-centered prerogatives is based on the assumption that each unit of pollution is equally harmful. But why should we believe this? Perhaps pollution is not harmful until it reaches a critical mass. To see that this is so, imagine a pristine Eden of a world where no action has ever sullied the environment, and people live in an enlightened and harmonious relationship with nature. Now imagine a typical early 21st century

automobile inserted into this world, turned on, allowed to idle in neutral gear for five seconds, and then turned off, never to start up again. Surely nothing has been harmed by the minor amount of pollution that has been inserted into this world. Or, if we want to look even harder for a case of harmless pollution, suppose that in this same pristine Eden, I travel far from human habitation to the peak of the world's highest mountain, and I release a single drop of crude oil from an eyedropper onto the mountain's rocky summit. This oil is far from any source of water, and there is no reasonable chance that it will ever make its way into the world's water supply. It seems obvious, then, that it is only in great quantities that pollution is harmful, which puts pressure on your assumption that each unit of pollution carries a small amount of harm, and each unit is equally harmful." If this objection is true, then there are some types of pollution which do not generate Polluter's Dilemmas for proponents of the agent-centered prerogative, and my general objection to the agent-centered prerogative is in danger.

I think it is possible to resist this objection. This objection seems to be saying that whether each unit of pollution is equally harmful or not depends on how it is experienced and/or whether it is experienced. I think that this objection is appealing to a kind of distinction that we can state thus: there is a difference between a small amount of pollution *experienced diffusely* and *experienced in concentration*. Let us say that we experience pollution in concentration when the amount of this type of pollution experienced has reached a critical mass in my environment such that it either directly endangers my health and well-being or produces environmental changes that indirectly endanger my health and well-being. We may then understand "experiencing pollution diffusely" as the negation of experiencing pollution in concentration: when I experience

pollution diffusely, it is not the case that the pollution directly endangers my health and well-being, and it is not the case that the changes in the environment produced by pollution indirectly endangers my health and well-being.

On this understanding of the harm of pollution, then, the potential of pollution to cause harm always exists, but that is not what matters. What matters is whether the harm is actually caused, and whether pollution actually causes harm or not depends on whether we experience it diffusely or in concentration. For example, I might experience the car's emissions as diffuse pollution, if I was standing next to it, breathing air that was overwhelmingly normal and only having a very minuscule amount of car emissions mixed in. Or I might experience it in concentration, if I wrapped my mouth around the car's exhaust pipe and inhaled deeply when it was running. If I experience even a small amount of pollution in concentration, it seems like it will be harmful to me. The proponent of this objection, though, would not necessarily find that result problematic; she, rather, would just want there to be some case in which diffusely-experienced pollution is not harmful.

Another way to understand this objection is as denying that we are at some kind of point where a consequentially optimal balance between economic and technological progress on one hand, and climate change on the other, is warranted. We are, thus, supposing that we are at some point further back on the timeline, before we have any incentive to bring our pollution regulations and our economic activities into balance with each other. Thus, there does not seem like there is anything wrong with exercising the agent-centered prerogative and generating some diffuse pollution as a result.

I would claim, though, that what the proponent of this objection overlooks is that diffuse pollution is not simply harmless when it is not in concentration; even in cases where the pollution generated as a result of other people's exercise of the agent-centered prerogative is experienced diffusely, the allegedly harmless diffuse pollution tends to linger and build up in the environment and eventually reach a point where it ceases to be diffuse. It is true that agents do not experience diffuse pollution as a direct harm, but it is not the case that diffuse pollution is not harmful *simpliciter*. When I experience pollution diffusely, the environment has not changed enough to indirectly endanger my health and well-being. But the point is that my environment is changing nonetheless, and it is changing in ways that are not conducive to my health and well-being. To drive the point home, consider the following situation. Suppose that I live in a tropical place that is often inhospitably stormy, but I am well-protected from storms by a strongly-built stone house. Suppose further that, unbeknownst to me, my neighbor is creating paintings with spray paint that has a chemical component that is not harmful to humans but has an eroding effect on stone. This component, over time, will wear down the walls of my house and leave me exposed to storms. I suspect that we would not want to count my neighbor's actions as not harmful to me, even though they do not directly injure me. His actions damage my environment by building up over time to the point where they reach critical mass (i.e., causing my house's stone walls to erode completely). But we should not say that his actions are *only* harmful at the point where they erode my walls completely. His actions are harmful because they constitute a growing threat to me in my environment that I depend on for health, safety, and well-being.

But because this threat to me grows slowly and incrementally, and because under a consequentialist moral theory containing an agent-centered prerogative, it seems like my neighbor is permitted to use his special spray paint. I am, presumably, permitted to take other kinds of actions which result in similar effects for him if I value my projects enough. And thus, we are well on our way to generating a Polluter's Dilemma. This tendency of pollution to build up is part of what makes the problem of pollution so intractable for the proponents of the agent-centered prerogative, and is a reason why we should be cautious of the notion that pollution is only harmful when it reaches critical mass. I do not want to be understood as denying that there are types of pollution like that; but I do want to be understood as claiming that in the actual world, such types of pollution are vanishingly rare. Canonical kinds of pollution, such as carbon dioxide emissions, are just the sort of thing that gets the Polluter's Dilemma going, and thus, provides us with all the reason we need to think that including agent-centered prerogatives in a consequentialist moral theory is unwise.

Chapter 2: Rule Consequentialism's Dilemma: Demandingness or Self-Defeat

1. What is Rule Consequentialism?

The standard version of consequentialist moral theory, act consequentialism, holds that an act is to be evaluated on the basis of whether it produces the best possible state of affairs, impartially considered. If it does produce the best possible state of affairs, then the act is counted as morally required (or morally permissible if the act is one of many possible acts that produce a state of affairs with as much good results as any other); if the act fails to produce the best possible state of affairs, then the act is counted as morally wrong. One of the standard objections to act consequentialism, of course, is the Demandingness Objection. It is easy enough to see why this is the case. Consider Tim Mulgan's case of Affluent:

Affluent's Tale. Affluent is, as her name implies, an affluent citizen of a developed country, sitting at her desk with a checkbook in one hand and two pamphlets in the other. One of the pamphlets is from a reputable international aid organization. The other is from a local theater company, advertising their latest production. Affluent has enough money to make a donation to the charity, or to buy tickets to the theater, but not enough money to do both. Because of her love for the theater, she buys the tickets, even though she knows the money would have done much more good in the hands of the charity.⁵¹

Surely, the convinced act consequentialist will say, it is the case that giving aid organizations that help people suffering from food shortages or extreme poverty, or lack of sufficient medical resources, or housing, and so forth, would bring about the greatest amount of good in the world. Affluent's taste for the theater has no doubt generated some

⁵¹ Mulgan 2001, 4. Adapted with minor modifications.

good, but that amount of good pales in comparison to the good she could easily produce for other people by donating her money to charity. Thus, act consequentialism must say that Affluent's action is morally wrong, because it fails to produce the best possible state of affairs, impartially considered. This is widely taken to be an extremely demanding result. Some philosophers have embraced this result.⁵² Others, however, have accepted this objection and departed from standard act consequentialism.

There are many ways of making such a departure. Some ethicists, for example, attempt to derive a more plausible version of consequentialism from the recognition of a special class of actions as *group acts* and the proposal that agents are sometimes allowed to defect from those group acts, while retaining direct consequentialist evaluation of actions.⁵³ Other ethicists reject the thesis that we are always required to produce the best possible state of affairs; we are, on such views, required to produce a state of affairs that is 'good enough.'⁵⁴ But an extremely common way of departing from standard act consequentialism is to shift to an *indirect* way of evaluating outcomes. One common way to indirectly evaluate outcomes is to think that we should not directly evaluate the outcomes of actions; perhaps instead we should evaluate the outcomes produced by the adoption of a set of rules. Such rules would be adopted precisely because they produce the best possible outcome. We would not need to evaluate the results of the actions of people who abide by this "ideal code": it would be enough to evaluate the ideal code of rules itself on consequentialist grounds. This is the thought behind a prominent family of modified indirect consequentialist theories, which are known as *rule consequentialist*

⁵² See, for example, Kagan 1989, Singer 1972, and Unger 1996.

⁵³ Mendola 2006.

⁵⁴ See, for example, Slote 1985.

theories of morality. These theories claim that an act is morally permissible if and only if it is recommended by the set of rules that would, if followed, produce the overall best results, impartially considered.

Rule consequentialism has had several prominent defenders.⁵⁵ The most influential recent formulation of rule consequentialism is due to Brad Hooker.⁵⁶ Because Hooker is the most prominent current defender of rule consequentialism, it is his version of rule consequentialism that I shall focus on here. Hooker formulates rule consequentialism in the following way:

RULE CONSEQUENTIALISM. An act is wrong if and only if it is forbidden by the code of rules whose internalization by the overwhelming majority of everyone everywhere in each new generation has maximum expected value in terms of well-being (with some priority for the worst off). The calculation of a code's expected value includes all costs of getting the code internalized. If in terms of expected value two or more codes are better than the rest but equal to one another, the one closest to conventional morality determines what acts are wrong.⁵⁷

It is worth being clear about two aspects of Hooker's position. First, Hooker is not interested in what happens if the ideal code is internalized, or accepted, by everyone. Hooker stipulates, instead, that the ideal code is that which would produce the best consequences if internalized by the *overwhelming majority* of the population. The rest may not accept the code, he says, because we should not expect every group to

⁵⁵ Richard Brandt, for example, has been an influential figure in the explanation and defense of rule consequentialism. An early version of Brandt's view may be found in Brandt 1963, and a fuller statement of his view may be found in Brandt 1998. John Harsanyi has been another influential defender of rule consequentialism: see Harsanyi 1977.

⁵⁶ Hooker 2000.

⁵⁷ *Ibid.*, 32.

internalize the code; young children, the mentally impaired, and the malevolent sociopaths among us, in particular, should not be expected to internalize the ideal moral code. After admitting the difficulty of identifying what an “overwhelming majority” is, Hooker proposes that we think of an “overwhelming majority” as ninety percent of the population.⁵⁸

As for the costs of internalizing a code that Hooker refers to in the above-quoted official definition, Hooker has us imagine that the cost of inculcating a moral code is distributed widely, and “family, teachers, and the broader culture” are responsible for inculcating each new generation with the ideal code of rules, just as the situation is in the actual world.⁵⁹ This, Hooker suggests, will render the ideal code far more likely to be adopted by an overwhelming majority, because humans are far more likely to respond positively to the “organic” teaching and internalization of the moral code, where by “organic” I mean “more akin to how conventional morality is currently taught” -- that is, by parents, teachers, other respected authority figures in society, and by observing a culture’s broader values. A more demanding moral code, taught by enlightened elites, Hooker argues, would not find easy acceptance, for they would not simply be more costly to the ordinary person, but the maintenance of a more demanding code of morality would be higher as well.⁶⁰

Finally, Hooker stipulates that closeness to conventional morality will be the tiebreaker between two competing candidates for the ideal moral code, in the event that two such superior candidates emerge that are the equal of each other in generating the

⁵⁸ Ibid., 84.

⁵⁹ Ibid., 79.

⁶⁰ Ibid.

optimal level of well-being. The resulting code should be, overall, quite close to conventional morality, which is a point that rule consequentialists like Hooker often claim is an advantage for their theory - it coheres well with what we already have. Indeed, Hooker even sets forward coherence with conventional morality as a desideratum for a moral theory. He gives a short rationale for this claim. Suppose we are trying to decide between two moral theories. Theory A is internally consistent, coherent with our moral convictions after thorough reflection, identifies a fundamental moral principle that explains our more specific judgments and justifies those judgments impartially, and helps us reach satisfactory answers to difficult moral problems. But Theory A is at odds with conventional morality. Theory B is like Theory A, insofar as Theory B satisfies all of the four qualities of Theory A just given. In addition, Theory B starts from and attempts to preserve attractive general beliefs about morality. For example, let us suppose that Theory B respects the common-sense judgment that morality does not constantly require heroically altruistic sacrifices from its adherents. One may live a morally good life without being a saint, according to Theory B. Theory A, however, claims that the saintly altruistic life is the only way to live a morally good life. For another example, let us suppose that Theory B respects the common-sense judgment that we are sometimes allowed to give special treatment to our kith and kin. Theory A, however, is rigorously impartial and demands that we consider our spouses, partners, siblings, parents, relatives, and friends no more or less than any other human being in the world. In both cases, Theory A is notably alien to our common-sense moral thinking, whereas Theory B is not. So there will be some psychological cost to adopting Theory A over Theory B. If we find, however, that Theory B meets all of Theory A's beneficial criteria and in addition lacks

the adoption costs of Theory A, then Theory Ba has a distinct and decisive advantage. As Hooker would say, Theory B would be “markedly superior.”⁶¹ Theory B would do a much better job “of matching and tying together our moral convictions.”⁶² Hooker painstakingly works to depict rule consequentialism as a theory that not only recommends much of what conventional morality familiarly also recommends to us, but has the additional advantage of providing a foundational rationale and justification for the code of rules that it recommends.⁶³

2. How Hooker Hopes to Avoid the Demandingness Objection

Having sketched Hooker’s view, it is time to consider just how Hooker’s rule consequentialism might avoid the Demandingness Objection. Perhaps it requires much sacrifice from us for the purposes of producing the best overall state of affairs for everyone. The ideal moral code might include a rule requiring us to donate a portion of our income to an effective charity. Perhaps even the ideal moral code will not prevent us from being required, in some cases, to sacrifice a great deal of our time and resources for the benefit of others, such that we are alienated from our own projects or might be required to give up something of moral value.⁶⁴

Hooker’s method of avoiding the Demandingness Objection is to argue for the thesis that rule consequentialism ought not to be based upon *compliance* with the moral rules, but instead upon *acceptance* of the moral rules. Acceptance of the moral rules, in Hooker’s view, consists in having the disposition to comply with the rules under certain

⁶¹ Ibid., 4.

⁶² Ibid., 101.

⁶³ Ibid., 101, footnote 13.

⁶⁴ See Williams 1973 and Singer 1972.

circumstances, to encourage others to comply with the rules, and to be favorably inclined toward those who comply with the rules (and vice versa for those who flout the rules).⁶⁵ Thus, in Hooker's view, the ideal moral code is the code that we could reasonably expect to produce the best overall state of affairs if the overwhelming majority of people adopted it. So if we are choosing between a more demanding code that cannot be expected to be widely adopted and a less-demanding code that is reasonably expected to be widely adopted, rule consequentialism would recommend that we adopt the less demanding code precisely because that is the code that can be reasonably expected to generate the best overall state of affairs. Call this Hooker's move to acceptance rule consequentialism.

This is a persuasive point. Even an advocate of a more demanding view, such as Peter Singer, seems to concede ground to something along the lines of Hooker's rationale for acceptance rule consequentialism. Despite the rather more extreme view he proposes in one of his most famous articles,⁶⁶ Peter Singer adopts a moderate view on what rule ordinary people actually ought to follow. On a Web site related to Singer's charitable work and his recent book, we find the following:

When did you last spend money on something to drink, when drinkable water was available for nothing? If the answer is "within the past week" then you are spending money on luxuries while children die from malnutrition or diseases that we know how to prevent or cure.⁶⁷

The implication here seems to be that spending money on, for example, a Coca-Cola when drinkable water was available for nothing is an action that morality condemns.

⁶⁵ Hooker 2000, 76.

⁶⁶ For example, in Singer 1972.

⁶⁷ Singer 2011a and Singer 2009.

That, of course, is a position in keeping with the strong version of Singer's principle of harm prevention.⁶⁸ Despite this implication, the Web site related to Singer's charitable work seems to concede that, despite what morality really requires of us, it is perhaps better if we did not really act just as morality requires; instead, we ought to adopt a norm that is easier for everyone to comply with, both in terms of personal sacrifice and helping others to feel comfortable with making a sacrifice:

If everyone who can afford to contribute to reducing extreme poverty were to give a modest proportion of their income to effective organizations fighting extreme poverty, the problem could be solved. It wouldn't take a huge sacrifice.

But first we need to change the culture of giving – to make giving to help the needy something that any normal decent person would do. To help bring about this change, we need to be upfront about our giving.⁶⁹

Hooker's move to acceptance rule consequentialism is designed to thwart the Demandingness Objection. Both rule consequentialism and act consequentialism tend to agree, as Hooker observes, in how agents should do their everyday moral thinking; they should apply rules of thumb, rather than seeking to calculate precise consequences.⁷⁰ But Hooker plausibly points out just how intractable the demands of complying with act consequentialism could be. One way in which people would find act consequentialism quite demanding is act consequentialism's insistence on impartial consideration of each person's well-being. This conflicts with our natural impulse to favor our close friends, family members, and perhaps members of our relevant communities. Hooker argues that

⁶⁸ Singer 1972, 241.

⁶⁹ Singer 2011a.

