

Consequentialism and the Demandingness Objection

By Jeff Heikkinen

**A Thesis submitted to the Faculty of Graduate Studies of
The University of Manitoba
in partial fulfilment of the requirements of the degree of
MASTER OF ARTS**

Department of Philosophy, University of Manitoba, Winnipeg

© 2007 Jeffrey Wayne Heikkinen

Abstract

Demandingness-based objections to utilitarianism and other consequentialist moral theories constitute the most important problem facing moral philosophers today. In this Thesis, I offer an explanation of what makes the demandingness objection compelling, namely, that utilitarianism alienates us from the projects and goals that define us as individual human beings (normally taken to be a *separate* objection). This suggests that solving the problems demandingness considerations present involves carving out a space for these projects and goals alongside the demands of a consequentialist morality; thus, we have two nearly independent sources of normative reasons, and the real question is how they interact. Various suggestions for answering this question are considered and rejected. I also discuss how Alastair Norcross' scalar utilitarianism "solves" the demandingness problem, what the costs of this solution are, and how it might be integrated into a theory concerning the aforementioned interaction.

Acknowledgments

Though authored by me (apart from borrowing a phrase here and there from some of those thanked below), there are many others without whom this thesis would not exist.

Between them, Tim Pechey, Zane Zalis, and Sam Kozlowich jointly convinced me that I wanted to do *some* sort of teaching, though only Sam consciously pushed me in that direction. I have no idea what I would be doing with my life if not for their inspiration. Sadly, I shall only have the opportunity to thank Zane personally; Sam passed away in 1998 and Tim in 2000.

My interest in philosophy was always there; every child asks numerous philosophically relevant questions about the world, and I, like most philosophers, merely failed to *stop* doing so. But my interest in philosophy *qua* academic discipline owes much to Carl Matheson. I hope I can instil *half* the enjoyment of the subject in my own future students that he does in his, including myself.

I likely would have drifted in the direction of ethics on my own, but in the event, Rob Shaver was the most immediately relevant factor in pushing me in that direction. When I was an undergraduate, the clarity and precision of his presentation of even some of the most complex material in the field did more than any one other thing to demolish my perception of ethics as a relatively “soft” subdiscipline of philosophy, with a lower intellectual bar than, say, philosophy of science. More recently, acting as my advisor on the present work, he has provided wave after wave of incisive comments without which this work would be completely unrecognizable, and far inferior to what it is. Whatever merits it has are Rob’s just as much as they are my own.

Rob was by no means the only such direct influence on the final form of this thesis, only the single largest. Joyce Jenkins’ perceptive multipage commentary required a rethink of not only many individual parts of this essay, but its very structure. Joyce, along with Margaret Ogradnick, also demonstrated the need for many important clarifications, most merely in my writing, but some in my very *thinking* as well. And much inspiration, as well as many examples and ideas (mostly, but not exclusively, in Chapter 4), owe an incalculable debt to two seminars and a number of personal discussions with Alastair Norcross. I hope Alastair can forgive the utter mangling his views received in earlier drafts of this work, and that with the present version, I have finally done them justice.

Other professors (not all of them philosophers) who have played a significant role in the development of my love of academia in general, philosophy in particular, and ethics in even-more-particular include, in roughly chronological order by my own

first encounter with them, Michael Stack, Wayne Nielsen, Bob Bright, Ken Warmbröd, Sinclair McRae, Arthur Schaeffer, Michael Feld, George Toles, Gerry Bowler, the late Fred Stambrook, Tina Chen, David Gunderson, Ben Caplan, Rhonda Martens, and most recently (and since this thesis was underway) Richard Grandy, George Sher, Nicoletta Orlandi, Melinda Fagan, and perhaps most of all, Baruch Brody.

Friends who have provided their support, not to mention influenced my thinking in some way, over the years, some formally trained in philosophy but most not, again arranged more or less chronologically, include but are almost certainly not limited to my mother Eleanore Chestley, my father George Heikkinen, my sister Lenora Rose Patrick, my stepmother Lolly Heikkinen, Tanya Misseghers, Mike Major, Holly “Branwen” Finnen-Stewart, Laurie Smith, Brenna Bacchus, Keith Grant, Adam Bacchus, Colin Patrick, Laurel Ralston, Jacob Kolman, Dani Wenner, Anthony Carreras, Ashley Allen, Jessie Armendt (who was particularly important in showing me that utilitarianism could be *inspiring*), and Lauren Thompson.

My apologies, and their fair share of gratitude, go out to anyone unintentionally omitted from either of the previous two lists. My shortcomings as a person are many, and while I like to think ingratitude is not one of them, forgetfulness most certainly is.

Jeff Heikkinen
Houston, Texas
November 17, 2007

Amended April 1, 2009

A dedication

Dedication on page iv as of April 1, 2009.

Table of Contents

CONSEQUENTIALISM AND THE DEMANDINGNESS OBJECTION.....	I
ABSTRACT.....	II
ACKNOWLEDGMENTS.....	III
A DEDICATION.....	V
TABLE OF CONTENTS.....	VI
INTRODUCTION.....	1
I – THE PROBLEM.....	6
CONSEQUENTIALISM.....	6
THE DEBATE ON CONSEQUENTIALISM.....	14
II – ALIENATION.....	21
WHY IS DEMANDINGNESS AN OBJECTION?.....	21
THE EXCESSIVE AND THE ALIENATING.....	30
III – SCHEFFLER-LIKE VIEWS.....	38
THE HYBRID THEORY.....	38
BEYOND THE HYBRID THEORY.....	46
IV – SCALARITY.....	58
AGAINST “BRIGHT LINE” VIEWS.....	59
SCALARITY AND THE RIGHT.....	64
THE RESTRICTED DELTA VIEW.....	69
SCALARITY AND SPECIFICITY.....	73
NON-COERCIVE MORAL THEORY.....	76
CONCLUSION.....	79
APPENDICES.....	82
APPENDIX A: METAETHICAL AND METHODOLOGICAL ASSUMPTIONS.....	82
APPENDIX B: CONTEXTUALISM?.....	94
BIBLIOGRAPHY.....	100

Introduction

What is required of us, if we wish to live morally decent lives? There is great diversity of opinion on the subject. Many philosophers, going back at least¹ to Jeremy Bentham and John Stuart Mill – and some few non-philosophers, in my experience mostly military sorts – have defended a view called *utilitarianism*, which is nowadays understood as a particularly important special case of a broader view, *consequentialism*. Utilitarian, and most other extant consequentialist, moral theories seem to be committed to a particularly demanding answer to this question. On these theories, at least in the forms most familiar today, we must always perform the action that produces the most good out of those available, without regard for the cost to ourselves (at least, not out of proportion to our own share of the overall good).²

This is often held to be too high a standard; surely, it is objected, acting optimally is *nice* of us, but any serious look at what such behaviour would entail will quickly

¹Or maybe going back *much* further. Plato has the character Socrates appear to defend a utilitarian moral theory in the *Protagoras* (at 351b-354e). This has sometimes been taken to mean that Plato (or the historical Socrates, at least in the relatively early years when the conversation depicted there is supposed to have taken place) held a utilitarian view. To me, however, this seems unlikely, or at least not well supported by this evidence. It is not altogether clear to me that Socrates is defending his own view rather than playing Devil's advocate in these passages, and even if he is doing the former, it is not a given that there is intended to be continuity of philosophical views between the Socrates of any given dialogue and either Plato or the real Socrates (or the Socrates of any given *other* dialogue, for that matter). This point seems especially relevant here as it is my understanding that these views are not mentioned anywhere else in Plato's works.