⁷⁰ Ibid.,144.

the cost of internalizing the act consequentialist moral principle of always acting so as to maximize the good would be incredibly, implausibly costly; we would have to be constantly concerned about the “suppression of strong affections and partiality.”⁷¹ The cost of suppressing individual partiality would not only be borne by all, but it would have to be continuously suppressed in each new generation. In a world, however, where a more modestly demanding rule is accepted by all, and in particular one that hews closer to conventional morality, the cost of acceptance is less high. Not only that, but (as the above quotes from the Web site related to Peter Singer’s charitable work indicate) the cost of acceptance might dwindle over time. As more people become acclimated to the idea of regular charitable giving, such giving becomes more widely seen as less of a burdensome sacrifice and more of a normal practice. To this end, Hooker thinks that a progressive rule making a moral duty out of donating at least one percent of your income, up to ten percent of your income for those who are wealthy, will be an excellent candidate for inclusion in the ideal moral code.⁷²

Another point that Hooker makes is that the rule consequentialist need not accept a fundamental commitment to maximize the good. The fundamental commitment, Hooker thinks, needs to be to the notion that rule consequentialism explains and ties together all of our previously held moral convictions, as well as offering us aid in resolving difficult

⁷¹ Hooker 1998, 24.

⁷² Hooker 2000, 163-164. Interestingly, a calculator on Peter Singer’s Web site (Singer 2011b) also uses the rule of one percent of one’s income to calculate how much money one ought to give to charity for people who earn less than \$105,000 USD. The Web site suggests, however, that if the person interested to find out how much she should give earns closer to \$105,000 USD, she should choose to give a larger percentage of her income. Additionally, a number of values larger than \$105,000 USD used with the calculator gave results notably larger than one percent. Entering a value of \$500,000 USD, for example, gave the result that I should donate \$48,500 USD to charity, a figure slightly less than ten percent of the total income originally entered. Up to this point, Singer’s recommendations seem to be largely in step with Hooker’s. Singer’s recommendations for highly wealthy people, however, exceed Hooker’s maximum; entering a value of \$1,000,000 USD into the calculator gives the result that I should donate \$143,450 USD to charity, which is significantly higher than Hooker’s maximum of ten percent.

ethical questions. Rule consequentialism itself must accept a fundamental commitment to select rules by the criteria of whether they produce the best overall state of affairs, but that is not the same thing as an agent being required to choose acts so as to maximize the good. An agent is not required to so act, Hooker thinks; an agent is simply required to abide by the rules that rule consequentialism recommends, and these rules will not be too demanding.

Such are the ways in which Hooker hopes to escape the Demandingness Objection. The ultimate basis for rule consequentialism is our intuitions about what the ideal moral code consists in, and how it would affect the society that it is widely internalized in.

3. Being Careful about Hypothetical Behavior and Belief

A couple of ways to argue that Hooker does not actually escape the Demandingness Objection exist. I want to start by discussing a way that I do not favor, and then proceed to a more promising way of criticizing rule consequentialism on this point.

I do not favor a line of criticism that seeks to rule out rule consequentialism on the grounds that it might require significant sacrifices in the actual world. I venture to say that any view worth calling a plausible view about morality would sometimes require significant sacrifices in the actual world. Virtue ethics, for example, would say that I have failed to develop the best kind of character, and it would say that I have done something morally blameworthy, if I refuse to wade in to the pond and save the drowning child. If I refuse, I have perhaps demonstrated a lack of courage, or a lack of appropriate empathy for the victim, or the vice of selfishness (or a combination of several or all of the above vices). Even if there are many drowning children, practical deliberation may require me

to spend a great deal of time wading in and pulling out one child after the other. That may add up to a significant cost to me in terms of my time and resources. But I think there would be something wrong if we criticized virtue ethics in this case by saying that it is too demanding. Some demands are entirely appropriate, and to demonstrate that a moral view sometimes requires significant demands in the actual world is not enough. Let me pursue this line of thought with regard to rule consequentialism.

In the actual world, nothing like overwhelming acceptance of an ideal moral code exists. The best we can hope for is partial acceptance. What does rule consequentialism require for those who accept the ideal moral code, but do so under conditions of less than overwhelming acceptance of the code? As a hint, when discussing a modified version of Peter Singer's famous drowning child case, Hooker argues that rule consequentialism would require a rescuer to save two children, despite the fact that saving one would be that rescuer's fair share (which he has already done).⁷³ Hooker introduces a more general rule later on in his discussion of what the ideal code requires under conditions of non-compliance, which reads as follows: "Over time agents should help those in greater need, especially the worst off, even if the personal sacrifices involved in helping them add up to a significant cost to the agents. The cost to the agents is to be assessed aggregatively, not iteratively."⁷⁴ By this, Hooker means that he does not want agents, for every time that a sacrifice is called for from them, to forget what they had to sacrifice on prior occasions. That would be to assess the personal cost to the agent iteratively. In other words, when an agent assesses the personal cost of sacrifice iteratively, every iteration of a choice of whether to make a sacrifice is considered in isolation from every other iteration of that

⁷³ Hooker 2000, 164-165.

⁷⁴ Ibid., 166.

choice. By contrast, Hooker means to say that agents ought to calculate the cost to themselves aggregatively when they are called upon to make sacrifices. They ought to, in other words, not forget all the previous times that they had to sacrifice; they will need to aggregate all the costs of each instance of a choice of whether to make a sacrifice. Such a way of assessing costs, Hooker thinks, will allow rule consequentialism to avoid a high degree of demandingness.

It is notable that the rule regarding sacrifice under conditions of widespread non-acceptance sometimes permits me to incur significant costs. That seems true because of what Hooker says with regard to the modification of Singer's drowning child case. It isn't clear that there is a way for me to escape such costs. Unlike Samuel Scheffler's Hybrid View and Tim Mulgan's Combined Consequentialism, Hooker's rule consequentialism does not allow me an Agent-Centered Prerogative to discontinue making sacrifices if I do a calculation and discover that the weighted cost to myself is greater than the unweighted cost to others generally.

If that is true, then perhaps aggregative assessment of the costs to affluent people living in the developed world will still permit significant costs to befall those people who accept with the ideal moral code's requirements, even if most people do not accept the code. But that would not be a good way to argue that rule consequentialism fails to avoid the Demandingness Objection. Even if the cost to someone who accepts the ideal code's requirements becomes significant, they might be incurring those costs in an attempt to prevent a catastrophe (about which Hooker has more to say, and which I cover in the next section of this chapter). In general, it seems intuitive to say that preventing a catastrophe is a good justification for a moral demand. And the Demandingness Objection was never

appealing to our intuitions about how much cost we had to bear to prevent a moral catastrophe from happening. The Demandingness Objection, rather, is the worry that living a moral life will be completely incompatible with living the kind of life that I want to lead, combined with the thought that anything worth being called the true theory of morality ought to be compatible with living the kind of life that it would be nice for me to lead. It is no great accomplishment to successfully argue that rule consequentialism is demanding in situations where we must make a sacrifice in order to prevent a catastrophe. If that is rule consequentialism's problem, then it is not *only* rule consequentialism's problem.

A better line of argument exists, however, and this line of argument does not depend on showing that rule consequentialism sometimes makes significant demands in the actual world. This line of reasoning about why rule consequentialism does not avoid the Demandingness Objection undermines one of Hooker's key assumptions. Hooker seems to be assuming that our intuitions about what counts as demanding under conditions of widespread acceptance of the moral code would be stable. In other words: in circumstances where overwhelming numbers of people accepted the ideal code proposed by rule consequentialism, Hooker thinks that those people would count as demanding the same things that we would count as demanding. I argue that there is no good rationale for this conclusion. If that is true, then for all Hooker says about how rule consequentialism is not demanding, rule consequentialism may still make demands that we would count as demanding.

Before I proceed to that argument, we can make an observation in favor of Hooker's assumption. A kind of conservatism in hypothetical theorizing seems like it

could be appropriate here, particularly as applied to the psychology of human beings in imagined conditions. If we want to plausibly think about what would happen in hypothetical circumstances where a certain moral code is widespread, we need to hold fixed certain factors about human psychology, including our notion of what humans consider highly demanding.

Still, we should be cautious of hypothetical theorizing, because our intuitions about what the society that follows the ideal moral code would look like are probably unreliable. For example, it is not clear that the notion of what humans consider highly demanding refers to any kind of universally acknowledged limit on sacrifice prior to which we think demands are permissible but after which we think demands are excessive. Peter Singer ably points this out by quoting Thomas Aquinas on the necessity of giving alms for the poor.⁷⁵ According to Aquinas, God provides material goods in order to satisfy human needs, and the divisions of property recommended by human law ought not to be in conflict with the divine law regarding what is to be done with material goods. Anyone's having a superabundance of goods, Aquinas thinks, is preventing others from using those same goods for the sustenance of their own needs, and thus is in conflict with the purpose of those goods as mandated by the divine law. Perhaps Aquinas might think it would be demanding, in some sense, for people to give away their superabundance of goods. But any such "demanding" sacrifice would likely pale in comparison to the importance of upholding the divine law. Naturally, Singer's goal in quoting Aquinas is not to say that divine law is the basis for his highly demanding results. It is, rather, to say that his conclusions are merely out of step with a modern conception of what is morally demanding. There is, Singer also thinks, no special reason to hold fixed the modern

⁷⁵ Singer 1972, 239.

conception of what is demanding, especially when influential intellectual leaders in the Western world hold very different ideas about what counts as demanding.

To give another example of obligatory charity not drawn from the Western tradition, consider the Islamic practice of *zakat*. *Zakat* is one of the five fundamental practices, or “pillars,” of Islam. According to this practice, a wealthy individual is required, as part of being a faithful Muslim, to donate a minimum amount of their material goods to charity. This practice is done, according to the American Islamic charity Baitulmaal, for many purposes. For example, it is meant to cleanse one’s own personality of unwholesome influences stemming from the hoarding of material goods. Muslims see *zakat* as a way of combating the pernicious effects of “baser instincts of greed, miserliness, and selfishness,” and it promotes attitudes of “generosity, love and care, and mutual help,” both within an individual life and within and across societal institutions and groups of people.⁷⁶ But Muslim scholars do not take “wealthy individual” to mean an extremely rich individual whose fortune runs into the millions of dollars. For the purposes of paying *zakat*, a wealthy individual is one whose personal wealth amount meets the appropriate *nisab*, a minimal level of wealth or its equivalent. The established convention is that a person meets the appropriate *nisab* when their personal wealth exceeds the value of three ounces of gold, an amount that Baitulmaal identifies as being roughly \$2500.⁷⁷ If the person who meets the *nisab* is adult, sane, and a free Muslim, then that person is required to pay *zakat*. Although the Qur’an does not appear to endorse a particular percentage as the required minimum *zakat*,⁷⁸ Islamic scholars routinely

⁷⁶ Baitulmaal 2011.

⁷⁷ Ibid. For this amount, upon visiting Baitulmaal’s Web site, one must click on “Zakat Calculator” in the left-hand navigation column of the Web site.

⁷⁸ Qur’an 2:267-274 (Abdel Haleem 2004, 31), for example, does not specify a percentage.

recommend a minimal *zakat* of 2.5% of one's personal wealth in one lunar year.

Although I am no Islamic scholar, it is worth noting, finally, that the Qur'an is pretty clear that the faithful Muslim must give *zakat* in the proper caring and charitable mindset; i.e., *zakat* ought to be willingly paid as part of the faithful person's duty to God, preferably in private (Qur'an 2:271) and without a public demonstration of the *zakat* payer's generosity or any hurtful words or reticence from the *zakat* payer (Qur'an 2:262-266).

I reference Singer's quotation of Aquinas and the Islamic practice of *zakat* to provide evidence for this claim: people who have internalized the ideal moral code may not see abandonment of their own projects in favor of promoting the general good as a great cost. What we consider to be unreasonable demands may seem to people living under conditions of widespread ideal moral code acceptance to be quite reasonable. Hooker believes that promotion of a highly demanding moral code would necessitate a highly costly and demanding education process, consuming a very high degree of time, energy, and resources. Naturally, that is possible. But we cannot know how education processes might change once we have actually reached conditions of widespread ideal moral code acceptance. Richard Arneson presses a similar point against Hooker: given future states of technological progress, we cannot be sure that some new development that renders moral education extremely easy and efficient, thus having a dramatic effect on the rule selection process that rule consequentialism carries out. In such situations, rule consequentialism would be likely to select rules that would appear very demanding to us, but not at all demanding to future people with access to advanced technology.⁷⁹

⁷⁹ Arneson 2005, 248-249. Hooker admits that he is unable to escape Arneson's counter-example without relativizing rule consequentialism to a present state of technology: see Hooker 2005, 268-269.

Whether we appeal to future state of technology, as Arneson does, or we appeal as I do to the ability of human beings, past and present, to revise their views on what is considered demanding, the reliance on what a hypothetical society does creates epistemic difficulties for Hooker. This is because Hooker's rule consequentialism makes the moral rightness of actions dependent on contingent and culturally specific facts about the society that lives under conditions of overwhelming ideal moral code acceptance. We simply cannot be sure that promoting widespread acceptance of the ideal code would have the result that everyone contributes a manageable amount of their resources and no one has to give up non-optimal projects that they care about. It seems equally possible that the society living under those conditions might be quite happy to give up non-optimal projects that they care about when morality requires it. Perhaps, as Singer thinks, more morally demanding practices may be adopted as the cost of inculcating them organically drops, much like the market price of new goods and services tends to fall over time.

Or perhaps people living under conditions of overwhelming ideal moral code acceptance would adopt more demanding moral codes when an effortless method of moral teaching becomes available. Perhaps a method of directly implanting the teachings of the most rigorously altruistic moral code known into a person's brain becomes available for widespread use. Call this rigorous code that I have just described the altruistic code, and call its competitor (the previously-accepted ideal moral code) the normal code. If this lowers the cost of the altruistic code's acceptance enough for the altruistic code to overtake the normal code's maximum expected value in terms of well-

being, then the good rule consequentialist society would have to look favorably upon the more-demanding altruistic code.

Hooker might object that this kind of direct moral code implantation technology, whatever form it takes — whether it looks like the “braincaps” from Arthur C. Clarke’s science-fiction novel *3001: The Final Odyssey*, the seemingly more malicious back-of-the-head socket familiar to us from *The Matrix*, or something else entirely — would be scary and unfamiliar enough to *increase* the cost of the altruistic code’s inculcation, not to decrease it, as I have suggested. It would alter the nature of a fundamental piece of parenting and teaching, and people would be very suspicious of such a technology. That’s a plausible objection, and I need a response. The answer is this: we do not have a compelling reason to think that people would *always* react negatively to such a cultural innovation. I must qualify this answer right away by saying that people very well may *initially* react negatively, because of its unfamiliarity and the perceived threat that it poses to standard ways of moral teaching. But given enough time, effective lobbying, and demonstration of benefits from real-world use cases, people may very well be won over.

Why is this result problematic for Hooker? Recall that Hooker wants to use the set of rules that, if overwhelmingly adopted by the vast majority of people, would produce the greatest expected value as what determines moral rightness and wrongness. But a basic difficulty in pinning down just what that set of rules is exists, because acceptance costs — or, more simply, our notions of what we think is too demanding — are subject to change. Hooker seems to rely on the thought that acceptance costs do not fluctuate very much, if at all. I do not think a compelling case can be made for this point. The point of this thought experiment is just to show that hypothetical conditions, like conditions of

overwhelming ideal moral code acceptance, are far removed enough from the actual world that our intuitions about this case are probably not reliable. We do not know what kinds of technological advances might make altruistic moral teachings more popular and effective, or how even relatively modest cultural projects (like Singer's anti-poverty program) might lead people to view what we would currently consider to be the very demanding option as less of a sacrifice. If all that is plausible, then the upshot is that for all Hooker has claimed about the moderation of rule consequentialism, it might still recommend something that we would take to be very demanding.

4. Rule Consequentialism and Polluter's Dilemmas

I want to consider, in this section, whether Hooker's rule consequentialism fares better when we consider Many-Person Dilemmas that involve small harms that accumulate over time and result in very large harms. In other words, I wish to investigate whether rule consequentialism is subject to the Polluter's Dilemma, which I developed in the previous chapter. I shall claim that rule consequentialism does not fare well in this regard. Before I proceed to argue for this claim, I want to take into account two more aspects of Hooker's formulation of rule consequentialism that are important.

One aspect requires me to be careful in formulating my objection. Hooker claims that much of rule consequentialism is amenable to reformulation so as to include protections for the environment. He does not elaborate, however, on what such reformulations might look like, but he is reasonably confident that such reformulations

can be made.⁸⁰ We can make some attempt to discern what rules rule consequentialism would recommend, and I will attempt to do so below.

Another aspect of rule consequentialism worth mentioning here is that it would endorse a rule that enjoins us to prevent disasters. Ordinarily, rule consequentialism does not allow deviation from rules, such as “avoid lying,” when doing so might allow us to gain a little extra utility on the side. Similarly, it does not recommend fine-grained exceptions to rules, even when such exceptions are apt to generate more utility. This feature of rule consequentialism, Hooker argues, enables him to avoid the Collapse Objection to rule consequentialism. This objection, roughly sketched, holds that rule consequentialism is extensionally equivalent to act consequentialism, but weaker because of its added theoretical machinery, including the indirect assessment of acts via its assessment of rules. Hooker, resisting the Collapse Objection (plausibly, in my view), writes:

How much confidence would you have in others if you knew they accepted such highly qualified rules? Mackie...observed, ‘We are rightly sceptical about a man of principle who has a new principle for every case.’ The same is true about someone who has too many exception clauses.⁸¹

So there is some reason to think that accepting such a highly complex set of rules would not engender trust regarding other people in us. Hooker resists the Collapse Objection further, stating that there is more reason to suppose that rule consequentialism would not accept a code of rules with so many epicycles. It would mean a higher cost of

⁸⁰ Hooker 2000, 70-71.

⁸¹ *Ibid.*, 96.

internalization. A better code of rules, Hooker thinks, would be the set that is limited in number and complexity.⁸²

But rule consequentialism will allow deviation from rules sometimes. Plausibly, I might not want to insist on scrupulously observing the “avoid lying” rule when an act of lying to the secret police might enable me to save the lives of innocent refugees I am hiding in my basement. Thus, Hooker argues that the ideal code of rules will include a rule that tells us to prevent disasters, and this rule will have the authority to override other rules in the ideal code.⁸³ Most obviously, the rule would apply to situations where we can prevent imminent, acute disasters, such as the drowning of a child. It would be better, all things considered, to break my promise to always get to work on time and save a drowning child rather than diligently follow the rule “keep the promises that you make.”

Taken together, these two aspects of rule consequentialism can generate rules that, at least at first glance, render it better prepared to face the challenge of the Polluter’s Dilemma. Rule consequentialism would recommend that we act according to rules that require us to prevent disaster. Presumably, this recommendation would include preventing environmental disasters. It does not allow agents to prefer their own non-optimal projects at the expense of the general good. And it does not allow agents to defect from a cooperative project to gain a little extra utility for themselves when the opportunity presents itself. So it would, we are let to believe, prevent disasters related to climate change.

I argue that this appearance is mistaken. Hooker’s rule consequentialism is subject to a dilemma. On one horn of the dilemma, rule consequentialism is directly collectively

⁸² Ibid.

⁸³ Ibid., 98-99.

self-defeating.⁸⁴ On the other horn of the dilemma, rule consequentialism endorses a set of rules with extremely high acceptance and maintenance costs that will alienate people from their projects and demand more of them than we might ordinarily have thought that morality demanded. And this endorsement may be necessary in the actual world.

The “prevent disaster” rule will not allow the society that accepts the ideal code to prevent the generation of Polluter’s Dilemmas. The “prevent disaster” rule seems designed, as I suggested above, to prevent imminent, acute disasters. Polluter’s Dilemmas are not disasters of this type. Polluter’s Dilemmas are generated by repeated instances of the same type of action, each involving a diffuse harm that, considered on its own, is negligible in its effect. But the repeated instances of the harm accumulate over time, and eventually a very great harm that applies to all individuals within the relevant ecosystem is generated by the accumulation of small harms. Polluter’s Dilemmas are disasters that are slow to occur and that inflict their great general harm over a long period of time, instead of being a disaster that occurs all at once and that might be prevented by an immediate action licensed by the “prevent disaster” rule. To prevent the generation of Polluter’s Dilemmas, there would need to be a rule in the ideal code that prevents us from creating the sort of small, accumulating harms that generate the Polluter’s Dilemma. But such a rule, by Hooker’s stipulation, is unlikely to be included in the ideal code. This rule would, at the very least, require us to think far more carefully about whether the benefit we derive from each individual car trip is worth the harm to the environment that the car trip involves, for example. It may require us to be far more scrupulous about taking care to properly dispose of waste and conserve water. This rule would be difficult to internalize and costly to maintain. People in wealthy Western nations are quite used to

⁸⁴ See Parfit 1984 for this restriction on acceptable moral theories.

being able to go to whatever places they wish, when they wish, if they have enough gas in their fuel tank. And they do so for many reasons, some of them worthwhile (a commute to their job, for example), but many of them unnecessary. A rule that abrogated this privilege would be difficult to accept. Remember Hooker's move to acceptance rule consequentialism. A rule that is likely to involve extremely high acceptance and maintenance costs is not a likely candidate for inclusion in the ideal code. In particular, a rule designed to prevent the generation of Polluter's Dilemmas would be far-reaching enough that it would undoubtedly greatly affect the total cost of acceptance and maintenance of the ideal moral code. Rule consequentialism would thus reject such a rule as too demanding. But perhaps rule consequentialism's "prevent disaster" rule would enable the society living under the ideal moral code to maintain a consequentially optimal level of pollution: more restriction and the code becomes too stringent to maintain, to say nothing of its dampening effects on economic activity; less restriction and the pain from climate change becomes too acute for us to readily adapt to. Suppose us to have such a story about how we go about generating this consequentially optimal level of pollution.

It is too optimistic, however, to suggest that rule consequentialism can ultimately avoid generating Polluter's Dilemmas while at the same time avoiding the Demandingness Objection. The point is this: no individual car trip is a disaster. The harm produced by a single car trip is negligible. It is only in the accumulation of these individually negligible harms that the harm of pollution reaches disaster proportions.⁸⁵ So there does not seem like any principled way for a rule consequentialist to apply the "prevent disaster" rule in a society where people are used to the freedom of movement

⁸⁵ Here I am ignoring the more typical sorts of disasters -- that is, fast-moving disasters with acute effects -- that involve pollution, such as accidents involving spills of crude oil, spills of industrial by-products, or even radioactive contamination of a region due to a nuclear meltdown.

that cars provide them. But let us suppose for a moment that our supercomputers which model the effects of pollution on the environment tell us that, sometime in the next week, our industrial and commercial activity will cause us to reach a “tipping point” in the process of climate change, such that reaching this tipping point will put us far enough above the consequentially optimal level of pollution to cause the process of climate change to accelerate beyond our capacity to adapt to it quickly enough. It’s unclear, as I have just argued, whether there is a principled rule consequentialist rationale for applying the “prevent disaster” rule in an effective way in “tipping-point” circumstances. But let us suppose that we ignore this problem and apply the “prevent disaster” rule effectively. We do so, out of necessity, on a very widespread basis, and so many ordinary rules are overridden. Non-essential car travel is heavily restricted, as is air transportation. Polluting industry is curtailed, including the generation of electrical power at coal-fired power plants. Economic activity suffers heavily. This is an extremely demanding result. And this demanding result may not just be in a possible world that has reached a climate tipping point. This may be the situation in the actual world. The scientific consensus has it that our world is either very near to, or has recently passed, a tipping point of this sort. If it is the situation in the actual world, Hooker’s rule consequentialism must either endorse an extremely demanding result in the actual world, thus rendering it unable to avoid the Demandingness Objection, or it must permit the generation of Polluter’s Dilemmas, in which case it is directly collectively self-defeating.

Before I conclude, let me consider an objection to this line of argument. The objection is this: a general directive enjoining us to be prudent is likely to be included in the ideal moral code. It is part of conventional morality, and it is furthermore a rule that

has high payoff in terms of expected well-being and low inculcation costs. Even though people sometimes have difficulty being prudent, inculcating a general practice of prudence is not very demanding because prudence appeals to an agent's self-interest. The cases I have pointed to in developing the Polluter's Dilemma could be headed off by a general rule of being prudent. Prudence might make significant demands in some circumstances, but those demands are always to my benefit. And if those demands are always to my benefit, then they cannot be the all-consuming, alienating kind of demands that motivate Bernard Williams's excellent formulation of the Demandingness Objection.⁸⁶ So I have failed to come up with a case demonstrating that Hooker has not avoided the Demandingness Objection.

This objection cannot shield Hooker entirely from the force of the Polluter's Dilemma. It may be prudent for the current generation to work together to avoid the consequences to us of global climate change. But if that is how we understand prudence, then prudence has nothing to say about why we ought to prevent a future disaster that would not affect us at all. Let us stipulate, in the thought experiment involving a consequentially optimal level of pollution given above, that the disasters resulting from the runaway effects of climate change shall not take effect until every currently living person is dead. In the very moment after the last currently living person dies, however, global climate change will begin to take its toll. In this situation, appeals to prudence will have no weight; no one's interests will be affected negatively due to climate change. But there will be future generations, and we will be in a position to prevent a disaster for them by regulating our own behavior and our collective appetite. So the "prevent disaster" rule will still be in effect, and will govern many actions that we take that have important

⁸⁶ In Williams 1973.

implications for future generations. Thus, rule consequentialism may require me to give up projects that I care about, and for reasons that have nothing to do with my own self-interest.⁸⁷

I conclude that there is no reason to suppose that rule consequentialism successfully escapes the Demandingness Objection. Generation of Polluter's Dilemmas is clearly unacceptable. And yet there is a very real danger that rule consequentialism will either generate Polluter's Dilemmas, or assign just the kind of alienating demands that the friend of the Demandingness Objection always thought was unacceptable.

Joseph Mendola observes, along with Derek Parfit, that common-sense morality probably is not suitable for the technologically advanced age we find ourselves in.⁸⁸ This is an acute problem for rule consequentialism, as it seeks to hew very closely to common-sense morality, even to the point where closeness to common-sense morality is the tie-breaker between ideal-code candidates. Coherence with common-sense morality is obviously important to some degree, and morality had better deliver familiar and expected results, such as the wrongness of murder, theft, and rape. But closeness to common-sense morality is clearly deficient when we consider large, nasty problems such as climate change. So much the worse for moral theories that attempt to stick close to common-sense morality.

⁸⁷ Thanks to David Sobel for this objection.

⁸⁸ Mendola 2006, 308.

Chapter 3: The Inconsistency of the Collective Principle

1. The Collective Principle and the Compliance Condition

There are several kinds of tactics that a philosopher interested in defending a version of consequentialism from the Demandingness Objection might use. Some of these tactics we have discussed already; namely, the Agent-Centered Prerogative and the reliance on intuitions about the cost of accepting and maintaining a principle of beneficence. In this chapter, I want to consider a third attempt to moderate consequentialism's demands enough to overcome the Demandingness Objection; that of replacing consequentialism's standard optimizing principle of beneficence with a principle of beneficence that incorporates the concept of a fair share of doing good that an agent is required to do. This is Liam Murphy's view. Murphy has done much in recent years to bring into focus issues of fairness in how consequentialism distributes demands among agents. His goal is to propose and defend what he believes is the most plausible candidate for a principle of beneficence: one that incorporates the concept of fairness.⁸⁹ In what follows, I present two arguments against Murphy's view. First, I argue that Murphy's view has been moderated so much that it will regularly fail to make appropriate, plausible demands of agents. Second, I argue that Murphy's version of moderate consequentialism cannot escape the Demandingness Objection in the actual world, because it recommends that each agent's fair share of responding to the problem of global climate change is extremely high. In the final section of this chapter, I explore a

⁸⁹ See Murphy 1993, 1997, 2000.

modification to Murphy's view that can enable him to avoid both objections.

Murphy's perspective on beneficence represents a departure from the standard act consequentialist conception of beneficence. Whereas the act consequentialist sees beneficence as a project for us to individually engage in, Murphy thinks it is better to see beneficence as "a *shared cooperative aim*....Each of us does not, strictly speaking, aim to promote the good. Each sees himself as *working with others* to promote the good."⁹⁰ Such a thought is appealing enough, of course. Suppose that an ordinary middle-class citizen of a wealthy nation acted on Peter Singer's moderate principle of beneficence and donated to an effective famine relief effort up to the point where she would have needed to give up something of moral significance to donate further.⁹¹ Although it is a large donation considered on its own, such a donation does not go very far. Upon realizing this, our donor might be filled with pessimism upon considering act consequentialism's individual requirement to maximize the good. At that point, our hypothetical agent would realize that a single charitable donation, despite it being a very large donation, and thus, a maximization of the good, her donation will not have done much to solve such an overwhelming problem. What might make the donor less pessimistic is the realization that every little bit does help, because she is part of a collective project to provide relief. This thought experiment, if plausible, shows that there is at least some reason to think that adopting a more moderate, collectively-distributed take on the requirements of beneficence might actually produce better consequences in the long run. If I think that all I need to do is my fair share of what needs to be done, and my fair share is manageable, I will be much more willing to do my part than if I am faced with a seemingly unlimited

⁹⁰ Murphy 1993, 285-286; emphasis in original.

⁹¹ Singer 1972.

demand that I promote the good wherever I am able in the course of solving an enormous problem.

With this departure from the standard consequentialist conception of beneficence, Murphy claims that the traditional principle of beneficence, embraced by the classical utilitarians and act consequentialists, is not the most plausible candidate for a principle of beneficence. This traditional principle he refers to as the *Optimizing Principle*. The Optimizing Principle's content is familiar. It holds that we are always required to act in the way that produces the greatest overall benefit. Murphy believes, however, that the Optimizing Principle is indefensible because of its unfairness to agents who comply with its demands.

What does this unfairness consist in? By now, most people are well aware that there are millions of desperately poor and needy people all over the world. Most people are also aware that numerous charities exist which work to alleviate the conditions of the desperately poor, and have had much success in doing so. And finally, most people recognize the fact that their contribution could help alleviate starvation, disease, or death. Recognizing this, some donate their time and/or money to charities. But these people, sadly, are most likely in the minority. While it seems plausible that most people have donated money to charity or volunteered for a charitable cause at some point, most people do not sacrifice a significant portion of their income for charitable purposes. Those people who are charitable donators know this. If their operative moral principle is the Optimizing Principle, then it seems like they are morally required to sacrifice as much of their time, energy, and resources as possible, at least up to the point where the charitable donator would be sacrificing something of moral significance (perhaps, for

example, her ability to pay for her own housing). This is because the most optimal action, under such non-ideal conditions, would be for the charitable donors to give as much of their own resources as possible so as to make up for the lack of compliance from the non-donors.

Murphy argues that this reveals the unfairness of the Optimizing Principle; it places disproportionate burdens on those who actually decide to comply with the demands of morality, so conceived. The donors, he suggests, might ask why they have to take on the additional burden of making up for the moral wrongdoings of others.⁹² It would be unfair to require me, he argues, to make up for the wrongdoings of others; this would allow others to exploit me.

It is for this reason that Murphy argues for the rejection of the Optimizing Principle, and the adoption of his *Collective Principle of Beneficence*. The Collective Principle, roughly stated, holds that we are required to perform the action that could reasonably be expected to produce the greatest overall benefit. We are only required to sacrifice, however, to the extent that we would be required to sacrifice under full compliance; that is, if everyone was complying with the demands of morality. Murphy's statement of the Collective Principle runs as follows:

Everyone is required to perform one of the actions that, of those available to her, is optimal in respect of expected aggregate weighted well-being, except in situations of partial compliance with this principle. In situations of partial compliance, a person's maximum level of required sacrifice is that which will reduce her level of expected well-being to the level it would be, all other aspects of her situation remaining the same, if there were to be full compliance from that point on. Under

⁹² Murphy 2000, 76.

partial compliance a person is required to perform either an action, of those requiring no more than the maximum level of required sacrifice, that is optimal in respect of expected weighted aggregate well-being, or any other action which is at least as good in respect of expected weighted aggregate well-being.

However, no one is required to act in a way that imposes a loss on some other person unless that other persons' expected level of well-being after the loss would be at least as high as it would be, all other aspects of the situation remaining the same, under full compliance from that point on.⁹³

Note that the crucial component of the Collective Principle is what Murphy has called the *Compliance Condition*. The Compliance Condition may be stated thus:

If other agents cease to comply with the demands of a principle of beneficence, an acceptable principle of beneficence will not increase its demands on agents who do comply with that principle.⁹⁴

Thus, if I am a member of a group which is collectively engaged in a beneficent project, such as providing aid to the victims of a natural disaster, the Compliance Condition protects me from being exploited in the unhappy circumstance where other members of the group defect from the beneficent project, whatever it is, and I am left to pick up the slack.

The way the Compliance Condition and the Collective Principle work together would look something like the following situation. Suppose that an earthquake has caused exactly one hundred billion dollars in damage to a city in a foreign country, and I

⁹³ Murphy 2000, 117-118.

⁹⁴ Murphy 1993, 278.

am a citizen of a country untouched by this disaster. These two nations -- my country and the foreign country -- are the only two nations in the world. The earthquake has destroyed the headquarters of the foreign country's stock exchange and major banks, so none of the foreign country's citizens can make relief donations within their own country. My fellow citizens and I think that it is morally obligatory to provide emergency aid to those living in the disaster zone, but as it happens a recent economic crisis is preventing the government from being able to do much but attend to its budget woes. That being the case, we citizens shall need to take matters into our own hands by figuring out just what our individual obligations are. Suppose now that my country has one hundred million citizens able to give aid. Full compliance with the Collective Principle would generate a grassroots, nationwide disaster relief program, such that every one of my country's citizens would give one thousand dollars to the stricken foreign country. But not everyone donates money to the relief effort. The Compliance Condition, however, prevents me from being morally required to donate any more than one thousand dollars in aid, since that is what my duty would be under conditions of full compliance. I would thus not be made to sacrifice more when non-cooperative agents sacrifice less. It is their moral responsibility to donate more, and we would be able to hold those non-cooperative agents morally accountable for their lack of willingness to donate.