Still, these passages do at the very least show that someone had *thought of* utilitarianism long before Bentham.

²These views also disregard any notion of rights (which Bentham called “nonsense upon stilts”) or other *constraints* on morally justifiable action.

reveal that it goes far beyond what most of us think can possibly be *required* of us. (Maybe Mother Teresa met this standard, some of the time. Certainly you and I don't.) And this is widely held to render consequentialism, at least in these forms, untenable as a moral theory.

I believe that this *demandingness objection* to consequentialist moral theories is the most important issue facing moral philosophers today – particularly those, such as myself, who are generally sympathetic to consequentialism. If the objection is correct then, assuming utilitarianism and its cousins really are as demanding as it suggests, an important family of moral theories stands revealed as false, or at the very least in need of serious rethinking. If the objection is not tenable, but the demands it attributes to utilitarianism are a real feature of that theory, then this has enormous implications for our moral obligations. And it seems to me that utilitarianism, at least in the optimizing form I have so far discussed, *does* in fact make these demands, and so on the face of it, one or the other of these two possibilities must be true – and either one would, if generally accepted, change the face of moral philosophy.

This work, then, is a long essay on the demandingness objection. I have seen fit to divide it into four chapters, and each of those into two or five sections.³

³ There are also two appendices, which could with little modification be essays in their own right. Appendix A addresses what my methodology presupposes, and might be best read at some point before reading chapter 2, though it is not essential to understanding anything in the main text. Appendix B discusses ethical contextualism, a subject germane to chapter 4, but which I could not fit elegantly into

In chapter 1, I will explain the objection in a little more detail, by means of attempting to place it in the larger context of the debate on consequentialist moral theories (the latter is mainly for the benefit of readers with limited or no philosophy backgrounds). This will merely set up the rest of the work.

Chapter 2 attempts to address the question of what, exactly, is objectionable about “excessive” demands, a crucial issue that received little discussion in the literature until fairly recently.⁴ A plausible answer, I will suggest, closely resembles what is normally taken to be a distinct objection to consequentialist theories – that they alienate us from too much of what we value in our own lives, failing to show due respect to our distinctive individual goals and projects.⁵ (I will not be directly arguing that this is the *best* possible answer, much less the only such answer, though I suspect it is the former.) Along the way I will pretty much give up on (optimizing) utilitarianism as a theory with much hope of being able to meet this objection.

Chapter 3 will be taken up by a search for a theory that does meet the objection while still remaining committed to what I take to be the core principles of

that chapter.

⁴ I may not have given the matter much thought myself if not for Liam Murphy’s recent book, *Moral Demands in Nonideal Theory* (Oxford: Clarendon, 2000). Despite his impressing the importance of this question upon me, however, this essay does not directly discuss Murphy and goes off in a very different direction from his book.

⁵Perhaps this is also the grain of truth in the otherwise bizarre objection that utilitarianism somehow ignores the separateness of persons.

consequentialism (what I shall call PTR, the *pro tanto* reason to promote the good,⁶ combined with the claim that no other reason matters morally). This search will take as its starting point and main inspiration the *hybrid view* introduced by Samuel Scheffler's *The Rejection of Consequentialism*,⁷ which is an earlier attempt at much the same goal I am trying to achieve. I will quickly find reason to reject the hybrid view as even more problematic than optimizing utilitarianism, yet accept that there is *something* basically right about the impulse behind it. The remainder of chapter 3 after that point will be taken up by an unsuccessful, but hopefully instructive, search for a theory that captures both PTR and that basically-right element of the hybrid view.

Chapter 4 will introduce a theory that, even if in a somewhat deflationary manner, *succeeds completely* at this goal – namely, the *scalar view* whose best-known defender is currently Alastair Norcross. This view holds that utilitarianism tells us only that some actions are better than various alternatives, but stops short of telling us we *must* do the better actions and avoid the worse. Yet, I will not be suggesting that the scalar view allows us to simply declare the project a success and have done with it. Rather, its introduction will occasion a hard look at how independently plausible the scalar view actually is. The feature that allows it to reconcile PTR with the alienation concerns of Scheffler and others is a lack of action-guidingness, and it is at the very

⁶This name is originally due to Shelly Kagan.

⁷Samuel Scheffler, *The Rejection of Consequentialism* (Oxford: Clarendon, 1982 [revised 1994]), 15-22.

least forgivable if this occasions scepticism about the point of scalarity as a moral theory. Much of chapter 4 will be taken up by attempts, with varying degrees of success, to allay these concerns. I will introduce a promising and, to the best of my knowledge, original bit of theory I refer to as the *restricted delta view*, but that will founder; I may develop it further in future work, but as of the end of the present essay it will have been (reluctantly) set aside. But I will go on to show that, in principle, it is *possible* to accept the scalar theory of *morality* and still incorporate that theory into a theory of *normativity in general* which does give us the specific advice the scalar view does not. I will close with some considerations that suggest that the scalar view's lack of action-guidance may not, in the final analysis, be such a bad thing. All of this seems to point to the following claim – that the scalar view offers some promising directions on reconciling utilitarianism and the demandingness objection, yet that work remains to be done before it is clear we should accept the scalar view. If this leaves things somewhat uncertain at the end of this essay, at least I have not seen the need to end it in complete despair, as a famous utilitarian⁸ once concluded a work with similar goals.

⁸I refer, of course, to Sidgwick's *The Methods of Ethics*.

I – The Problem

Consequentialism

One important family of moral theories goes by the name **consequentialism**. As the name suggests, these theories (as they are generally understood) judge the rightness and wrongness of actions by their consequences – the state of the world that results from them. In slightly more technical language, a consequentialist – again, as most people use the term – is one who holds that the right is best understood in terms of the good.

I should explain what this means in slightly more detail. People in general, and philosophers in particular, differ considerably on which things they consider good, but they share enough common ground to discuss the subject meaningfully. Nearly everyone thinks that it is good, for example, if people are cured of potentially fatal illnesses, bad if people become injured. This agreement will be enough to get us started.

It follows from these points of general agreement that, all else being equal, a world in which people who now have fatal illnesses are cured of them is superior to a world in which they are not (though inferior, perhaps, to one in which they never had those illnesses in the first place). To the consequentialist, such comparisons

define, or failing that, at least uniquely determine, all of morality. States of affairs⁹ are ranked according to such comparisons, and the consequentialist goes on to claim that, at least ideally, we ought to act¹⁰ so as to bring about the most highly ranked state of affairs we find ourselves in a position to bring about. Most simply, a consequentialist might claim that one action is preferable, morally, to another if and only if the first action leaves the world in a state in which there is a greater net balance of the good things minus the bad things, compared to the second action. Even if a particular consequentialist values something that does not at first glance seem strictly quantitative, such as equality,¹¹ the ranking is in any event based only on features of the states of affairs being compared, as opposed, for example, to how the state of affairs was brought about (save to the extent that this may, itself, be considered a good or bad feature of them).¹²

⁹ I use “state of affairs” and “state of the world” interchangeably, and the “consequences” of an action shall mean the state of affairs/the world that results from performing that action (including the fact that that action was performed).