2. Does the Collective Principle Fairly Evaluate Rescue Cases?

Toward the end of his *Moral Demands in Nonideal Theory*, Murphy attempts to address James Rachels' two-rescuer case.⁹⁵ Rachels says that this shows "the fallacy of

⁹⁵ Rachels 1979, 162-163.

supposing that one's duty is only to do one's fair share [...]."⁹⁶ The two-rescuer case proceeds as follows. Suppose you and a perfect stranger are walking alongside a pond, and you come across two children drowning in a pond. Obviously, each person's fair share under conditions of full compliance with the Collective Principle is to save one child apiece. You wade in and do your fair share, saving one child. But the stranger simply walks away. Are you required to save the remaining child?

Murphy's answer is complex. Obviously, the Collective Principle does not directly require me to save the remaining child, because to require me to wade back into the pond would be to require me to do more than my fair share, and that would violate the Compliance Condition. There is no question that I would be permitted to save the remaining child; agents would always be permitted to sacrifice further if they wished to do so. But that is not the point. It is our goal to figure out exactly what the Collective Principle requires us to do in this situation. And the answer that Murphy gives on this question seems to be that it depends on the situation whether the Collective Principle gives the "right" answer in this case.⁹⁷ By "the right answer," I mean that you would be required to wade back in to the pond to save the other child, even though you have done your fair share.⁹⁸ Murphy says that it follows that a certain level of sacrifice would be required of us in such a two-rescuer case. But it does not follow directly that I should wade back into the pond. The Collective Principle, he writes, will not treat needs generated by emergencies differently than other cases of required beneficence. It is simply a matter of where applying my required level of sacrifice would do the most

⁹⁶ For Murphy's response to Rachels, see Murphy 2000, 127-133.

⁹⁷ Ibid., 128.

⁹⁸ Regular posing of this dilemma to undergraduates in introductory ethics courses seems to support this assumption. I admit, however, that this is anecdotal evidence.

good.⁹⁹ Here, the Collective Principle may come up with the right answer. For consider how much good I would be doing by hauling another child out of the pond. In exchange for the low cost of spending a few more seconds in the water and getting my already-muddy shoes and pants even dirtier, the payoff of saving a life seems for all the world like a cost-effective sacrifice. Consequently, the collective principle of beneficence, Murphy claims, may support a rule of thumb that requires you to perform rescues if the cost of doing so is very low; furthermore, it may increase overall well-being if we cultivate dispositions to perform rescues when they are needed.

But sometimes, the Collective Principle will give the “wrong” answer in a rescue case. Murphy’s example here is a case in which some astronomers have been planning for years an important and expensive experiment, which must be conducted at sea on a specific night. At the requisite time and place, the astronomers get a distress call from a sinking ship that is many miles away. The astronomers could abandon their experiment and set course for the sinking ship to rescue that ship’s crew. But if they abandon their experiment and steam to the rescue, they will lose forever their chance to conduct their experiment. Moreover, if they abandon their experiment, all the money spent in preparation for that experiment would be wasted, leaving not only the astronomers themselves but also the government agencies, private donors, and private research companies worse off, for they would have nothing to show for all the money spent. The Collective Principle, as stated in the previous section, does not require the astronomers to perform the rescue, for doing so would require the astronomers to both impose a loss on those agencies, companies, and donors and be unable to compensate them enough to make up for the loss. In other words, it violates what Murphy refers to as the *Third-*

⁹⁹ Ibid.

Person Rider.¹⁰⁰ What the astronomers might be required to do is figure out what level of sacrifice they might have needed to make to carry out the rescue, and factor that in to their general duty of beneficence required by the Collective Principle. They would, in other words, have to make a donation “in memory of those lost at sea.” But it would be permissible to let the sailors drown.

This makes a hard kind of sense. At some point, heroic rescue efforts, whatever the situation, become fruitless and too costly to continue. Doctors sometimes cease resuscitation efforts when continuing could save a life. Searches for missing hikers or for missing persons buried in disaster debris are called off eventually, even when it is theoretically possible that the unfortunate ones might yet continue to live. Common sense might seem to adhere to an unlimited duty of rescue at times, but that is not necessarily a point in favor of common sense. Murphy’s argument that not all rescues are cheap or easy, and therefore a highly demanding duty of rescue might force the rescuer to exceed his fair share of sacrifice, may go some way toward alleviating the demandingness of the need to perform rescues.

But those who hold that the Demandingness Objection has some force could object that Murphy has failed to show that his account spreads the demands of morality evenly. For example, Murphy admits that extremely destitute people are not required to perform a rescue at all, even if it is quite easy to do. This is because they have so little to sacrifice that *any* level of sacrifice might exceed their fair share.¹⁰¹ But surely, Murphy’s critic might say, even a very poor and destitute person could be expected to wade into the pond and save the drowning child.

¹⁰⁰ Ibid., 121-122.

¹⁰¹ Ibid., 131-133.

Consider the following case. Suppose you are homeless, and have only a single dollar in change in your pocket and the clothes you are wearing as resources. You typically reside at a homeless shelter at a church, but in your daily wanderings you have crossed to the other side of town, relative to the shelter. Suppose then that you come across another homeless person, who has been walking for some time toward this shelter from a town many miles away. This person is extremely tired and has not had a meal in several days. He has eighty cents, just twenty cents too few to buy a package of trail mix from a vending machine, and he asks you if you could give him twenty cents. You have twenty cents, and you conclude from his appearance that he will probably collapse and die from starvation if he does not have this package of trail mix right away.

Given what Murphy has said, the Collective Principle seems to not require you to give this other homeless person twenty cents. But why is that? The Collective Principle, all other things being equal, essentially says that the maximum level of sacrifice is equivalent to the total amount of good to be done divided by the number of agents available to produce the good. Since the total amount of good to be done is equivalent to twenty cents in this case, and you are the only agent available to do the good, it follows that you are required to give twenty cents. It isn't clear why you *wouldn't* be required to give twenty cents in this case. Even if twenty cents is equivalent to twenty percent of one's total financial resources, twenty cents is still a relatively negligible amount of money because its purchasing power is so low. If you lived in a nation where the purchasing power of twenty cents was very high, then to demand a twenty-cent sacrifice from you would perhaps seem extreme. But the impact of twenty cents upon you in this

situation is negligible, and your sacrifice could save a life. A strict emphasis on fairness here thus does not seem appropriate.

What seems more appropriate to consider is the position you are in are to bring about something good, and how efficiently you could produce that good. In this case, you are uniquely situation to bring about some good, and because of the relative purchasing power of twenty cents, you are also in a position to efficiently bring about that good. That is to say, a relatively small sacrifice on your part would bring about a great deal of good. We would say the same of Superman, who can very easily rescue a great many people. Superman, with his incredible abilities, is uniquely situated to bring about a great deal of good, and has the ability to efficiently bring about that good. He could easily perform rescues and bring about good far beyond his fair share, as Murphy would understand that notion.

Imagine, for a moment, that Superman tells us that he is not morally required to save any more people today because he has done his fair share of good. Metropolis is threatened by numerous calamities, and Superman has done his part to save the people of Metropolis. But the rest of the Justice League, uncharacteristically, has declined to do their part. Superman is not interested in picking up the slack, though he could easily do so. Intuitively, our reaction would be puzzlement. What explains this puzzlement? Just the considerations I have listed. I suggest that one's fair share is not determined simply by the amount of good to be done divided by the number of agents available to produce the good. One's fair share is also determined, although perhaps not completely, by how well someone is situated to do good, and how easily they could pull off their contribution to the good. Although it is unfair, in Murphy's sense, for us to ask Superman to pick up

the slack from the rest of the Justice League, I suggest that it does not disrespect common sense for us to ask for more sacrifice from Superman on this point. It would be fair to ask this of him, we think, just because he is so uniquely situated to do good and so efficient at bringing it about. For the same reasons, the general public would be correct to castigate me if I were to refuse to save the second child in Rachels' two-rescuer case.

Murphy thinks that our disgust with you, should we find out that you allowed someone else to suffer because you think that saving him or her would have violated the Compliance Condition, consists in not wanting to have anything to do with you rather than an especially strong judgment of wrongdoing.¹⁰² That holds true, it seems, whether you are the rescuer in Rachels' case who has already complied with morality and refuses to save the second child, or the unwilling Superman, or a homeless person with a dollar in change. I do not believe Murphy's result is correct. We can criticize you for having an incomplete understanding of what fairness consists in. Your wrongdoing in the situation of the homeless person, for example, stems from not taking important contingent facts into account, namely the difference between the good that the twenty cents would do you and the good that the twenty cents would do the other homeless person. It also stems from not incorporating that contingent information into an assessment of how efficiently you could perform the good to be done. Now, Murphy's Collective Principle requires you to know a lot of contingent information before you can know what your fair share is; it requires you to know, for example, how much good needs to be done and how many people are available to do the good. But there is more to determining a fair share than that. In this case, you have failed to be fair in appraising how much twenty cents is worth to you and your homeless counterpart, respectively. You have thus failed to be

¹⁰² Ibid., 133.

completely impartial in your judgment about where the place is that your resources would do the most good. And that means you have failed to fully capture the concept of fairness in how your moral theory assigns demands. Having a fair principle of beneficence isn't just a matter of having a principle that allows one to figure out fair shares. A fully fair principle of beneficence will need to take into account contingent information like who is uniquely situated to produce the good and how efficient they are at producing that good. That principle of beneficence would respect our common-sense notion that greater obligation to do good goes along with greater ability to do good. It seems perverse to consider just what your fair share, as Murphy understands that notion, might be in situation like this. This is true not just because we watch what you do and make a negative judgment about your character.

Up to this point, I have been concerned with what the Collective Principle requires of a destitute person. Even if my result in this discussion is correct, one might think that this was a case of going after low-hanging fruit. My case might cause trouble for Murphy, but the trouble I have caused so far is trouble on the fringes. It's not clear why demands morality might make of a destitute person are important when the paradigm case of a person that the Demandingness Objection seeks to defend from the overwhelming demands of consequentialism is an affluent person seeking to not be alienated from the kinds of projects and pursuits that we might have thought a moral life contained room for.

I think, however, that what we have learned from considering the case of what demands might be appropriate for a destitute person can be generalized from the case of a homeless person to the case of an affluent middle-class citizen of a developed nation.

Suppose that we select any member of the class of affluent middle-class citizens in developed nations. Suppose further that whatever the required fair share that Murphy's Collective Principle generates for her is, she has complied with it. Let's stipulate that ten percent of her money has gone to effective charitable organizations that relieve acute suffering and famine in war-torn, far-flung places. In every sense, she has effectively and faithfully followed through on her moral obligations, as Murphy understands them. Now she is out for a walk and comes across a starved homeless person seeking twenty cents to buy a bag of trail mix from a vending machine. As before, the appearance of this homeless person suggests that he will collapse and die of starvation if he does not have a bag of trail mix immediately. But our affluent citizen, having some leisure time to read philosophy, knows that the Collective Principle says that she has met her moral obligations to alleviate famine in the world and that she is not required to give a single cent more, much less twenty more. So she walks away, and allows the homeless person to die of starvation. Whatever Murphy says about the necessity of having a solid motive to render help in cases of extreme need, the Collective Principle does not require our affluent citizen to render aid in this case. This is, I think, a *reductio ad absurdum* on the Collective Principle.

It does matter whether our affluent citizen in this case is uniquely situated to bring about the good, and how efficient she is at bringing it about. As I argued previously, fairness in being beneficent is not just a matter of knowing our fair shares, abstractly considered. Knowing how uniquely situated we are to do the good, and how efficiently we are able to do the good, is important for us to figuring out what a fair demand on us is. If our affluent citizen here is not uniquely situated to bring about the good, then her

wrongdoing is greatly reduced. Perhaps other people are nearby, and they have not done their fair share of alleviating poverty and homelessness. If she refuses to give twenty cents, she has still done something wrong, because she would be at least as efficiently as any of the other affluent onlookers at bringing about the good. If, however, she is not as efficient as some of the onlookers would be at bringing about the good, then she is absolved of wrongdoing. If a reduction in her financial resources of twenty cents leaves her unable to afford shelter for herself, for example, then she would not be efficient at bringing about the good to be done. Then we could excuse her for asking one of the onlookers to start doing their fair share.

What if our affluent citizen was uniquely situated to save the homeless person's life but was not efficient at bringing that about? Suppose now that there are no onlookers, but giving twenty cents would leave her unable to afford shelter for herself. It would, I think, still be wrong if she were to refuse to give twenty cents. The loss of a life seems weightier than a temporary inability to pay a bill. It would, moreover, not be much of a burden on someone else if our affluent citizen were to ask for assistance totaling twenty cents so that she could pay her bill.

Keith Horton has recently come to a similar conclusion when assessing Murphy's view about what fairness consists in.¹⁰³ He argues that considerations of fairness are stronger in some contexts and weaker in others, depending on how strong opposing factors are (such as when a life is at stake). I agree with Horton's assessment, but I think we can say more about why the Collective Principle has failed to produce an acceptable answer in what Horton calls "low-cost-high-benefit cases." The Collective Principle permits this absurd result because it is concerned only with what one's fair share is,

¹⁰³ Horton 2011, 91-92.

abstractly considered. It would not permit this result if it required us to fairly evaluate the benefit to an affluent citizen holding on to their resources against the benefit to a destitute person of an affluent citizen making a sacrifice of their resources and make decisions on that basis. That, for Murphy, might be a step too far in the direction of the Optimizing Principle. But it's possible that such a modification might be keeping in the original spirit of the Collective Principle, as the modification attempts to flesh out the concept of fairness that Murphy is relying upon.

As we have seen, the Collective Principle has been moderated so much that it will regularly fail to make appropriate, plausible demands. A modification to the Compliance Condition may make things better for Murphy. I will consider this modification in a subsequent section of this chapter. But this is not the only trouble that the Collective Principle faces.

3. Mulgan's Wrong Facts Objection and the Problems of Global Climate Change

Tim Mulgan has forcefully objected to Murphy's Collective Principle in the hopes of showing that all forms of Collective Consequentialism are unacceptable.¹⁰⁴ One of his arguments that he wields against Murphy is called the *Wrong Facts Objection*. The objection, roughly sketched, is that in using the Collective Principle to decide what to do in a given situation, I base my conclusion too much upon contingent facts about that situation. These contingent facts affect the extent of my obligation very strongly. Consider the following version of the Wrong Facts Objection. Suppose that an affluent citizen of the developed world is prepared to do her fair share, but does not know how

¹⁰⁴ The discussion of Mulgan's objection which follows is from Mulgan 2001, chapter 4.

much is required of her. She only knows that she is one of three possible situations. In situation A, there are only one million people living in famine conditions. In situation B, there are fifty million people living in famine conditions. And finally, in situation C, there are 2.5 billion people living in famine conditions.

Mulgan stipulates that situation B is the actual world and that full compliance with the Collective Principle would require a donation of ten percent of the affluent citizen's income. When we examine what happens in situations A and C, however, we notice that our affluent citizen's fair share varies wildly. Decreasing the amount of starving people in the world by a factor of fifty, as in situation A, reduces our affluent citizen's fair share to a fifth of one percent of her income. Increasing the amount of starving people in the world by a factor of fifty, as in situation C, increases our affluent citizen's fair share to five hundred percent of her income. Since this obviously could not be sustained in practice, we should simply say that she is required to give all of her income away to charity.

Mulgan's objection to the results that the Collective Principle generates is that the large difference between the sizes of the fair shares is unreasonable. Why is this so? He thinks that the Collective Principle is responding merely to two facts: how much aid needs to be given, and how many people are available to produce that aid. But it generates odd results in two ways. Dramatically reducing the amount of good that needs to be done generates the result that our affluent citizen is required to sacrifice a relative pittance to the cause of famine relief, even when she understands that compliance with morality, as Murphy conceives of it, is not going to reach an ideal level because some of her fellow citizens are not going to donate. The facts that would be right for a principle of

beneficence to respond to, Mulgan says, is where her contribution would do the most good and about how much good she would be able to do.

Mulgan's objection is not very persuasive in situation A. If one is interested in promoting a culture that values the virtues of giving, generosity, and going above and beyond the call of duty, as Murphy does, it just isn't clear why we should object to someone trying to figure out what the fair share of everyone able to donate to charity is. If we figure out that everyone's fair share is relatively small and manageable, as it is in Mulgan's situation A, then it would be quite a bit easier to convince everyone that they ought to be donating to charity, and quite a bit easier to sound reasonable when (gently) castigating someone for failing to live up to their obligation.

But Mulgan's situation C is very problematic indeed. It suggests that if the size of the good that needs to be done is great enough, Murphy will be unable to avoid the Demandingness Objection. Perhaps it is unrealistic to suppose that situation C is anything like the real world. There are parts of the world that are in famine conditions, to be sure, but to have roughly every third person in the world starving would be an unheard-of disaster.

Unfortunately for us, there is a problem in the actual world on the scale of situation C. The effects of global climate change cannot be said to be a single unified natural disaster, such as in the case of situation C's world-historic famine. But there is no denying that the multitude of global climate change's effects -- including sea level rises, extreme temperatures, desertification, and increased drought and storm frequency, as well as their secondary effects, which include widespread food shortages, water crises, infectious disease outbreaks, and displacements of entire populations, -- makes it a

phenomenon on the scale of situation C's famine. It is, as Dale Jamieson has said in a recent paper, "the world's biggest collective action problem."¹⁰⁵

What would Murphy's Collective Principle require that we do about the harms of climate change that await us in the future? Recall what the Collective Principle says about situations of partial compliance:

In situations of partial compliance, a person's maximum level of required sacrifice is that which will reduce her level of expected well-being to the level it would be, all other aspects of her situation remaining the same, if there were to be full compliance from that point on.¹⁰⁶

We are currently, I suggest, in a situation of partial compliance. Our collective failure to act optimally has generated the greenhouse gases necessary for the process of global climate change and its attendant effects to come about. In that case, Murphy's principle says that we shall need to assume that everybody complies with the Collective Principle from this point forward. In other words, everyone shall be required to act so that no future climate-related disaster is any further abetted than it already is by human action. A moment's reflection on the extent to which our daily behaviors abet the process of climate change should illustrate the far-reaching nature of this result. To be sure, an immediate and radical change such as the Collective Principle would appear to require in our case may not be practical. Still, the recommendation is clear enough: given our current ecological situation, the Collective Principle recommends a drastic change to our daily habits, diets, methods of transportation, trade regulation, and level of support for

¹⁰⁵ Jamieson 2007, 165.

¹⁰⁶ Murphy 2000, 117.

ecologically sound technologies. Furthermore, future disasters produced by the process of climate change would require a worldwide disaster relief response from everyone unaffected by the disaster. These future disasters are not the result of any future person's partial compliance; thus, future persons would have to ready themselves to do their fair share of work to alleviate suffering from these disasters.

This is an extremely demanding result, and it can be derived directly from Murphy's Collective Principle plus some plausible assumptions about what conditions in the actual world are like. The Collective Principle thus fails to avoid the Demandingness Objection; there is no question that the Collective Principle, in the actual world, would alienate us from many of the projects and pursuits that we might have thought a moral life was compatible with.

One might be tempted to reply that my just-reached conclusion is not an effective objection against the Collective Principle. All I may have demonstrated here is that *prudence*, not the Collective Principle, can turn out to be extremely demanding. It is not, therefore, clear that the Collective Principle has fallen to the Objection. The point may be illustrated as follows:

The Selfish Will Suffer. Unless everyone adopts a very demanding, altruistic system of morality, everyone's immortal soul will be sent to a realm of anguish and torture for all eternity after their physical death. Suppose, furthermore, that we have indubitable, independently verifiable proof that this is the case: we have all heard a booming voice from the sky proclaiming the moral law, and the Gateway to the Underworld is a real, physical object from which one can easily hear the pitiful cries of the deceased selfish if one stands near enough.

In this case, whatever non-prudential reasons we have for being altruistic, we have a very excellent prudential reason to observe a highly demanding morality. The same is true in the case of climate change. If we want to avoid the worst effects of climate change, we should change our behavior immediately; but doing so is a matter of prudence, not of morality.¹⁰⁷

But an important difference exists between *The Selfish Will Suffer* and the case of global climate change. That important difference is that prudence has a very clear role to play in *The Selfish Will Suffer*. Its role is not nearly as clear in the case of global climate change. The effects of global climate change are predicted to intensify over the course of the 21st century, and many who are alive today will not be alive to experience the projected more intense effects of climate change.¹⁰⁸ However, many who are very young today and have had relatively little impact on the climate through their personal behavior *will* be alive to experience more of the intensified effects of global climate change. As already discussed in chapter 1, many of the harms that contribute to global climate change are individually negligible. The negative impact on an individual polluter of that individual polluter's behavior, moreover, is not nearly as clear for that individual as the negative impact of selfish behavior would be for an individual in *The Selfish Will Suffer* who is thinking about committing a selfish act. In *The Selfish Will Suffer*, it is beyond question that people have an excellent and very general prudential reason to adhere to a

¹⁰⁷ Thanks to David Sobel for raising this objection.

¹⁰⁸ Table SPM.2 in the Summary for Policymakers of the Intergovernmental Panel on Climate Change's Fourth Assessment Report (IPCC 2007) details a number of possible effects of climate change. The IPCC does not include specific ranges for projected effects. The IPCC does, however, have a high degree of confidence in their findings that climate change effects as diverse as increased risk of drought, heat wave, desertification, wildfire risk, decreased cereal production, and reduced access to water resources will increase, and progressively so, over the course of the 21st century.

very demanding altruistic moral code. In the actual world-case of global climate change, however, it is less clear that people have a general prudential reason to avoid climate change.

The upshot is that those generations (and perhaps in some cases, individuals) most responsible for the current increase in average global temperatures, whoever they are, will not directly suffer the negative consequences of their climate-affecting actions. The negative consequences of those actions contributing to global climate change will most heavily fall on young people alive today and on future generations. Avoiding global climate change, in moral terms, *would* be a matter of prudence for current generations if the negative consequences of pollution fell squarely on those responsible for polluting acts. But that is not the case; it is not the case that carbon emissions from driving one's car are, for example, directly pumped into one's own lungs and into no one else's; there is no system by which individual punishments for individual acts of pollution are given out; and we have no independent confirmation that anything like punishment for pollution exists in the hereafter. This means that whatever reasons exist for me to avoid polluting acts where I can and reduce their impact where I cannot avoid them, they must not be solely prudential reasons. I must have moral reasons to reduce my contribution to climate change. So my case does, I think, show more than just that prudence alone might turn out to be greatly demanding. I agree with the person who argues that prudence, in some cases, will be extremely demanding. In those cases, it will perhaps not be a mark against prudence that it is so demanding. But I think this thought experiment shows that, whatever role prudence plays in our response to global climate change, it is not clear that it plays the dominant role. Morality must, then, have something to say about our response

to global climate change, because prudence alone has thus far proven insufficient to address the problem.

Does the Collective Principle also generate the Polluter's Dilemma? To see whether this is so, let us run the same test that we ran in the previous two chapters against rule consequentialism and theories containing an agent-centered prerogative. Let's suppose that we are in a world where there is a consequentially optimal level of pollution, such that if there is more environmental regulation then scientific and economic activity will be hampered, and such that if there is less environmental regulation, climate change will accelerate beyond our ability to adapt to it. In such a situation, the Collective Principle requires us to assume full compliance with a plan of action that is optimal in terms of expected aggregate well-being from that point forward, taking into account the effects of actual failures of beneficence in the past. It would thus appear that the Collective Principle, if actually followed, would not generate the runaway result of the Polluter's Dilemma, as (for example) a consequentialist theory containing an agent-centered prerogative would.

Here, then, is a way in which Murphy's Collective Principle escapes the troubling result that hobbled rule consequentialism and agent-centered prerogative-containing theories: it is not directly collectively self-defeating. At least, it is not directly collectively self-defeating in this situation. It may be just that in other situations. But the troubling results we have uncovered so far are reason enough to reject the Collective Principle as an acceptable solution to the Demandingness Objection. In the previous section of this chapter, we learned that the Collective Principle regularly fails to make appropriate, plausible demands of agents. This is so because the Collective Principle depends on a

narrowly specified notion of fairness that is tailored to avoid the result in which those agents who do not comply with the demands of beneficence exploit me. In this section, we learned that there are circumstances in the actual world that would result in significant demands upon an agent if she complied fully with the Collective Principle. Even if we think that result is appropriate because it is necessary to avoid disaster, the Collective Principle's recommendations for low-cost-high-benefit cases seem notably inconsistent with its extreme demandingness in the case where we are forced to respond to the problem of global climate change.

4. The Modified Compliance Condition

To conclude this chapter, I want to consider whether we can use the fairness considerations I pointed to earlier to construct a modified Compliance Condition that would solve the problems the Compliance Condition has generated for Murphy. Let the new Compliance Condition read as follows:

Modified Compliance Condition (MCC). If other agents cease to comply with the demands of a principle of beneficence, an acceptable principle of beneficence will not increase its demands on agents who do comply with that principle unless agents who have already complied are either uniquely situated to bring about beneficent effects or particularly efficient at bringing beneficent effects about.

Let me now be explicit about what “uniquely situated” and “particularly efficient” mean. If I am *uniquely situated to bring about beneficent effects*, I occupy a place of optimal physical proximity and effective ability to the situation that no one else does, such that I am the only one who could bring about the good to be done. If I am *particularly efficient at bringing about beneficent effects*, my abilities or resources are such that a small sacrifice on my part brings about a beneficent effect that greatly outweighs my sacrifice.

Making this change to the Compliance Condition requires a modification of the Collective Principle. The modification is italicized:

Modified Collective Principle (MCP). Everyone is required to perform one of the actions that, of those available to her, is optimal in respect of expected aggregate weighted well-being, except in situations of partial compliance with this principle. In situations of partial compliance, a person's maximum level of required sacrifice is that which will reduce her level of expected well-being to the level it would be, all other aspects of her situation remaining the same, if there were to be full compliance from that point on, *unless that person is uniquely situated to bring about beneficent effects or particularly efficient at bringing beneficent effects about.* Under partial compliance a person is required to perform either an action that is optimal in respect of expected weighted aggregate well-being, or any other action which is at least as good in respect of expected weighted aggregate well-being, *if they are uniquely able to do so.*

However, no one is required to act in a way that imposes a loss on some other person unless that other persons' expected level of well-being after the loss would be at least as high as it would be, all other aspects of the situation remaining the same, under full compliance from that point on.

I indicated that, in section 2 of this chapter, that such modifications would be desirable because they would prevent absurd results such as an affluent person being permitted to deny a starving homeless person twenty cents if that affluent person has already done their fair share of bringing about beneficent effects. We also have an adequate rationale for making such a modification. The original motivation for the Collective Principle was the recognition that beneficence is a shared group project that needs to be fair in how it parcels out its demands. MCC is not an *ad hoc* modification; it responds to the same recognition that originally motivated the Compliance Condition, but differs in that it has a more nuanced understanding of fairness.

The main problem with the Collective Principle, as I understand it, is not that it generates Polluter's Dilemmas. The central problem was its failure to make appropriate,

plausible demands of agents who had already done their fair share. MCP is far less likely to permit such failures.

Consider again Rachels' two-rescuer case. Even if I have already complied with the demands of morality by saving one of the two drowning children, MCP requires me to wade in and save the second child if I am either uniquely positioned to do so or particularly efficient at bringing about the saving of the second child. I am the only bystander, and it would not be difficult for me to wade back in to the pond. I satisfy both conditions, so MCP does not let me off the hook. I must wade in and save the second child.

Let's modify the two-rescuer case slightly so that wading back in involves a danger to my health, where I am the rescuer. Suppose that the situation is the same as before; two children are drowning in a body of water, and I and another person are on hand to witness this. However, we are not faced with a pond. Instead, the two children are drowning in a deep lake, and the lake, like everything else in the immediate vicinity, is in the icy grip of winter. Perhaps because of this, the other potential rescuer walks away without doing anything to save one of the drowning children. I, however, comply with the demands of morality, and I plunge into the frigid waters to drag one of the children out. I am in some danger of hypothermia; after all, my clothes are now wet with freezing cold water. Does MCP tell me to wade back in to the water? That depends on how reasonably able I am to bring about the saving of the second child's life. If I am already suffering horribly from the shock of the cold, then I am not required to wade back in. I would not be uniquely situated to perform a beneficent action here, because I would not be physically able to withstand another trip into the water. Nor would I be particularly

efficient at bringing about the intended beneficent effects; it would not be a small sacrifice on my part to wade back into the lake, because the sacrifice of making another trip into the water might actually kill me, rather than effectively result in the saving of the second child's life and my continued survival. If I am able to withstand a second trip into the water, however, then MCP may tell me to plunge back into the lake. That is because I would still be uniquely situated to perform the beneficent action, even if I wasn't particularly efficient at performing the action. I am still the only other person around, and I would be effectively able to pull the action off.

And what if I find myself in the situation of the affluent citizen who has to decide whether to give a homeless person twenty cents? Again, the answer depends on whether I am uniquely situated or particularly efficient with regard to my ability to bring about beneficent effects. If there are bystanders who have not done their fair share and sacrificing more would leave me unable to pay for my home, then I am not required to give twenty cents. But if I am the only one around who can help, or if twenty cents would not represent a significant burden to me, then MCP would require me to give twenty cents to the homeless person.

These results seem to be consistent with common sense. MCP requires that we contribute our fair share toward a shared beneficent goal in normal circumstances. That is a feature shared with Murphy's original Collective Principle. The new feature, however, is that MCP requires that we take special care to evaluate whether we are uniquely able to make a difference in a particular situation. MCP is still concerned with protecting us from being exploited by people who do not do their fair share, but not obsessively so, and particularly not at the expense of our common-sense intuition that, in many cases, a duty

of rescue trumps considerations of fairness. So MCP is a little more demanding than Murphy's unmodified Collective Principle in this way, but it doesn't appear excessively demanding. It would be excessively demanding if it recommended that an affluent citizen must give twenty cents to a homeless person in the situation where (1) that affluent citizen has already done her fair share of donating to charity, (2) there are onlookers who have not complied with morality by doing their fair share of donating to charity, and (3) giving twenty cents would seriously damage the affluent citizen's prospects and interests. It seems, in that situation, that there would not be many limits to the demands that morality would make. But MCP would not recommend that this affluent citizen give twenty cents. It would, instead, require that the onlookers start doing their fair share.

The Collective Principle also faced the objection that it was unable to escape the Demandingness Objection in the actual world. Recall, from section three, the result of Mulgan's Wrong Facts Objection. If the size of the good to be done is large enough, then each individual's fair share of the good to be done will be overwhelming. In responding to the problem of climate change, the Collective Principle required us all to immediately shift to full compliance with a system of environmentally friendly behavior, with drastic consequences. The Collective Principle avoids generating the Polluter's Dilemma, but does not appear to avoid the Demandingness Objection. Does MCP do any better in this regard?

In a situation of partial compliance, like the unmodified Collective Principle, MCP requires us to come into full compliance with the demands of beneficence. It assigns fair shares of the good to be done to each individual. So MCP requires some sacrifice of most people. But there is some reason to think that MCP is less demanding in

the case of global climate change than the unmodified Collective Principle. MCP requires more sacrifice from people who are uniquely able to bring about beneficent effects. To be sure, there is no one unique person who is in a position to perform all the beneficent acts needed to solve the problem of global climate change. But there are many people who are in a better position than others to ensure that beneficent effects are brought about; politicians, policymakers, public intellectuals, thought leaders, important economic and business leaders, and the like. Those individuals are more able to shape the way the public perceives the threat of climate change, and more able to ensure that the changes necessary to mitigate and adapt to the effects of climate change are brought about. MCP requires them to do their fair share to mitigate the effects of climate change. But MCP requires them to do more than they have been doing to help the rest of the population and change. Perhaps, then, everyone else's fair share will be more manageable.

I think MCP is a modification that Murphy could accept, because there is a solid rationale for doing so and MCP retains Murphy's sense of the importance of protecting people who care about complying with morality from exploitation by those who do not. I do not, myself, know whether MCP is true, and I don't claim that it is true. But I do claim that it is an improvement to the original Collective Principle, and that it is a better way of avoiding the Demandingness Objection than the original Collective Principle was. Remember that the original Collective Principle, even though it avoided the Polluter's Dilemma, was moderated so much that it regularly failed to make appropriate, plausible demands. I argued that a deeper analysis of the notion of fairness could help the Collective Principle out, and MCP is the result. There is some reason to think that MCP avoids both of the objections that I raised against the unmodified Collective Principle. I

conclude that if the Collective Principle is modified to reflect a more nuanced understanding of the concept of fairness, it is a better response to the Demandingness Objection than the previous alternatives we have surveyed.

In the next chapter, I wish to survey another version of consequentialism that also represents a promising answer to the Demandingness Objection. It shares with MCP the ability to avoid both the Demandingness Objection and the Polluter's Dilemma, and it does so without placing enormous burdens on those who care about morality. Once we have considered that version of consequentialism, we shall be in an even better position to see what features are desirable in an acceptable version of consequentialism.