¹⁰ I shall assume that individual acts are the loci of moral judgment throughout this work.

¹¹ More precisely, there are at least two ways to value equality, as a good in its own right (so that greater equality - and there are statistical methods by which this can be measured - means that, *ceteris paribus*, a state of affairs contains more good) or as a distributive principle (a constraint, independent of the *amount* of good in a given world, on how it is acceptable to distribute that good). The former version, which is unfamiliar but seems perfectly coherent and possibly even attractive, is the one I am saying is compatible with consequentialism; the latter, which is the more familiar approach, obviously is not. The extent to which the worlds endorsed by these views will differ *in practice* is not a subject I will be discussing here and, at any rate, I am not even sure whether the question is best answered by philosophers or statisticians.

¹² I place any view that values only states of affairs under the consequentialist umbrella; in other words, I use “consequentialist” the way Rawls and Scanlon use “teleological”. However, the view I use to represent consequentialism most of the time - namely utilitarianism - will be consequentialist in the narrower sense more

Another way to define consequentialism, hopefully one equivalent to or at least coextensive with the above, is to say a view is consequentialist if it accepts the following claim, *and* does not accept that there are any other moral principles which ever outweigh it:

PTR: All else being equal, we ought to choose the option that produces more good overall.

Kagan¹³ calls this the *pro tanto*¹⁴ **reason to promote the good** (from which my abbreviation derives). Intuitively it seems right to most of us; if someone questions it, one (momentarily) wonders whether the objector understands what “good” means.

usually found in the literature and touched on in the following paragraphs as well. That is to say, it will be consequentialist even if “goods” like equality, which do not simply depend upon summing up the amount of straightforward goods such as happiness in a given world, are excluded from consideration as consequentialist theories.

Another means of drawing this line has been suggested to me by Robert Shaver. On Shaver's view, a theory is consequentialist if, for any two states of affairs differing only in the amount of some good that they contain, it is always and everywhere preferable to bring about the world with more of it. (Note that the value or disvalue of the means of bringing it about needs to be taken into consideration – if these have moral weight of their own, then this is another respect in which the states of affairs differ). If it is possible to construct a counterexample to this suggested principle under a moral theory, that theory is not consequentialist. I find this suggestion plausible and believe, or hope, that it admits all and only views that would fall under my own broad definition.

¹³ Shelly Kagan, *The Limits of Morality* (Oxford: Clarendon), 1989, 16-17.

¹⁴Kagan borrows this term from law, explicitly choosing it over the more common *prima facie*. To call something a *prima facie* reason suggests something which only appears to be a reason, but might not turn out to actually be one in some cases; *pro tanto*, according to Kagan (*Limits* 17), connotes a reason which is always present, but leaves open the possibility that it is sometimes outweighed.

Contrast this with the view that one *simply ought not* to do a thing – lie or steal, say. One who accepts consequentialism, by my or almost any other definition, is committed to lying if doing so will reduce the overall amount of lying that takes place, even slightly, assuming a theory of the good that directly or indirectly claims lying has disvalue in the first place. Non-consequentialists tend not to agree. Though very few follow Kant in saying one ought not to lie even if someone’s life is at stake, such thinkers nevertheless claim that you need a better reason than, say, the mere prevention of two other such lies.

Kagan takes the issue between the consequentialist and the non-consequentialist to be primarily a dispute over when the “all else” PTR mentions is equal; the consequentialist (Kagan uses the term **extremist**,¹⁵ to emphasize his unapologetic embrace of the demandingness that will be described momentarily) says always, while other moral theorists think there are cases, perhaps quite a lot of them, where PTR is outweighed by other considerations. In making this claim, Kagan certainly gets *some* non-consequentialists right (**moderates**, in Kagan’s terminology, those who defend what he calls “ordinary morality”), but it is open to others to deny that the PTR has any force at all. Perhaps we could follow Scanlon¹⁶ in assigning moral

¹⁵ I am being a little fast and loose here. For our purposes, though, it will be enough to define an extremist as a consequentialist who accepts a very demanding version of the theory, such as that briefly outlined in the next section.

¹⁶ T.M. Scanlon, *What We Owe to Each Other* (Cambridge, MA: Belknap), 1998, especially the notorious discussion of aggregation (see, for example, the solution to the drowning problem presented at 232 to see what seems to be this kind of reasoning in action). Scanlon, like many non-consequentialists, runs into the problem

relevance only to the single largest complaint that would be produced by a given action. But in either case, the non-consequentialist does not give such a central place to PTR as the consequentialist; and this central role can be taken to define consequentialism. We can describe a given version of consequentialism, then, by stating PTR, together with the supplementary claim that no other principle ever outweighs it, and then going on to answer the question “how do we define what is good?”.

Let me introduce another bit of terminology that will be useful later, chiefly in chapter 3. A “pseudo-consequentialist” is one who agrees that PTR is centrally important to morality, but supplements it with additional, non-consequentialist that his theory appears to imply that we should be neutral between, for example, saving one person from certain death in some disaster and saving two others from the same disaster, should we be unable to do both – a case where consequentialist theories clearly give us the intuitively-right answer. He handles this by adding the following wrinkle: if we find rescuing the smaller group (in this case the one) acceptable, then one arbitrarily chosen member of the group of two can say that, not only will he die if we proceed in this fashion, but he is also being treated as though he were not there – his presence is making no difference to our decision. This seems like reasonable grounds to reject the principle on which we are acting – and one very useful way to cash out *that* claim is to say that this makes his complaint “larger” than that of the one person we would leave to die by rescuing the pair of which the complainant is a member. On this way of understanding Scanlon, if (but probably not *only* if) some act would ignore the “largest” such complaint any of our available options would generate, we can reasonably reject the principle behind that action, and for that reason, we can’t do an action which would ignore such a complaint. While it has much initial plausibility, this move has opened Scanlon up to many objections (there is no need to discuss them here) and he no longer defends this account of aggregation.