Chapter 4: Multiple-Act Consequentialism and Constraints on Acceptable Forms of Consequentialism

1. Introduction

We have reviewed several prominent forms of consequentialism that have failed to successfully avoid the Demandingness Objection. Three varieties -- Scheffler's Hybrid Theory, Mulgan's Combined Consequentialism, and Hooker's rule consequentialism -- fall victim to the Polluter's dilemma. Murphy's Collective Principle of Beneficence avoids the Polluter's Dilemma, yet in doing so generates an extremely demanding result. None of these theories, then, are acceptable forms of consequentialism.

I must say more about why this is so. What I propose to do in this chapter is not simply to review the results of the negative arguments of previous chapters, but to propose a general analysis of why these forms of consequentialism go awry. Using the results of that analysis, I shall propose and argue in favor of a proper constraint on any acceptable form of consequentialism. But the task of this chapter is not merely to sum up the negative account of this project. Once that task is finished, I shall analyze a new version of consequentialism, called Multiple-Act Consequentialism, or MAC for short.¹⁰⁹ I shall argue that MAC avoids both the Demandingness Objection and the Polluter's Dilemma, and thus represents one promising path for future research on consequentialist moral theory to take.

¹⁰⁹ Mendola 2006.

2. Consequentialism, Its Rivals, and Climate Change

I have been presuming, up to this point, that climate change counts as a moral problem. It seems appropriate to defend that position here. Doing so will strengthen my argumentative position when I later argue that failure to deal with climate change adequately is a proper constraint on acceptable forms of consequentialism.

I start with a very general observation: pollution is everyone's problem, but not everyone agrees whether it is, properly conceived, a moral problem. The facts are striking in their scope: human industrial and economic activity, beginning with the Industrial Revolution and proceeding increasingly quickly as decades pass, is transforming Earth in ways that have severely taxing implications for many contemporary humans, future humans, and other species present and future. Depending on what they think the nature of morality is, philosophers have not universally responded to the problem of climate change as a moral issue.

For example, Christine Korsgaard, a prominent defender of Kantian ethics, holds that “bringing something about” is a feature of consequentialist thought that “distorts our thinking” about ethics.¹¹⁰ Utilitarians, she holds, are obsessed with the preservation of the environment.¹¹¹ According to Korsgaard and Kant, the domain in which morality may rightly sanction us for doing something wrong is much narrower than the consequentialist suggests.

It is tempting to take away the impression that Kantians do not have much to say about environmental protection, at least directly. But Kantians have many resources to draw on when considering how to be beneficent toward the least well off in society.

¹¹⁰ Korsgaard 1996, 275.

¹¹¹ *Ibid.*, 300.

Onora O'Neill has argued that Kantians do have a basis for beneficent action toward those in famine conditions: by not aiding people bedeviled by famine, we fail to respect their autonomy and treat them as ends in themselves.¹¹² We actually work to lessen their autonomy by refusing to provide them with food.

The same argument could be extended to the case of those suffering due to global climate change. When we fail to pursue environmentally sound policies and instead satisfy our trivial desires that involve pollution as a side effect, we are essentially treating other human beings, especially those most vulnerable to climate change, as a mere means to an end. Their raw materials and the capacity of their ecosystems on which they depend to absorb the pollution we generate provide the means to our own idiosyncratic end. By adopting policies that create refugees from climate change, whether they are citizens of the Maldives escaping rising sea levels, Bangladeshi citizens fleeing destructive floods, or low-income Americans whose inexpensive homes are devastated by severe thunderstorms and tornadoes, we have done serious harm to the autonomy of those refugees. If that is true, then a Kantian argument exists which implies that we should examine our own attitude toward the permissibility of polluting activities far more closely and harshly than we currently do, though it remains to be spelled out exactly what kind of practical consequences this argument would have.

Kantianism, however, is in better shape than contractarianism and common-sense ethical pluralism.¹¹³ Dale Jamieson has recently taken these two ethical theories to task for their laxity on the issue of climate change and I largely agree with his assessments,

¹¹² O'Neill 2007, 558.

¹¹³ For an example of a work in the contractarian tradition, see Gauthier 1985.

though, as noted above, I am less negative about the prospects for Kantianism.¹¹⁴

Contractarianism, Jamieson argues, excludes many types of beings from moral consideration who are not party to moral agreements being made, and yet we believe that many of these types of beings are just the sort of beings who deserve consideration when dealing with environmental problems -- future people, non-rational humans such as infants and those suffering from dementia, and non-rational animals.¹¹⁵ Common-sense pluralists are generally averse to consequentialism's tendency toward revisionism, but a moral theory which deals adequately with climate change's dangers would be very revisionist in scope.¹¹⁶

Consequentialism, because of its orientation toward the future and its insistence on evaluating outcomes as the determining factor in what makes an action right or wrong, seems ideally positioned to have an excellent response to the problem of climate change. But things are not so simple. Perhaps it is the case that the paradigm case of consequentialism, which is classical act utilitarianism, is not a plausible way of dealing with our environmental problems. Indeed, it may have contributed to them. Environmental philosophers are at least as skeptical toward consequentialism as they are toward Kantianism and contractarianism.¹¹⁷

Indeed, consider the Polluter's Dilemma applied to act utilitarianism. Suppose, once again, that we are in a situation where we have reached a consequentially optimal level of pollution. If we restrict pollution more tightly than we already do, we hamper economic productivity; if we restrict pollution less, the runaway process of climate

¹¹⁴ Jamieson 2007.

¹¹⁵ See Jamieson 1981 for more.

¹¹⁶ Jamieson cites Wolf 1982 as an example of revisionism-aversion.

¹¹⁷ For criticisms from prominent environmental philosophers, see Callicot 1980, Rolston 1999, and Sagoff 1984.

change gets going and we lose our ability to successfully adapt to it. Now suppose you are a good act utilitarian trying to decide whether to drive yourself to the mall to go shopping or to stay at home and be content with what you have. If you drive yourself to the mall, you exceed the pollution limit and the Polluter's Dilemma gets going. But as we have seen, the harm of a single polluting act is extremely small; as a good utilitarian, I may thus cynically drive myself to the mall, generating some utility for myself (and perhaps for others, since I am now in a good mood and less apt to be a sourpuss toward everyone else) while contributing to the biggest collective action problem in human history. Classical act utilitarianism thus looks doubly bad. Not only does it fail to avoid the Demandingness Objection, but also it fails to avoid the Polluter's Dilemma. Jamieson concurs:

But when it comes to large-scale collective action problems, calculation invites cynicism...because it appears that both morality and self-interest demand that 'I get mine', since whatever others do, it appears that both the world and I are better off if I fail to cooperate. Indeed, it is even possible that in some circumstances the best outcome would be one in which I cause you to cooperate and me to defect. Joyriding in my '57 Chevy will not in itself change the climate, nor will my refraining from driving stabilize the climate...Since everyone, both individuals and nations, can reason in this way, it appears that calculation leads to a downward spiral of non-cooperation.¹¹⁸

3. The Polluter's Dilemma as a Constraint on Acceptable Forms of Consequentialism

Consequentialists have been historically most interested in avoiding the Demandingness Objection. Perhaps this is unnecessary; after all, there is some precedent suggesting that standard objections to most moral theories are question begging in subtle

¹¹⁸ Jamieson 2007, 167.

ways.¹¹⁹ But consequentialists cannot afford to be lax in combating the Polluter's Dilemma. Failing to recommend widespread revisions in the way we ordinarily behave, and thereby failing to resolve the problem of global climate change, would demonstrate the complete inadequacy of consequentialism as a plausible ethical system. This is because consequentialism would be shown to suffer from a deep incoherency. Here is the incoherency: consequentialism is fundamentally about producing the best state of affairs possible for the greatest number, impartially judged; but if consequentialism fails to solve the Polluter's Dilemma, then consequentialism will be, in essence, telling us to produce the best state of affairs possible, and to do so *in a way that will fail to produce the best state of affairs possible*. Consequentialism, then, would be self-refuting -- proven false on its own terms.

The objection that a theory is self-refuting is, of course, not a new kind of objection.¹²⁰ The most influential modern treatment of this topic is Derek Parfit's treatment of directly collectively self-defeating moral theories in chapter 2 of *Reasons and Persons*.

According to Parfit, the "T-given aims" of a moral theory T are the goals that a moral theory tells us to seek. And we "successfully follow T" whenever we manage to do that act which best accomplishes the T-given aims of our moral theory. With that terminology in hand, Parfit proposes the following disjunctive definition of direct

¹¹⁹ Korsgaard makes this point in Korsgaard 1996, xiii. For a sustained defense of consequentialism against the question-begging nature of the Demandingness Objection, see Sobel 2007 and for extra support of Sobel's view, see Tedesco 2011.

¹²⁰ Indeed, as with much of interest in philosophy, one of the earliest cases of such an objection may be found in Plato; specifically, in the *Theatetus*, where Socrates strenuously objects to Protagoras's relativistic theory of truth. See Plato 1997, 197-198 (178b-179d).

collective self-defeat. A moral theory T that gives the same goals to every agent is directly collectively self-defeating when

- 1) It is certain that, if we all successfully follow T, we will thereby cause our T-given aims to be worse achieved than if none of us had successfully followed T; or
- 2) our acts will cause our T-given aims to be best achieved only if we do not successfully follow T.¹²¹

Parfit is motivated to make this a proper constraint on acceptable theories of morality because most philosophers accept that morality is ultimately a collective project that requires the possibility of success at the collective level.¹²² And although Parfit does not at length defend this motivation, it is a fairly plausible thesis which other authors are concerned to defend. As an example, Alan Gibbard has pursued a similar thesis in a recent series of Tanner Lectures by arguing that planning how to live together on terms that no one could reasonably reject requires us to have some kind of method by which we balance out competing projects. Thus, a consistent theory of how to live together would have to be, in a crucial sense, consequentialist.¹²³

There is some disagreement over whether being directly collectively self-defeating should actually be a proper constraint on moral theories. Joseph Mendola has objected to this proposal, arguing that any plausible moral theory would violate the constraint given the correct circumstances.¹²⁴ Circumstances might be such that, if we were to all successfully act on our dearly held consequentialist moral theory, some disastrous event would befall everyone. Mendola's thought experiment involves an extraterrestrial attack on a community of consequentialist humans. Whenever the humans successfully follow

¹²¹ Parfit 1984, 54; italics Parfit's. Parfit includes a definition of direct collective self-defeat for moral theories which give different goals to each agent on pp. 55, but I ignore that here as I am concerned only with consequentialism, which gives the same goals to each agent.

¹²² Ibid., 106-108; see also Parfit's footnote 58.

¹²³ Gibbard 2008, lectures 2 and 3.

¹²⁴ Mendola 1984.

their theory by seeking to promote the impartial good of all, the extraterrestrials inflict grievous harm on everyone. It is only by abandoning consequentialism that the extraterrestrials are placated. In this case, consequentialism would be directly collectively self-defeating, as would a variety of other moral theories that counted beneficence as important. The troubling upshot, then, is that a wide swath of morality is completely ruled out by the constraint on self-defeat. So Mendola would perhaps recommend that I cease wielding the self-defeat constraint as a weapon against other normative theories, because using the self-defeat constraint is a kind of philosophical Mutually Assured Destruction.

Mendola considers and responds to many possible objections to his argument. One of those possible responses is that his case is not only fictional, but too implausible to take seriously for the purposes of moral theorizing. Mendola rejects this rebuttal. He makes two points in response.

The first point is that the case is plausible, despite its initial reliance on stringently deontological extraterrestrials. It is possible, he claims, that we could develop the requisite technologies that would enable us to make life miserable for a specified group of humans who behave in consequentialist ways.¹²⁵ Given some basic knowledge about the extent of spying capabilities of intelligence agencies, it is hard to dispute this point.

Mendola's second point is to insist that saying his case is "irrelevant because implausible is to insist that moral theories are contingently true or false."¹²⁶ Maybe consequentialism is true right now, in the actual world, but if the deontological extraterrestrials descend from the sky, consequentialism will suddenly become false.

¹²⁵ Ibid., 161

¹²⁶ Ibid.

Indeed, this is just what Mendola claims: consequentialism will be false if the evil extraterrestrials exist.¹²⁷ Surely, he thinks, the truth of moral theories cannot vary when contingent facts of the situation change.

What is it for a moral theory to be false? One way a moral theory can be false is to be take arbitrary facts into account. In relevantly similar situations, with all morally relevant facts staying the same, a moral theory might be false if it told us to do different things based on some morally irrelevant fact like the color of a person's skin. Here is a novel form of consequentialism, which I shall call *orange socks on alternate Wednesdays act utilitarianism*. Combine a hedonic theory of value with the classical act consequentialist theory of right action, but jettison the classical utilitarian emphasis on strict impartiality. Instead, stipulate that the people whose utility are to be maximized are those people who consistently wear orange socks on alternate Wednesdays. This theory is false because there is no morally significant justification for promoting the value of such people.

A second way that a theory could be false because it is internally inconsistent, such as when it affirms both of the following statements:

(1) "All forms of life are infinitely morally valuable and a moral person ought to never harm any form of life."

(2) "It is morally permissible to clean my shower with a bleach-containing agent."

Obviously, if I use bleach in my shower, I kill many microorganisms living in my shower drain. I would have to believe that doing so is morally permissible to carry out my act. But that belief is inconsistent with the first statement. This moral theory is false.

¹²⁷ Ibid.

A third way for justifiably concluding that a theory is false is for that theory to produce wildly counterintuitive results, either on a very crucial point on which most reasonable people are agreed, or to consistently produce counter-intuitive results on a wide variety of issues. Classical act utilitarianism has been thought by many to be false because there are situations in which classical act utilitarianism recommends that we violate someone's rights in order to maximize utility. This is, of course, a less reliable method of showing a moral theory to be false. An act utilitarian might bite the bullet and argue that the theory is correct, regardless of what our intuitions tell us about a particular case. But we might nevertheless *justifiably* conclude that the act utilitarian's argument is insufficient to save the moral theory from the objection.

Suppose that Mendola's extraterrestrials do exist. Is consequentialism false? More precisely, is consequentialism false in this case because it commits one of the errors just listed? So far as we can tell, consequentialists are not focusing on morally irrelevant details, nor are they behaving in obviously counterintuitive ways. Moreover, consequentialism is not obviously internally inconsistent. It isn't affirming two mutually incompatible theses.

What Mendola has shown is that we might say that consequentialism is being shown to be a theory that cannot be successfully followed. It recommends that we follow a course of action that should in theory lead to great benefit but in practice leads to horrible suffering. If Prisoner's Dilemmas are as common as Mendola thinks, then his extraterrestrials case is the merely the tip of the spear.

I accept the contention that being directly collectively self-defeating shows that a moral theory is false. However, Mendola's results are limited. He has only shown that

one version of consequentialism is false. There may be others that are not directly collectively self-defeating. The extraterrestrials in his case stipulate that they will do their worst only if the consequentialist humans universally and successfully follow their moral theory. But if our consequentialist theory is versatile enough to tell most people to continue maximizing the good and to simultaneously permit a defection from the theory such that at least one person becomes and behaves as a non-consequentialist, then that theory will avoid being directly-collectively self-defeating. Collectively, we will have acted so as to achieve the goal of the theory (i.e., to maximize the good), but we will not have universally followed the theory. I think this is defensible as being compatible with consequentialism because it is far more important to a consequentialist that an act satisfy the consequentialist standard of rightness than it is to have an action carried out according to a formal procedure from which no act may be allowed to deviate. Such a constraint is not properly part of consequentialism, because it has no consistent consequentialist rationale. It may be right for a consequentialist to tell someone else that he should not behave as a consequentialist, if not behaving as a consequentialist has better consequences than not behaving as a consequentialist.

Another reason exists for me to not give up the line of attack that relies upon direct collective self-defeat. The problem of global climate change is a special problem for consequentialism, and the test of the Polluter's Dilemma can tell us whether consequentialism is even worthy of consideration, on its own terms. Consider the revisionist nature of consequentialism, alluded to in chapter 1. Consequentialism's (or rather, utilitarianism's) future orientation, impartiality, and dogged insistence on examining the actual effects of moral decision-making on the lives of the members of the

moral community has historically made it a moral theory especially well-suited for proposing intuitively much-needed change. Thus, the use of the self-defeat argument could be read less as an insistence on a formal rational constraint, and instead as a test of whether consequentialism can solve its problems and live up to its promise. If consequentialism cannot answer the test of the Polluter's Dilemma, then traditional ethical theory would be deprived of the most straightforward explanation possible as to why global climate change is a colossal moral problem: because it imposes widespread, temporally-extended, and dramatic suffering on the widest possible variety of Earthly creatures. That explanation is practically valuable to our moral thinking.