I admit that the emphasis on the “size” of the complaint is somewhat idiosyncratic. The idiosyncrasy is not originally mine but Robert Shaver's, having proved useful in his moral philosophy seminar on Scanlon.

principles. Essentially, such a thinker is a moral pluralist who thinks the consequentialist is onto something, and gets it right much of the time, but who still thinks that view is importantly incomplete. The line between the pseudo-consequentialist and the non-consequentialist is not sharp; thinkers ranging from Peter Unger to Frances Kamm have defended consequentialist thinking in particular cases and at the same time acknowledged that other principles are morally relevant. Unger is the clearest example of a pseudo-consequentialist I can think of, being just a few passing acknowledgements of special obligations to our families away from being a utilitarian. Kamm is *not* a pseudo-consequentialist by any stretch of the imagination. At least conceptually, there would seem to be a significant grey area somewhere between these two thinkers; however, this thesis will not discuss any views that fall into it. Wherever it matters, it should be quite clear whether this label applies to someone or not.¹⁷

There are many ways to produce differing versions of consequentialism, some of which I will touch upon later in this work, but historically the differences most discussed in the literature have concerned the criteria for the good. For example, **utilitarianism** (or one important version of it, at any rate) holds that pleasure in the

¹⁷ Indeed, I deal with just three such thinkers – Unger (mostly in the appendices), one who accepts Scheffler’s Hybrid Theory (a category which doesn’t actually include Scheffler), and *myself*, to whatever extent I accept my own conclusions in the present work.

broadest possible sense¹⁸ is good, pain in the broadest possible sense is bad, and that these are simply to be summed up over all the people concerned; another, perhaps more often defended today, performs these sums using the satisfaction of preferences, rather than pleasure. The theory of the good which holds that *only* pleasure and pain are intrinsically good and bad is called **hedonism**, and so we might define the former version of utilitarianism as the view that combines hedonism (as its theory of the good) with consequentialism (as its theory of the right).¹⁹

I shall not be defending a particular theory of the good in the present work, but it will not always be possible (at least without trying the reader's patience unduly) to keep the discussion entirely neutral, or even free of *controversial* assumptions, as to what the good is either. When forced to make claims that are not wholly independent of what the good is, I shall use a hedonistic framework unless otherwise stated. (When using this shortcut, I shall endeavour not to claim anything that could not, with minor and reasonably obvious changes, be applied to other

¹⁸ I do not hereby commit myself on exactly what "the broadest possible sense" is, but one candidate can be ruled out. It is doubtful that watching a well-crafted but intense and depressing play, for example, gives rise to pleasurable *sensations*, yet it is something we would want to say is a pleasure, for those who appreciate such things. Whatever pleasure is on this construal, it is not plausible that it is exclusively a sort of sensation.

¹⁹ This definition is slightly oversimplified, for we can change other features of utilitarianism – for example, optimization, or acts as the loci of moral judgment rather than, say, dispositions – and get theories that are hedonistic and consequentialist, but not utilitarian as I shall use the term.

forms of consequentialism as well, so long as their theory of the good is not too divorced from hedonism.) This seems the obvious choice, for a few interrelated reasons. Hedonistic utilitarianism is a particularly important version of consequentialism, partly for historical reasons that need not concern us, but also partly because virtually all consequentialist views take hedonism to be, if not itself the correct theory of the good, then at least an important part of that theory.²⁰ As a result, claims that are true under utilitarianism will generally be true under many other forms of consequentialism as well, and so will usually be transferable to other forms of consequentialism directly – or at least with little shoehorning – in a way that may not be true if I were to choose some other version of consequentialism as my default theory. A final point in favour of hedonistic utilitarianism as a default, though I believe it is a consequence of the other reasons I have already given, is the relative familiarity of this theory. Far more people defend, or otherwise write about, utilitarianism than any one version of consequentialism that is clearly distinct from it, and most of those have either accepted hedonism themselves or used it for ease of discussion in precisely the same way I am doing.

It is important not to allow this explanatory shortcut to create the impression that utilitarianism and consequentialism are fully interchangeable. A moral theory might grant intrinsic value to multiple things, such as knowledge or the cultivation of

²⁰ And it is my view that many proposed candidates for other intrinsic goods really are only valued or valuable because, and to the extent that, they produce or are correlated with happiness.

talents, instead of or in addition to pleasure, while holding that some good things are not necessarily good in proportion to their quantity due, for example, to distributive considerations. Such a theory would not look very much like utilitarianism, but would still be consequentialist as I use the term. But it would be neither possible nor, at least until late in my discussion, desirable to make my discussion cover *all* possible theories of which that is true.

The Debate on Consequentialism

What remains constant among all versions of consequentialism is that actions are evaluated only by the world, or state of affairs, that results from them, as compared to those that would have resulted from other possible actions. This leads directly to claims that many people find counterintuitive. For example, most consequentialists would agree that the mere fact that something is an instance of theft does not count morally against doing it,²¹ though such considerations as the despair of the victim and any inconvenience the loss of the stolen good causes him certainly would so count. It will probably turn out that, due to considerations such as these, *most* thefts are wrong after all; but this will not be *because* they are thefts, nor would this be true of a theft that went unnoticed, inconvenienced no-one, and greatly benefited the perpetrator.

²¹ It is possible to imagine a version of consequentialism on which the number of thefts that occur is one of the aspects of outcomes that we measure. I do not, however, know of such a view that is seriously defended by anyone. Most consequentialists measure good by appeal to much broader principles.

Some will consider such a derivative and non-absolute prohibition on theft (and promise-breaking, and lying, and so on even up to acts such as torture and murder) to be grounds for an objection to consequentialism and others will not; I daresay some would even regard this as a point in consequentialism's favour (I, for example, represent the latter group). But even those in the former camp must admit that consequentialism has its attractions. Most well-known consequential theories reduce morality to one, or at worst a fairly small number of, relatively simple, intuitively appealing principle(s). Consequentialism offers, at least in principle, a determinate way of deciding any moral conflict; simply see which option scores the highest on a particular consequentialist ranking (with ties, presumably, implying that any of the tied options is acceptable). The most well-known of these theories place value on many of the same things most of us do in our daily lives; if eating chocolate cake gives us pleasure (or satisfies a preference we have), the Utilitarian, at least, says all else being equal, that's what we should do. All of these points are in sharp contrast to the views held by many non-consequentialists, which end up in tangled messes of highly abstract, sometimes conflicting principles, often far removed from the way we actually make decisions in real life and generally featuring innumerable dubiously motivated special cases.²²

Stacked up against these (and other) points in consequentialism's favour are several types of objections, of which two broad categories seem particularly noteworthy to

²² I am thinking in particular of Frances Kamm's methodology here.

me. The first is that, in many respects, consequentialism allows too much. The second is that, in other respects, it demands too much.

I will not be spending much time on the first sort of objection in the current work, though it is certainly lurking in the background and will need to be mentioned again. This is of course the line of reasoning I alluded to at the beginning of this section; consequentialist theories generally want no truck with the idea of certain categories of actions *just being wrong* regardless of their results, and many will see this as a problem with consequentialism. The consequentialist says that, for example, if murdering one innocent person is what will have the best consequences (for example, if it will save five others, and no other available act will), then that is what we should do. If you think we have inviolable rights against being harmed merely to benefit others – or hold any of several other fairly common moral views – this is likely to strike you as outrageous. While I find myself unable to agree with this line of thinking, I feel just enough of an intuitive tug in that general direction to be somewhat confused on how to go about entering *either* side of this particular debate. So, to the greatest extent possible, I will refrain from even trying. In any event, the second sort of objection strikes me as more important, and yet, has been far less discussed in the moral theory literature. So for these reasons, it is on that type that I will focus in the present essay.

This other major objection is that consequentialism is too demanding a moral theory. To make the basis for this objection clear, I must bring out a feature of utilitarianism, in the forms most often seen in the modern literature, that I have not discussed in any detail so far (this feature holds for the overwhelming majority of other versions of consequentialism that have serious defenders as well). Utilitarianism does not merely ask that we produce some good, or even that we produce a certain level of it. Rather, it insists that the act, of those available to us, which produces the *most* good is right, and that anything else is wrong. I shall call this the **maximization principle**.

It may already be obvious to many readers why consequentialist theories that accept this principle are viewed as excessively demanding, but allow me to bring the point out explicitly. There are many areas of the world stricken by famine and preventable diseases, affecting massive numbers of people, disproportionately children. As it happens, in these parts of the world, Western currency goes a long way, and in any event some of the foods and medicines needed in these places are cheap by any standard. Now consider almost anyone in the Western world who leads a reasonably comfortable life – you, for example, if you’re at all typical of people who might have occasion to read a Master’s thesis in philosophy – making a decision about how to spend some portion of your income. On any plausible way of measuring such things, you could, by making suitable donations to the right charities, produce, or at least be statistically expected to produce, far more good than you could reasonably expect to

produce by spending the same money on yourself or your nearest and dearest. This means that the former is almost certainly, from a consequentialist perspective, the optimal use of these particular resources; and even if it is not (i.e. there is a still better use for them), surely it is far closer than the latter.

But this amounts to more than a mere call for Westerners to spend more on famine relief and the like, though it is certainly that; for this decision will need to be repeated for every portion of your income, and the same decision will be reached nearly every time. The exceptions will only be those strictly necessary for you to continue contributing the rest. You need to be able to feed yourself and those dependant on you, but from a utilitarian standpoint, this is true chiefly because your capacity to continue making such contributions to the overall good depends on it; that *your* happiness would suffer (indeed, cease altogether) from doing otherwise is only a tiny part of the problem with such a course of action. You are entitled to whatever expenditures are necessary to continue holding down your job, but again, only because this is what allows you to contribute to the overall good at a high level – and may not apply if you could plausibly get a higher-paying job that wasn't itself egregiously harmful to the world at large. Plausibly, a reasonable level of comforts might be needed in order to maintain your capacity to contribute; but surely this will fall far short of entitling you to the home, clothing and recreations you would ideally prefer. (Speaking for myself, a brisk walk is free and often serves admirably to

maintain my own capacity to work, such as it is.) Already this is looking very unlike the way most of us live our lives.

But I have barely even begun to describe the difference this would make in most of our lives. It is plausible that spending time at our jobs would be permitted, even required, as this allows us to gather the resources needed in order to make our contributions to the overall good. But what about our *choice* of job? Of hobbies and recreations? Of foods? Of friends? Almost all of our decisions would need to be made in light of their contribution to the overall good, and I imagine that for most of us, very few of them would survive in any recognizable form.

This, then, is what the **demandingness objection** is reacting against. In its simplest possible form, this objection simply states that if a moral theory demands what consequentialism, on the above account, must demand, then something has gone wrong; that such a theory cannot possibly be acceptable. This is, of course, justified and further elaborated in a variety of ways. This thesis will consider a fairly new account of what justifies a demandingness-based objection. It is my view that consequentialists should take the demandingness objection seriously. Though of primarily consequentialist sympathies myself, I will spend most of this thesis playing Devil's advocate, forcefully presenting the non-consequentialist case and exploring various consequentialist and pseudo-consequentialist responses. All of

these will, unfortunately by my lights, come up short. My ultimate goal in doing this is to clarify the issues so that such a consequentialist response to the demandingness objection can be found (or failing that, one that is largely in a consequentialist spirit, in the sense that, for example, Scheffler's hybrid theory, discussed in chapter 3, could be said to be consequentialist in spirit). But honesty compels me to point out that, as far as the current work is concerned, I have not found much reason to maintain my current optimism about the prospects of finding such a response.

II – Alienation

Why is demandingness an objection?

So far, I have simply assumed that both consequentialism and the demandingness objection to it are plausible, which seems to leave us at an impasse. For the discussion to continue in any reasonably interesting manner, I need to say more about at least one of these, my two main subjects. And a natural place to start would be to discuss their justification.

I shall not be offering much in the way of justification for consequentialism, simply noting that PTR, which is the core of (though not unique to) consequentialism, seems to be independently plausible. The strength of its intuitive support is high, it yields to nothing in its generality, and it seems to be very resilient against attempts to explain it away nonmorally.²³ We can proceed on the assumption, then, that PTR is a plausible part of any complete moral theory. I have noted that there are significant thinkers, such as Scanlon, who disagree; I acknowledge this and hope to address the problems they raise in future work, but for now I simply invite the reader to join me in taking PTR as given, an invitation which I imagine most such readers, including many non-consequentialists, will in any case be inclined to accept.

²³ In appendix A, I sketch out a case that these are the factors we should be looking at when evaluating such things.

So let us turn to how the demandingness objection might be justified – what, to break this question into its component parts, might be the thought underlying it, and whether, to what extent, and why we ought to take that thought seriously. That doing what some theory tells us to do is too demanding does seem like a compelling reason why we should not *like* for it to be binding upon us; but this is essentially a selfish thought, at least on its face, so we might not find it immediately clear why this should seem like a compelling *objection* to a theory, much less a *morally* compelling one. The consequentialist, in other words, can be forgiven for being tempted to say at this juncture – if consequentialism is a more demanding moral theory than somebody wishes to follow, so much the worse for that person; this shows their lack of moral fortitude, not any problem with consequentialism.

What might the non-consequentialist (or, indeed, the Devil's advocate of a consequentialist who does feel the intuitive pull of the demandingness objection, say, the present author) say in response to the above line of reasoning? To just flatly deny the above consequentialist claim is tempting, but unsatisfying; what we need to do, then, is show some *morally* compelling grounds for accepting that being too demanding is a real objection to a theory.

A superficially attractive first attempt at an answer might be – ought implies can, and we as human beings cannot live up to the moral standard utilitarianism and cognate theories demand of us, therefore there must be something flawed about

them as theories of what we ought to do. But the problem with this suggestion is obvious. The utilitarian need not say we act wrongly if we fail to do the impossible; he can simply say we act wrongly if we fail to do the best we can. So this cannot be the answer.

And even if this fatal problem did not exist, the suggestion wouldn't quite capture the impulse behind the demandingness objection. Suppose you receive a request for a \$100 donation to some agency devoted to fighting famine and disease in far-off parts of the world.²⁴ Most of us want to say, as utilitarianism claims we may not, that we are at least sometimes allowed to simply keep our money, even if we grant that it is in some sense *better* morally to contribute it to such a cause. This is, at least at the level of individual decisions, nowhere *near* the purview of the "ought implies can" principle.

At the risk of apparent triviality, one way of expressing this idea that we are allowed to keep our money is to simply claim that we are allowed to keep our money. That is to say, we could defend the claim that we have a right to keep what we have earned, or some similar claim which entails that a morality which says we *must* donate it oversteps its bounds. This, of course, involves the issues of utilitarianism's inattention to such concepts as rights, which I said I would be avoiding where possible, but it does seem I need to say a few words about the subject at this point.

²⁴ This, of course, is essentially Unger's "Envelope" example.