I am thus not worried about the Polluter's Dilemma on the ground of formal rationality. Consistency is important, but the worry I am pressing is *not* the insistence, rooted in that ground of formal rationality, that consequentialism should not demand its own abandonment. Taking that worry seriously would be to have decamped for non-consequentialist territory already. Rather, the insistence on consequentialism being able to avoid self-defeat is in pursuit of the goal of shoring up consequentialism against one of the traditional complaints people make against it; that consequentialism is too permissive in certain perverse ways and thereby permits me to cause grievous harm in pursuit of overall maximization of the good. This worry is especially pressing in the case of climate change. If a concern for maximizing the good leads directly to catastrophic results, then it isn't the case that the theory is recommending its own abandonment. Rather, we are demonstrating the theory to be completely inadequate to achieve the very goal we hope it accomplishes.

It's worth asking at this point just *how* other versions of consequentialism have proven themselves to be completely inadequate. I think the answer has something to do with how consequentialists have responded to the Demandingness Objection's insistence on space for the personal point of view. Recall that Bernard Williams's excellent statement of the Demandingness Objection involved an explicit appeal to the need to defend the integrity of a person:

[A person] is identified with his actions as flowing from projects and attitudes which in some cases he takes seriously at the deepest level, as what his life is about....It is absurd to demand...that he should just step aside...and acknowledge the decision which utilitarian calculation requires. It is to alienate him in a real sense from his actions and the source of his action in his own convictions....It is thus, in the most literal sense, an attack on his integrity.¹²⁸

This, combined with Rawls's insistence that utilitarianism, due to its aggregative nature, does not respect the difference between persons, has left consequentialists understandably anxious to demonstrate that their theory avoids this flaw. But two theories that I examined in a previous chapter have committed the mistake of combining too much room for the personal point of view with the contingent fact that some harms are such that they are imperceptibly tiny, yet of an aggregative nature. In chapter 1, I considered two forms of consequentialism that shared an Agent-Centered Prerogative. I rejected them on the grounds that the Agent-Centered Prerogative grants far too much weight to the personal point of view, thus allowing for the possibility that individual idiosyncrasies would easily be able to trump the small harms generated by polluting acts. So that version of consequentialism fails, and it does so instructively. The lesson to learn from those two theories' failures is that too much space for the personal point of view is counterproductive. And yet, it is hard to deny that integrity is important. So, a better form

¹²⁸ Williams 1973, 116-117.

of consequentialism would grant some integrity to persons without allowing them room to defect from cooperative acts due to their desires to pursue their own projects, regardless of what those projects are.

In chapter 2, I considered how Brad Hooker's form of rule consequentialism hews close to common-sense morality. It hews so closely to common-sense morality that, in fact, it would either not recommend adoption of a "prevent disaster" rule that would prevent the generation of Polluter's Dilemmas because Hooker wishes to respect our intuitive notions about how hard a moral theory should be to follow, or it would adopt the "prevent disaster" rule and endorse an extremely demanding result. But I showed that Hooker's view rests on unstable notions of what would count as too demanding for hypothetical people. The instability of such intuitions is enough to reject Hooker's rule consequentialism. So a better form of consequentialism would not rest on appeals to counterfactual situations, or what some suitably idealized actors would tell us to do.

And in chapter 3, we saw how Liam Murphy's Collective Principle of Beneficence is problematic. This is due to its reliance on facts about what my fair share would be in a situation where everyone was complying with the demands of morality. Murphy's thought seems to be that I am not allowed to get away with not doing anything beneficent, but there will come a point at which, despite whatever demands of morality that remain, I will be permitted to stand aside and cease cooperating. The unique flaw of Murphy's view is that it is apparently insensitive to the kinds of beneficent demands being made. We can perhaps make sense of a fair share in the case of famine relief. But other kinds of beneficence -- for example, taking action to avoid climate change -- do not seem to admit of a stopping point. Murphy's view also shares with Hooker's view the

insistence on making the criteria of right action dependent upon what obtains in counterfactual situations, and such an insistence causes problems for Murphy just as it does for Hooker. A better version of consequentialism, finally, would be sensitive to the kinds of beneficent demands being made.

Ideally, our hypothetical better version of consequentialism would be able to deploy very general theoretical mechanisms in order to accomplish these tasks. The agent-centered prerogative, for instance, is introduced seemingly for the sole purpose of defending an agent's personal point of view. But its specialized nature makes it the perfect belt loop with which to catch oneself on the outstretched hook of a counterexample. Such counterexamples would not work against more generalized theoretical machinery. There is at least one version of this kind of theoretical machinery available to consider, and we would do well to consider it here.

4. Mendola's Multiple-Act Consequentialism

In previous chapters, I considered and rejected three moderate versions of consequentialism. In this chapter, I considered and rejected classical act utilitarianism as equally inadequate. But none of these are the best version of consequentialism. Perhaps we don't know what that version of consequentialism is. We know a little bit about what it looks like. And if we went about trying to build it, we would know where to start. A version of consequentialism exists which seems to satisfy the list of desiderata from the last section of this chapter, and that is Multiple-Act Consequentialism. MAC is a recently proposed novel version of consequentialism that is designed to avoid both the familiar objections to standard act-consequentialism and the unstable consequentialist rationale of

indirect forms of consequentialism.¹²⁹ In the following sections, my goal is to make out of MAC a case study in how consequentialism may avoid both the Demandingness Objection and the Polluter's Dilemma.

MAC is a variety of classical act utilitarianism which retains two crucial components of that theory: (1) its insistence on direct consequentialist analysis of the rightness or wrongness of acts, and (2) its injunction that maximization of the good determines what is morally right and wrong. But MAC departs from classical act utilitarianism in a number of interesting ways. The most interesting of these departures, which, I believe is the central reason why MAC avoids the Polluter's Dilemma, is Mendola's account of the metaphysics of group action.

A central reason why many philosophers have objected to act utilitarianism is this: act utilitarianism appears to insist that we must directly and thoroughly evaluate the consequences of each individual act we propose to perform in order to determine whether that act is morally obligatory or forbidden,¹³⁰ and a variety of ways exist in which our calculations may lead us astray. In practice, it is unlikely that such a calculative effort would be successful, for it requires that we first arm ourselves with a relatively precise quantitative account of the values of pleasure and pain. Such an account is difficult to produce. Aside from this, as Peter Railton has pointed out, the direct evaluation of an

¹²⁹ Mendola 2006.

¹³⁰ This is a point on which consequentialists disagree. Peter Railton (in Railton 1984) rejects the notion that the act consequentialist must subjectively evaluate the consequences of his or her actions. However, Jeremy Bentham appears to advocate that a utilitarian should do exactly what Railton thinks the consequentialist need not do. Bentham writes, "A man may be said to be a partizan of the principle of utility, when the approbation or disapprobation he annexes to any action, or to any measure, is determined by and proportioned to the tendency which he conceives it to have to augment or to diminish the happiness of the community" (Bentham 1948, 3). And, of course, Bentham lists many separate criteria by which the utilitarian may evaluate the goodness of an act, as well as the procedure for going about the evaluation (Ibid., 29-30). Bentham does not think that such an extensive evaluation needs to be done for every act, but he does appear to think that agents trying to figure out what to do need to be able to carry out this task when necessary and need to match their actual decision procedures to this ideal decision procedure as much as possible (Ibid., 30).

individual act's consequences is subject to a regress.¹³¹ Before I decide what to do in a given situation, I must decide on the optimal time to take in making my decision. But I might have to decide how much time is optimal to allocate to deciding how much time to allocate for making my decision. And so on.

Familiar to us by now, act utilitarianism will often recommend that we sacrifice our time and resources for the general welfare, leaving us unable to pursue our own projects, because much of our own income is spent in trivial ways that are not nearly as efficient at generating maximal impersonal utility as a charitable donation to someone in desperate need would be.¹³²

Finally, classical utilitarianism is insensitive to the way in which overall utility is distributed. It only enjoins us to generate as much utility as possible; it does not require that we distribute it in any particular way.

Philosophers have responded to this problem in different ways. Derek Parfit, for example, demonstrated that classical utilitarianism generates the recommendation that we ought to favor the creation of a nightmarish maximally-populated world in which people live lives that are on balance only barely worth living, if that is, in fact, what would generate maximal utility, and Parfit argues that it is highly likely that any theory which accepts a requirement to maximize value without any concern for how it is distributed generates this nightmare world. This recommendation Parfit refers to as the Repugnant Conclusion.¹³³ John Rawls developed his explicitly non-consequentialist theory of justice in response to this problem of classical utilitarianism, claiming that a proper understanding of justice needs to respect the difference between persons, and part of

¹³¹ Railton 1984, 153-154.

¹³² Singer 1972.

¹³³ Parfit 1984, 388.

respecting this difference is to combine the desire to maximize value with prioritizing the well-being of the least-well off people in society.¹³⁴

A crucial insight of MAC is that many of the ills of classical utilitarianism are traceable to its conception of acts as standalone choices about what maximally promotes the good. According to MAC, this conception of acts is incomplete. One of MAC's key principles is the thesis that group actions and group agency exist, both within a human life and across different human lives.¹³⁵ And furthermore, Mendola argues, there is often no fact of the matter that any single action we perform is *only* part of one group action, rather than being a part of one or more group actions. Thus, group actions may overlap in one single momentary action. As an example of the kind of group action that exists within a human life, consider any action that cannot be done in the space of a moment; for example, the action of writing a research paper. Such an action - or, perhaps, it is better to call it a project - is something that can only effectively be done over time. It will require the cooperation of yourself today, such as when you go to the library to conduct a literature search, with yourself tomorrow, when you begin reviewing the literature, and of yourself at those times with yourself at a number of other further future points, when you are engaged in the actual act of writing the paper.

Your literature review tomorrow, however, is not simply a part of the single group act of writing the research paper. It may be a part of more than one group act within your own life. If you hope to achieve tenure, then conducting a literature review for your soon-to-be-written paper is not merely a part of your paper-writing project. It is part of the larger group act that constitutes your earning tenure, which requires the cooperation of

¹³⁴ Rawls 1971.

¹³⁵ A similar idea is proposed in Jackson 1987.

yourself today with yourself tomorrow, and the next day, and so forth -- until the tenure committee gives you a favorable report, long after your research paper has been published.

Things are more complicated still, however. For, if we stipulate that you are in a stable, permanent relationship, your literature review tomorrow is not simply part of the relatively small group act of writing the research paper, nor is it both part of writing the paper and part of the larger group act of earning tenure. Your conducting a literature review is, plausibly, part of the much larger group act involving both yourself and another person at multiple times, whose overall group action is to succeed at building a life together. Key to accomplishing this task is your gaining tenure. And key to gaining tenure is writing a research paper. And key to writing a research paper is conducting a literature review. When you conduct a literature review, there is very often no fact of the matter about what single group act you are doing, according to MAC. Your single action is often part of many group actions all at once, just as you, a single momentary agent, are part of many group agents all at once. So not only do group actions exist, but also any given action may be a part of a multiplicity of overlapping group acts.

A notion at work in the background here is Mendola's position that group acts are performed by group agents, which are in any case composed out of cooperating "atomic" agents over time. Atomic agents, Mendola says, are the "basic cells" of the group agent, whether the group agent is contained within and persists over a single life or the group agent is comprised of cooperating atomic agents ranging over many separate lives.¹³⁶ If you are trying to earn tenure, then the multiplicity of individual moments of yourself over the relevant timespan where you are working toward tenure is a group agent of the former

¹³⁶ Mendola 2006, 33.

sort. If you are trying to build a successful life with your spouse or partner, then you and your partner jointly constitute a group agent of the latter sort.

Under what conditions do group agents exist? According to Mendola, a group agent exists if and only if a collection of atomic agents take common action toward a goal they share, when this action is rooted in shared true belief that a common goal exists, and when there is a reason to continue pursuing this goal until the goal is either accomplished or abandoned.¹³⁷ The key point, Mendola stresses repeatedly, is that group agents must be “literally try[ing]” to accomplish some shared goal over time, whether the group agent in question persists over a single life or ranges over several lives.¹³⁸ Two strangers merely walking side-by-side do not, as Margaret Gilbert writes, constitute an act of going for a walk together. There is a real difference between that situation and a different one in which you and your partner both have the goal of going on a stroll and both believe that you share this goal with each other. If you start speeding up in the former case, your speeding up will occasion no criticism from the stranger walking beside you. But if you start speeding up in the latter case, you may, rightly, come in for some criticism from your partner for interfering with the work-in-progress of going on a walk together.¹³⁹

5. Atomic Agency Within a Life?

The fact that group acts may overlap is the feature of this metaphysics of group action that, I believe, allows Mendola to escape the Polluter’s Dilemma. But for a moment, I want to return for a moment to an aspect of Mendola’s view that Elinor Mason

¹³⁷ Ibid., 35-36.

¹³⁸ Ibid., 38.

¹³⁹ Gilbert 1996, 178-180.

has briefly criticized. This aspect is the proposal that group agency exists within a life. Mason worries that Mendola's conception of group agency runs aground on considerations that favor Mason suggests that there is reason to favor the typical conception of agency (on which an agent is just an ordinary human being directing its life) over Mendola's "overlapping" view of agency. The reason is that I can defect from many entities that would be group agents in Mendola's sense, but I cannot defect from, as Mason writes, "the group agent that consists of the group of temporally distinct atomic Elinors."¹⁴⁰ Mason's point that I cannot defect from my own life is well-taken; I cannot "opt out" of living my existence. Actually, I can, but that is on pain of ceasing my existence. By contrast, I go on existing if I renounce all existing ties to partnerships or institutions, mean existence though the result may be. But all the same, some question-begging is taking place here. It is an unjustified assumption that I cannot defect from a group *project* within my own life. To see this, consider: I am writing a dissertation now. Three years ago, I was only working toward the writing of a dissertation. Were I to cease writing a dissertation now, I would be defecting from a sincerely held goal shared by many past time-slices of myself. Grant the gift of foresight to myself exactly three years ago, and suppose for a moment that younger version of myself sees that I have grown weary and decided to cease writing the dissertation. It is highly plausible that shock, followed by righteous anger and withering criticism directed at me, would be my younger self's reaction. And rightly so: I would be defecting from a shared goal that has great importance not only for myself now but also for myself at many previous (and future) times. Hypothetical cases involving familiar science-fiction plot devices are always

¹⁴⁰ Mason 2007.

suspect, but this intuition is plausible enough. Mason's suggestion does not carry as much weight as it would first appear to.

Perhaps another reading of Mason's worry might be appropriate. This second reading runs as follows: I can defect from the ordinary kinds of group projects involving partnerships and institutions, but I cannot defect from the group agent that ranges over my own life because it is not the case that I am cooperating with past or future versions of myself. That is because there is no one else that I am now cooperating with when I am writing a dissertation. Only myself right now and the projects that I care about exist. What Mendola identifies as "cooperation" between atomic agents within a life is better thought of as a causal story about how I come to have the goals and plans for achieving those goals, and the folk phenomenon of "making an agreement with oneself" -- for example, to avoid eating ice cream excessively -- is not literally a mutual recognition that there is a shared goal and a reason to continue cooperating. It is, rather, a simple matter of placing a constraint on oneself about what constitutes acceptable behavior. But the constraint is never something that multiple atomic agents are agreeing to; there is only one agent, myself.

Mendola's response is that the notion of group agency within a life is not really controversial. Atomic agency within a life is meant to be an "analytic convenience," not reflective of any deep metaphysical underlying reality.¹⁴¹ All Mendola means to suggest, it seems, is that we can recognize more or less discrete stages within a life by looking at how the set of a given agent's options is filled out.

Still, we can say more. I have a pretty good idea of what I will be like tomorrow, because I cannot reasonably foresee any plausible scenario in which I lose my familiar

¹⁴¹ Mendola 2006, 33.

desires for chocolate, having a dog, going on five-mile runs, and finishing my dissertation. But I don't know what I will be like in the distant future. Perhaps I will become a couch potato who loves cats and has no desire to eat chocolate. But I am at least weakly connected to my future self.¹⁴² It seems permissible to assume that long-term projects like writing a dissertation and saving for retirement will be valuable to myself-plus-forty-years.

Let's examine the example of saving for retirement in this context. When I make the decision to save ten percent of my income for retirement, that decision has an attendant hope, accompanied by the threat of self-criticism, that I will not become profligate and change my savings plan such that I am saving almost nothing and spending all my money in the present moment. So when I am trying to make decisions about how best to benefit my distant future self, my thought process is more akin to working out how best to cooperate with a different person removed in space from myself, and it is less akin to placing a special constraint on myself now. If that is true, then the notion of an atomic agent existing within a life seems to make perfectly good sense. Because this is how my psychology operates, it also seems like the notion of an atomic agent within a life is less an analytic convenience and more like an accurate reflection of my psychology.

Of course, this could merely be a report about my own psychology, and thus not a very strong argument. Maybe all I have established with any degree of reliability is that atomic agency exists within my own life, but no one else's. After all, it is not a very strong inductive argument to reason from one's own case to a general truth. But my case is not special. We can make relatively good sense of many types of widespread

¹⁴² Derek Parfit has deployed a similar notion of identity across time; see Parfit 1971.

phenomena, such as self-criticism for failing to live up to one's own past judgments and failure to achieve sincerely held goals because of momentary *akrasia*, with the explanatory device of atomic agency within a life.