Such a right could, with some initial plausibility, be grounded in the idea that we are *entitled* to what we have earned.²⁵ What we have done the work involved in acquiring, we should be able to enjoy the fruits of, or so the argument might go, at least provided we have not done so under any false pretences or otherwise mistreated anyone on the way to getting it. And though we consequentialists are not exactly known for our uncompromising defence of rights, even a consequentialist might, and in my view should, agree that there is *something* to this idea. We might, for example, point to the commonplace observation that in the latter days of the Soviet Union, people saw no incentive to work hard if they were not going to end up any better off than those who didn't. There seems to be a consensus that this significantly hurt the Soviet economy – an effect which no utilitarian can take lying down. Therefore, even a utilitarian should recognize some degree of entitlements. The disagreement on this point between a utilitarian and a rights theorist should be, not over whether such entitlements exist at all, but over how important such entitlements are. (And, of course, over the meta-issue of whether they are moral ends in their own right or a means to another end, but that is not so immediately relevant at this juncture).

²⁵ This is not dissimilar to Robert Nozick's view.

A common criticism of the entitlement view is that we *haven't* in fact earned much of our income, in at least one sense that it is plausible to regard as morally relevant. Anyone likely to be reading this essay was almost certainly born into relative privilege and has without a doubt taken advantage of many opportunities – educational ones at the very least, since no-one who doesn't know me personally will ever see this without access to a good academic library – that simply aren't available to starving people elsewhere on the globe. And we didn't do anything to earn those opportunities, so the degree to which we are entitled to their fruits is questionable. And this could be thought to undermine the justification for the entitlement view.

But this argument is not available to me *qua* consequentialist, at least not without a great deal more work, which I shall not be doing here. From a consequentialist point of view, whether we should recognize such entitlements has nothing to do with their origins, even setting aside the worry that the whole objection might be a genetic fallacy. We appeal instead to the *results* of recognizing them; their justification comes, not from the past, but from the future.

Consequentialists, then, should endorse such entitlements because and to the extent that they lead to a greater overall amount of good existing in the world. We are entitled to keep a slice of the pie, beyond merely meeting our strict needs, because

respecting such an entitlement makes the overall pie bigger (e.g. through its beneficial side-effects on the economy). But at the same time, if endorsing such entitlements costs lives, as can be the case in the sorts of scenarios I have been focusing on, we obviously cannot endorse just any amount of entitlement. A utilitarian must be prepared to say that large-scale economic benefits are worth *some* lives – that is an unavoidable consequence of the theory (and, at any rate, is independently plausible when it is remembered that it leaves us better-equipped to *save* more such lives). How many is a difficult question the present work makes no attempt to answer, but surely there is *some* limit. Therefore, the utilitarian is committed to recognizing some sort of entitlement but *not* to the view that we can turn down just any request for assistance on behalf of the needy abroad.

Which of the two clauses of this conclusion is operative in a given instance is a difficult question that seems to be as much about economics and psychology as it is about philosophy; what we can be certain of is that from a consequentialist standpoint, there *does*, morally speaking, need to be a “protected sphere” of some kind, but that it cannot possibly be so large as a Nozick would have us make it. This is important to know but, unfortunately for the present project, raises at least as many questions as it answers. While it gives us some idea of what sense can be made of the demandingness objection, and at the same time points us in a direction that

might be explored if we want to find a consequentialist solution to it, as stated it is annoyingly vague on both counts.

One might at this point think that, vague though it may be, the result of the above discussion seems in reasonably close harmony with most people's pre-philosophical intuitions. By common-sense lights, surely the idea of being entitled to do just anything we want with our resources is almost as implausible as the idea that we have no say at all in this. If the demandingness objection is interpreted as simply saying that we are entitled to spend our money in whatever frivolous way we like even with lives at stake, then the objection doesn't look so compelling.²⁶ Far better to see it as asking us to put limits on the extent to which morality may tell us to make such sacrifices than to depict it as calling for a complete ban on such sacrifices. There certainly are those who have held the latter position. But contemporary moral philosophers take few of them at all seriously.

But this observation brings out a point the non-consequentialist objector should, at this point, pounce upon. While the objector might be happy that the above line of argument moves the consequentialist *some* distance toward recognizing entitlements, it does so for the wrong reasons, or at least, fails to capture the objector's *primary* reasons. The objection is not supposed to be that, as a whole, there is more good

²⁶ This remark is taken, almost verbatim, from a comment by Joyce Jenkins on a previous draft of this essay.

produced if we can keep at least some of the goods we “earn”, welcome though that observation is. The objection is that there is some manner of *inherent* value in allowing us to do this – that a theory which claims we may not is *ipso facto* in some way morally obnoxious. There is, the objector will want to say, some value attached to being able to exercise our capacities, or decide the course of our own lives, or to express who we are as individuals, in ways of our own choosing. The above argument, while correct as far as it goes, simply *ignores* this claim.

It would be most implausible, intuitively, for a moral theory to deny that something of this general sort is true.²⁷ And indeed, most consequentialists *don't* deny this. I have noted that preference-satisfaction is at the center of one widely held theory of the good, and it seems like the ability to choose and act upon precisely these sorts of “projects”²⁸ is in turn at the center of preference-satisfaction. If consequentialism nevertheless demands that we should help others to the exclusion of actually getting to do this, this reflects the deplorable state, not of consequentialism, but of the world. The consequentialist wants a world in which we are all free to pursue our own

²⁷ Consider the three principles I mention at the beginning of this chapter, and describe in (a little) more detail at the end of Appendix A, for evaluating intuitions. While I believe that most of the conclusions such principles lead us to will be consequentialist-friendly, it *also* seems to me that they will end up endorsing some place for “autonomy”, as I shall be calling the cluster of claims defended here.

²⁸ I use the word “project” in the extremely broad sense now customary in moral philosophy. My projects might include, not only things with reasonably determinate endpoints like completing this thesis or cleaning my room, but things like being a more generous person or living a life of high intellectual achievement, which seem outside the scope of the word’s everyday sense.

projects as much as anyone – indeed, rather more, a consequentialist might pointedly add, than those deontologists who *don't* think we have strong obligations to the needy. But the present state of the world – in which we may include not only injustices across the globe but the lack of moral fortitude of most of those in a position to help – just does not afford those of us who would follow the consequentialist creed such luxuries.

But if consequentialists allow for valuing of personal projects – the *ability* to pursue them as well as their actual pursuit and fulfillment – there are reasons to be suspicious of the *ways* in which they do so, and that is the subject the remainder of this chapter sets out to explore. While I am not without sympathy for the consequentialist reasoning rehearsed above, if I fully believed it, it would be starkest hypocrisy for me to even be sitting here typing these words. Consequentialists recognize the value of such projects' being available and pursued, but not of their belonging to particular people. If the critics are wrong to suggest that consequentialism ignores what makes us human, they are nonetheless on more solid ground when they claim that they ignore much of what makes us *particular* humans, however silly I may find the “ignoring the separateness of persons” language in which they often do so. (The issue is not with the *separateness* of persons, but with the particular differences between them, and what the appropriate response to that is.)

Consequentialism, then, can be objected to on the grounds that it *alienates* us from our projects, goals and individuality as people.

The Excessive and the Alienating

Let us now consider in more detail exactly what, on this interpretation, the demandingness objection – or as we might now also call it, the alienation objection – must be claiming.

Let us call the cluster of values identified near the end of the previous section “autonomy”, for it seems to me most people who write about autonomy have something very much like such a space for personal projects in mind. The non-consequentialist objector, then, is claiming that making room for autonomy is an adequacy condition for an acceptable morality. The problem with the utilitarian dystopia sketched in the previous chapter is precisely that it has so little room for this; the people in it are valued but not respected,²⁹ their ability to contribute to the overall good seemingly regarded as their only morally salient feature. Something is missing here. The demandingness objection can then be recast: the fundamental failure of utilitarianism is that it fails to respect this adequacy condition for an acceptable morality, and our feeling that it is excessively demanding just manifests our intuitive sense of this fact.

²⁹ Perhaps revealingly, the phrase “valued but not respected” is one I first heard in reference to the prevailing attitude toward female “talent” within the pornography industry!

Having said all this, the non-consequentialist can continue to press the attack; not only is the consequentialist marking out too small a protected sphere, he does so for the wrong reasons. Even if we can drag the utilitarian (kicking and screaming all the way) to the point of admitting the permissibility of pursuing our own ends, the thinking behind this grudging admission is deeply wrongheaded. Allowing us to pursue our own projects is, for the utilitarian, a means of contributing to a larger end, that of human happiness generally, directly because of its contribution to our own happiness, and more importantly, indirectly because it allows us to contribute to a greater extent to *other* people's happiness (mainly by keeping us psychologically healthy enough to do so). This is perverse.

The point is especially clear for such long-term projects as being a good parent or spouse, and here we come to a fairly common objection to consequentialism, one which is normally taken to be a *separate* objection from any that concerns demandingness; and while the objections are indeed distinct from one another, I am here going to make a case that they overlap considerably. Suppose you ask someone how he finds the patience to be such a good husband, and he replies: "Well, I believe people should produce as much good as possible – basically, it's best if we help each other out wherever we can. As it happens, I know better than anyone how to meet her needs, and it's gotten to the point where it's not that much effort either, and

actually pretty rewarding. So, clearly treating her well is for the best." I, for one, would smile nervously and begin backing toward the nearest exit.³⁰

The utilitarian can, and often does, make a certain *partially* successful response to this charge. Every seriously-defended consequentialist view with which I am familiar holds that hedonism is at least part of the truth, in other words, that human happiness is at least part of what makes a set of consequences good and thus what PTR tells us to promote. And I concur with Parfit in taking it that no consequentialist view which *didn't* hold this could be plausible.³¹ But inextricably entwined with such happiness is the ability to choose, within limits, our own destinies that autonomy captures. The consequentialist, then, can simply point out that a world in which we uphold the respect for autonomy that common-sense morality dictates is, in general, a happier world than one in which we do not. If such respect is indispensable to happiness, then it is *ipso facto* indispensable to the following of the consequentialist moral creed. According to this view, the consequentialist may, with perfect consistency, agree that the utilitarian husband's answer to your question is perverse and still hold that the consequentialist view that underlies it is the correct moral theory. The husband's problem is not that he does what he does for consequentialist reasons; it is that he does what he does for *poorly thought-out* consequentialist

³⁰ The example is adapted from Peter Railton, "Alienation, Consequentialism and the Demands of Morality", *Philosophy and Public Affairs* 13 (1984), 135-6; and the general line of thinking I will momentarily be using to respond to it also owes much to that seminal paper.

³¹ *Reasons and Persons*, 4.

reasons. He is, on the view I am sketching – that of what Railton and others call a “sophisticated consequentialist”³² – following his basically sound moral theory in a superficial way, apparently unaware of its full implications; even if he is ultimately correct about what it tells him to do, he is getting the right answer for the wrong reason, or for only a poorly-understood *part* of the right reason.

This argument is persuasive in its original context, that of defending special treatment for those closest to us, though questions might yet be raised about just how much special treatment it justifies. But it can be counted only a miserable failure as a full consequentialist defence of the way most of us live our lives. I just cannot see how anything I have said gives utilitarian reasons to prefer small increments in the happiness of those nearest to us over the *lives* of strangers. Perhaps the sophisticated consequentialist need not think taking his son to a baseball game every so often is evil, but he probably shouldn't get season tickets; the extra money would be too useful in Africa. Doing well by one's children may be important, but any plausible consequentialist doctrine must nevertheless hold that saving lives by the dozen (or even more), as we in many cases have the economic power to do, is *more* important.

³² I dislike the phrase, as it implies that there is some other legitimate consequentialist view. I believe that so-called “sophisticated consequentialism” is just consequentialism *simpliciter* – in other words, it seems to me *any* consequentialist (who values happiness at all) is committed to this view. I mean this both as a statement about the logic of consequentialist theories, *and* as an empirical observation about the views of actual existing consequentialists.

In itself, this might not be so bad, for either the consequentialist or her opponent. As mentioned in the previous section, to some extent, it's probably true on *any* plausible moral theory that we should do more to feed the hungry, cure preventable diseases, and remove murderous, thieving dictators from power. And we've at least found a reason to put *some* limits on our obligations to do so, though many questions would still need to be answered to determine exactly what those limits are. But I suspect those limits are very low; and even if they're not, the suspicion may remain that these limitations are there for the wrong reasons. Even if the utilitarian husband from my earlier example becomes a "sophisticated utilitarian" and comes to treat his wife as he does because he loves her, what makes that the right thing to do is still its results, not the inherent value of their relationship and its place in their lives. Even if his conscious aims no longer reflect that fact, this still seems wrong.

The careful reader will have noticed, however, that the arguments I have just discussed are pretty specifically directed at the *utilitarian*. And it might appear to such a reader that we have reached a point where utilitarianism and consequentialism more generally come apart, or in other words, one where the correctness of my claims depends on what theory of the good we hold. For, if it is indeed plausible to treat autonomy as having value, then why can't the consequentialist do just that by adopting a pluralist theory which values both

autonomy and, say, happiness?³³ I am not sure whether there are any serious defenders of such a view,³⁴ but I actually find that fact surprising; it strikes me as an interesting and plausible position. But though it has a number of things to recommend it, a good intuitive fit with the idea that we are entitled to moral space in which to pursue our own projects is not one of them. For by claiming that autonomy is to be valued in a consequentialist (that is, teleological) fashion, such a pluralist theory gives me, not a moral space in which to pursue my own projects, but a reason to enable personal projects to be pursued, full stop – without regard for whether they are mine, yours or those of the Queen of England.

Now, it may be true that I am in a better position to actually advance some project of mine (designing a role-playing game, say) than any project of yours (say, improving your performance at some sport in which I take no interest) or the Queen's (getting some rest, no doubt). And we often have as projects, not merely that certain things be accomplished, but that *particular people* accomplish them – most commonly ourselves, as with Parfit's contortionists,³⁵ but sometimes our children or other loved ones as well, as with a father who strongly wants his son to make it onto some

³³ It is even possible, of course, to have autonomy as the *only* value in such a theory, but I take this possibility to be merely theoretical; I would be extremely surprised to see anyone seriously defend such a theory. (This is *not* what, say, the Nozick of *Anarchy, State and Utopia* is up to, as he does not want us to *maximize* autonomy, but to *respect* it.)

³⁴ Though Shaver has pointed out to me that Amartya Sen and Martha Nussbaum may hold something like it.