6. How Overlapping Group Agency Plausibly Solves the Polluter's Dilemma

The main thread of this discussion, however, is to see how MAC avoids both the Demandingness Objection and the Polluter's Dilemma. And now, having done some work to explicate and defend the notion of group agency within a life, we are in a position to see how MAC avoids both obstacles.

As we have seen, there are a variety of mechanisms that consequentialist philosophers have attempted to make use of. Mendola's preferred mechanism is direct consequentialist evaluation of overlapping group agency, which is sufficiently general enough that it may actually be used to avoid both the Demandingness Objection and the Polluter's Dilemma. The crucial mechanism is this:

We should compare a first situation in which the atomic agent achieves what it can by defection but in which the various other atomic agents that in fact constitute group agent do not constitute such an agent, to a second situation in which the group agent acts as it does and the atomic agent does not defect. If the first situation is better, then MAC says to defect. If the second situation is better, MAC says not to defect....*one should defect from a group act with good consequences only if one can achieve better consequences by the defecting act alone than the entire group act achieves.*¹⁴³

Consider first how it avoids the Demandingness Objection. The first ingredient is an acceptance of Williams's objection; the integrity of the personal point of view must be defended, and MAC appears to do this well. Mendola argues that MAC will directly support the continued existence of our own moral agency through its recognition of a

¹⁴³ Mendola 2006, 46; italics his.

persisting group agent over time. This is because better consequences result from allowing the various atomic agents of which I am composed to constitute an effective group agent over time, as opposed to allowing individual time-slices of myself to defect at will from overall projects with which I am supposed to cooperate. The worthwhile projects are the difficult, long-term ones (e.g., writing a book, earning tenure), and pulling those off would be better for me than any momentary option that might present itself. But cooperation with myself over time requires being fair to my future selves; not piling excessive demands upon them that they cannot handle seems part and parcel of having respect for myself over time.¹⁴⁴

It might be the case that, if I have excess time and effort not being devoted to some beneficent group act, MAC would recommend that I join a beneficent group act.¹⁴⁵ Perhaps MAC is more threatening to the integrity of the personal point of view than I suggest here. I don't believe this is a serious problem, however. It will be necessary for my personal integrity to reserve the right to occasionally have some free time for my own projects. Perhaps I don't care about this; perhaps I am an altruistic workaholic who loves nothing more than building houses 14 hours per day with Habitat for Humanity, and I don't perceive the work as a threat to my integrity because I identify so completely with the view. In that case, I would not regard the work as demanding, and MAC would not tell me to defect. But if I did regard that kind of work as extremely demanding, MAC would underwrite the effort to avoid being fatigued by promoting good works. It would do so by the afore-mentioned mechanism of promoting my sense of myself as an effective agent over time. MAC might not allow me to similarly defect from an

¹⁴⁴ Ibid., 94-96.

¹⁴⁵ Mendola (personal correspondence) suggests this problem for his own view.

immediately pressing group act involving a catastrophe, such as saving as many lives as possible from the wreckage in the wake of a severe outbreak of tornadoes. In that case, the need to be beneficent is especially weighty, and complaining that I need some time for myself seems especially inappropriate. But in ordinary circumstances, I will be allowed time to pursue personal projects that enhance my own sense of effective agency over time. So MAC strikes a balance between defending the personal point of view and recommending beneficent acts. The same general mechanism by which MAC recommends beneficent group acts, the mechanism of overlapping group agency, is the same general mechanism by which it preserves us from the Demandingness Objection.

This mechanism furthermore governs our choices about the kinds of group actions we join. MAC requires that we join group agents that have sufficiently good consequences, stay inside beneficent group agents that we are already a part of, and defect from group agents whose directly-evaluated actions are negative overall if we can achieve some good that would not have otherwise been achieved by the negative group act. For the first two sorts of group agents, Mendola claims that there is a weighty group agent which holds that people ought to give at least 2.5% of their income to charity. This group agent derives from traditions of beneficent alms giving rooted in several religious traditions' understanding of what obligation to the needy requires.¹⁴⁶ For the latter sort of group agent, consider a sort of situation in which you find yourself as an employee of a company taking advantage of the desperately needy in a far-off land. You have access to the secret archives that conclusively demonstrate the company's wrongdoing. The value of the company's involvement in the far-off land is a net negative. You can defect and

¹⁴⁶ Ibid., 98-102.

bring down the entire operation. MAC says to defect, as you can accomplish more good with this action than the entire group agent put together.

And MAC easily deploys the very same machinery to avoid the Polluter's Dilemma, in the following ways.

First, there are various group agents that have as their goal the avoidance of the catastrophe of climate change. If the group agents successfully act, then much good will be the long-term result. Human beings will be able to sustainably enjoy a higher quality of life for a much longer period of time, instead of rapidly depleting the earth's resources and ecosystems with the attendant deleterious effects on human health and society. There are also various group acts that do not respect the goal of avoiding catastrophic climate change. While these group acts will often generate a great deal of consequentially weighty outcomes, alternative group acts that accomplish the same goal *and* that are better with respect to the environment will often be available. To get this distinction in clearer view, consider the group act of demanding, producing, and consuming paper products such as facial tissue, toilet paper, paper towels, and sheets of paper for office use. These group acts may be carried out with no heed paid to the sustainable forestry practices of the paper products companies, or they may be carried out in a way that promotes sustainable forestry. The second way of carrying out the group acts gets us what we want, and preserves the environment for the long term. This results in better consequences overall. MAC recommends that we join the second kind of group act.

Second, once inside a group act that is beneficent with regard to the environment like the one just described above, MAC recommends that we not defect from it except in extremely exceptional circumstances. We generally cannot generate as much good as the

entire group act of non-pollution when we individually defect to pursue our own projects that are non-optimal from the point of view of the environment and the beings dependent upon it.

Third, if you are part of group acts whose actions are harmful to the environment and you could achieve some good by defecting, you should defect from the harmful group act.¹⁴⁷ And the good consequences that could result from the successful pursuit of the minimization of global climate change's harsh effects are extremely weighty. There is little good excuse, according to MAC, to defect from pursuing environmentally sound action. MAC thus avoids the Polluter's Dilemma.¹⁴⁸

MAC, then, represents a real improvement over previous forms of consequentialism. The theoretical machinery that sustains it is complicated, and it retains several controversial features -- its insistence on direct evaluation of group acts, its metaphysics of group agency, and its hedonistic value theory, for example. MAC may prove to have faults. But it meets the constraint on self-defeat that I set out in previous sections of this chapter, and it solves both the Demandingness Objection and the Polluter's Dilemma with the same piece of theoretical machinery. We have, then, at least some reason to prefer MAC to rival forms of consequentialism.

¹⁴⁷ This formulation reflects the version of MAC on offer in Mendola 2006. A forthcoming version of MAC requires that we defect from all harmful group acts whatsoever, without regard to what kind of good we could accomplish by our defection (Mendola forthcoming).

¹⁴⁸ Ibid., 308-311.

References

- Abdel Haleem, M.A.S., trans. 2004. *The Qur'an*. Oxford: Oxford University Press.
- Arneson, Richard. 2005. "Sophisticated Rule Consequentialism: Some Simple Objections." *Philosophical Issues* 15:235-251.
- Baitulmaal. 2011. "About Zakat." Accessed May 22, 2011. <http://www.baitulmaal.org/us/about-zakat>.
- Bentham, Jeremy. 1948. *An Introduction to the Principles of Morals and Legislation*. New York: Hafner.
- Brandt, Richard. 1963. "Toward a Credible Form of Utilitarianism." In *Morality and the Language of Conduct*, edited by Hector-Neri Castaneda and George Nakmikian. Detroit: Wayne State University Press.
- . 1979. *A Theory of the Good and the Right*. Amherst, NY: Prometheus Books.
- Broome, John. 1991. *Weighing Lives*. Oxford: Basil Blackwell.
- Brown, Campbell. 2011. "Consequentialise This." *Ethics* 121:4, 749-771.
- Callicot, J. Baird. 1980. "Animal Liberation: A Triangular Affair." *Environmental Ethics* 2, 311-338.
- Das, Ramon. 2000. "Prerogatives Without Restrictions?" *Philosophical Studies* 99:3, 347-371.
- Dreier, James. 1993. "The Structure of Normative Theories." *The Monist* 76, 22-40.
- . 2011. "In defense of consequentializing." In *Oxford Studies in Normative Ethics*, vol. 1, edited by Mark Timmons. New York: Oxford University Press.
- Gauthier, David. 1985. *Morals By Agreement*. New York: Oxford University Press.
- Gibbard, Alan. 2008. *Reconciling Our Aims: In Search of Bases for Ethics*. New York: Oxford University Press.
- Gilbert, Margaret. 1996. "Walking Together." In *Living Together*. Lanham, MD: Rowman & Littlefield, 177-194.
- Jackson, Frank. 1987. "Group Morality." In *Metaphysics and Morality*, edited by Phillip Pettit, Richard Sylvan, and Jean Norman. Oxford: Blackwell, 91-110.

- Harsanyi, John. 1977. "Rule Utilitarianism and Decision Theory." *Erkenntnis* 11, 25–53.
- Hooker, Brad. 1998. "Rule-Consequentialism and Obligations toward the Needy." *Pacific Philosophical Quarterly* 79, 19–33.
- . 2000. *Ideal Code, Real World: A Rule-consequentialist Theory of Morality*. New York: Oxford University Press.
- . 2005. "Reply to Arneson and McIntyre." *Philosophical Issues* 15, 264–281.
- Horton, Keith. 2011. "Fairness and Fair Shares." *Utilitas* 23:1, 88–93.
- IPCC. 2007. "Synthesis Report Summary for Policymakers." *Climate Change 2007: Synthesis Report. Contribution of Working Groups I, II, and III to the Fourth Assessment Report of the Intergovernmental Panel on Climate Change*, edited by the Core Writing Team, R.K. Pachauri, and A. Reisinger. Print version: IPCC: Geneva, Switzerland. Online version: http://www.ipcc.ch/publications_and_data/ar4/syr/en/spm.html. Accessed June 5, 2011.
- Jamieson, Dale. 1981. "Rational Egoism and Animal Rights." *Environmental Ethics* 3:2, 167–171.
- . 2007. "When Utilitarians Should Be Virtue Theorists." *Utilitas* 19:2, 160–183.
- Kagan, Shelly. 1984. "Does Consequentialism Demand Too Much?" *Philosophy and Public Affairs* 13, pp. 239–254.
- . 1989. *The Limits of Morality*. New York: Oxford University Press.
- Korsgaard, Christine. 1996. *Creating the Kingdom of Ends*. New York: Cambridge University Press.
- Mason, Elinor. 2007. Review of *Goodness and Justice: A Consequentialist Moral Theory*, by Joseph Mendola. *Notre Dame Philosophical Reviews*, August 4, 2007. Accessed April 29, 2011. <http://ndpr.nd.edu/review.cfm?id=10583>.
- McNaughton, David, and Rawling, Piers. 1991. "Agent-Relativity and the Doing-Happening Distinction." *Philosophical Studies* 63:2, 167–185.
- Mendola, Joseph. 1986. "Parfit on Directly Collectively Self-Defeating Moral Theories." *Philosophical Studies* 50:1, 153–166.
- . 2006. *Goodness and Justice*. New York: Cambridge University Press.

- . Forthcoming. *Human Events*.
- Mulgan, Tim. 2001. *The Demands of Consequentialism*. New York: Oxford University Press.
- Murphy, Liam. 1993. "The Demands of Beneficence." *Philosophy & Public Affairs* 22:4, 267-292.
- . 1997. "A Relatively Plausible Principle of Beneficence." *Philosophy & Public Affairs* 26:1, 80-86.
- . 2000. *Moral Demands in Nonideal Theory*. New York: Oxford University Press.
- Nicholls, R.J., P.P. Wong, V.R. Burkett, J.O. Codignotto, J.E. Hay, R.F. McLean, S. Ragoonaden and C.D. Woodroffe. 2007. Coastal systems and low-lying areas. *Climate Change 2007: Impacts, Adaptation and Vulnerability. Contribution of Working Group II to the Fourth Assessment Report of the Intergovernmental Panel on Climate Change*, edited by M.L. Parry, O.F. Canziani, J.P. Palutikof, P.J. van der Linden and C.E. Hanson. Print version: Cambridge University Press, Cambridge, UK, 315-356. Online version: http://www.ipcc.ch/publications_and_data/ar4/wg2/en/ch6s6-4-2.html. Accessed October 18, 2010.
- O'Neill, Onora. 2007. "Kantian Approaches to Some Famine Problems." In *Ethical Theory: An Anthology*, edited by Russ Shafer-Landau. Oxford: Blackwell, 553-564.
- Parfit, Derek. 1971. "Personal Identity." *Philosophical Review* 80, 3-27.
- . 1984. *Reasons and Persons*. New York: Oxford University Press.
- Pettit, Phillip. 1997. "The Consequentialist Perspective." In *Three Methods of Ethics*, edited by Marcia Baron, Phillip Pettit, and Michael Slote. Oxford: Blackwell.
- Plato. 1997. *Theatetus*. Translated by M.J. Levett, revised by Myles Burnyeat. In *Plato: Complete Works*, edited by John M. Cooper and D.S. Hutchinson. Indianapolis: Hackett Publishing Company.
- Portmore, Douglas. 2001. "Can an Act-Consequentialist Theory Be Agent-Relative?" *American Philosophical Quarterly* 38, 363-377.
- . 2005. "Combining Teleological Ethics with Evaluator Relativism: A Promising Result." *Pacific Philosophical Quarterly* 86, 95-113.
- . 2011. *Commonsense Consequentialism: Wherein Morality Meets Rationality*. New York: Oxford University Press.

- Rachels, James. 1979. "Killing and Starving to Death." *Philosophy* 54, 159-171.
- Railton, Peter. 1984. "Alienation, Consequentialism, and the Demands of Morality." *Philosophy & Public Affairs* 13:2, 134-171.
- Rawls, John. 1971. *A Theory of Justice*. Cambridge, MA: Harvard University Press.
- Rolston, Holmes, III. 1999. "Respect for Life: Counting What Singer Finds of No Account." In *Singer and his Critics*, edited by Dale Jamieson. New York: Oxford University Press.
- Sagoff, Mark. 1984. "Animal Liberation and Environmental Ethics: Bad Marriage, Quick Divorce." *Osgood Hall Law Journal* 22, 297-307.
- Scheffler, Samuel. 1992. "Prerogatives Without Restrictions," *Philosophical Perspectives* 6, 377-97. Reprinted in Scheffler, *The Rejection of Consequentialism*, revised edition, New York: Oxford University Press, 1994, pp. 167-192.
- . 1994. *The Rejection of Consequentialism*, revised edition. New York: Oxford University Press.
- Schroeder, Mark. 2006. "Not So Promising After All: Evaluator-Relative Teleology and Common-Sense Morality." *Pacific Philosophical Quarterly* 87, 348-356.
- . 2007. "Teleology, Agent-Relative Value, and 'Good'." *Ethics* 117, 265-295.
- Shafer-Landau, Russ. 2010. *The Fundamentals of Ethics*, first edition. New York: Oxford University Press.
- Singer, Peter. 1972. "Famine, Affluence, and Morality," *Philosophy & Public Affairs* 1, pp. 229-243.
- . 2009. *The Life You Can Save: Acting Now to End World Poverty*. New York: Random House.
- . 2011a. "The idea." Accessed March 25, 2011.
<http://www.thelifeyoucansave.com/idea>.
- . 2011b. "How much?" Accessed March 25, 2011.
<http://www.thelifeyoucansave.com/calculator>.
- Skorupski, John. 1995. "Agent-Neutrality, Consequentialism, Utilitarianism...A Terminological Note." *Utilitas* 7, 49-54.

- Slote, Michael. 1985. *Common-Sense Morality and Consequentialism*. London: Routledge & Kegan Paul.
- Sobel, David. 2007. "The Impotence of the Demandingness Objection." *Philosopher's Imprint* 7 (8). Accessed Jan. 26, 2011. <http://www.philosophersimprint.org/007008/>.
- Tedesco, Matthew. 2011. "Intuitions and the Demands of Consequentialism." *Utilitas* 23:1, 94-104.
- Unger, Peter. 1996. *Living High and Letting Die: Our Illusion of Innocence*. New York: Oxford University Press.
- United Nations Development Programme. "Fast Facts – Climate Change." Accessed September 14, 2011. <http://www.undp.org/climatechange/docs/English/FF-climate.pdf>.
- Williams, Bernard. 1973. "A Critique of Utilitarianism." In *Utilitarianism: For and Against*, edited by J.J.C. Smart and Bernard Williams. New York: Cambridge University Press.
- Wolf, Susan. 1982. "Moral Saints." *Journal of Philosophy* 79, 419-439.