³⁵ *Reasons and Persons*, 90.

hockey team. In such cases, there may be little or nothing I can do to help some such project along, unless it is mine or that of someone very close to me. But even so, this does not seem to be a defence of the same ground the non-consequentialist is staking out. If I *am* in a position to either help my own projects or to help yours slightly more, surely the non-consequentialist impulse under discussion says I am permitted to help mine. Pointing out that such cases are unusual is beside the point; surely they do happen, and the consequentialist theory being discussed would seem to get the wrong result in such situations.

Now, we do need to be careful here – notice that autonomy is not the *actual* pursuit of projects, but roughly, the possibility of it. And it does seem a little more plausible for a theory to claim that I have reason to increase your autonomy than to actually pursue your goals. But the theory on the table still implies that I have *just as much* reason to enable your projects as mine, and this seems almost as mistaken, if we wish to *respect* rather than merely granting consequentialist value to autonomy, as giving the two of us equal reason to pursue those projects. Once again, we seem to be ignoring the fact that, whatever autonomy is carving out space for, it is by its very nature something unique to each of us.

In short, while this form of consequentialism might (for example) contain the seeds of an interesting defence of political liberalism, it just does not address the alienation

objection in any form that a non-consequentialist is likely to find plausible. The consequentialist should look for a different strategy, on pain of (at best) reaching an impasse. Indeed, the foregoing suggests that we *must* depart from consequentialism in order to address the point about autonomy. Consequentialism has the wrong structure, and is motivated by the wrong concerns, to address the worry on the table correctly. PTR must be supplemented by some other principle, one that is *not* straightforwardly consequentialist. In other words, we must settle for a theory that is what I earlier called “pseudo-consequentialist”.

It is only appropriate, then, to turn our attention to perhaps the best-known such theory at this point.

III – Scheffler-like Views

The Hybrid Theory

This chapter will look at various initially plausible-looking strategies that the “pseudo-consequentialist”, in the sense described in chapter 1, might use to accommodate the intuition about autonomy. I will begin with one of the most famous of the theories that have been proposed to this end, the so-called hybrid theory first suggested, to the best of my knowledge, by Samuel Scheffler³⁶ (though not actually endorsed by him). This is a pseudo-consequentialist theory, that is, one which is not strictly a consequentialist theory, even in my broad sense, but nonetheless seeks to preserve what is attractive in views such as utilitarianism in light of objections similar to those I have been discussing. On a hybrid view, we consider two independent³⁷ scales when deciding whether we have reason to take a particular action – the value of the action from an impersonal standpoint, or in other words, what would be required by the best consequentialist theory (whatever that turns out to be), and its value from the personal point of view, or in something closer to the terms I have been using, its contribution to the personal projects of the acting agent. (We can say that objectionable alienation exists if we are required to act in a way that does sufficiently poorly on this scale.) Supposing for the moment that each

³⁶ *The Rejection of Consequentialism* (Oxford: Clarendon), 1982, 20-22.

³⁷ Independent, that is, except for the contribution one’s own good makes to the impersonal scale. But we may assume this will be negligible in most instances, particularly those we will be most concerned with.

of these can be measured quantitatively,³⁸ we first multiply the latter value by some (presumably very large) scaling factor, such that the personal value of the personally best options times this scaling factor tends to be of similar magnitude to, and often exceeds, the impersonal value of the impersonally best option. We then rank various proposed actions by the sum of their impersonal value and their scaled-up personal value. Next, we see where the impersonally best option ranks on the combined scale; this action is always permissible. But *so is any action that ranks higher than this action on the combined scale.* (And no other constraints on acceptable actions exist, at least on the original and simplest hybrid theory, that of *The Rejection of Consequentialism*; this is one of the features that makes me comfortable calling the theory pseudo-consequentialist rather than outright non-consequentialist.)

Obviously, this allows us to give greatly disproportionate weight to our own projects when judging possible actions, yet still grants that acting as the utilitarian would demand is a perfectly acceptable option – for it would be uncomfortable at best to be committed to saying we *cannot* do the action which is, overall, for the best.³⁹ Despite this, we can nearly always promote our own projects at the expense of the overall good; only when the impersonally best action also helps our own projects along in

³⁸Even if they can't, discussing them for the moment as though they could will be helpful in making it clear what Scheffler's suggestion is actually endorsing.

³⁹For someone with consequentialist sympathies, that is. Of course, if we accept the existence of rights, this may not be so uncomfortable. But even if we do, surely mere self-interest is not the sort of reason we want to appeal to for barring the best overall action, when we do so.

some fairly significant respect, or when we have no options that significantly help those projects, is it likely the utilitarian-endorsed action will come out obligatory (i.e. will itself be the highest-ranking action on the combined scale), and these seem likely to be cases in which the demandingness objection is not a compelling one in the first place. We will never be *required* to make great sacrifices to help others, though we practically always have at least one option to do so if we so choose (assuming that the impersonally best option will require such sacrifices, which presumably would usually be true). When described in these very broad terms, this view seems to comport with most people's intuitions, though I shall soon raise doubts as to whether this is really true.

We might also hope that a hybrid view would disallow any actions that are, by ordinary moral standards, blatantly unacceptable, on the grounds that the (negative) contribution to the combined scale from the impersonal value of such actions would drag them down below the acceptable level. Unfortunately, this is almost certainly not the case, and this is a serious problem for hybrid views.⁴⁰ Suppose the best way to advance an important project of mine, say, finishing this thesis, involves murdering you; from the personal point of view, this is simply the best possible action for me at some juncture, even taking into consideration whatever distress I would feel at performing such an act, the risk of getting caught, which we may

⁴⁰ Kagan has raised largely the same points I am about to in "Does Consequentialism Demand Too Much?" in *Philosophy and Public Affairs* 13 (1984), at 251.

assume I'm clever enough to minimize, and so on. (The reader is invited to amuse him- or herself by filling in the details.) Now, from the impersonal viewpoint, obviously your death is a strike against this action, but is it a large enough strike against it to render it impermissible? The answer appears to be no. For consider a fairly routine act – I buy an overpriced coffee drink at a local coffee shop, where I like to go to read philosophy and work on the present essay. (The latter activity is one I find easier in this environment, and one which advances various projects of mine, obviously including this thesis.) This is just the sort of thing whose defensibility is called into question by the Demandingness Objection and presumably defended by the non-consequentialist, which makes it precisely the kind of case we're interested in. But even supposing these things are a worthwhile use of my time in the first place, I could do them at home (at least, if I were a little less easily distracted than is in fact the case), so going out is not necessary; so let us consider an alternate use of the same money such as stuffing it in a UNICEF box at the convenience store across the street. The latter might very well save multiple lives; let's say two (and if, as a matter of empirical fact, this is not plausible, I need not change the example all that much to *make* it plausible; anyone who is reading this in the first place can no doubt fill in the details him- or herself). If we maintain that I am nevertheless allowed to go out for coffee when I do my reading, then it must be okay for me to do such a thing at the expense of two lives. But now look at the case involving murdering you again. We are talking about a presumably much larger

advance of my own interests at the expense of just *one* life, namely yours.⁴¹ All else being equal, it seems that murdering you is better than going for coffee from both the personal and the impersonal point of view, and so must also be acceptable under our hybrid theory.⁴²

Biting the bullet – accepting that both actions are in fact permissible – is obviously a very unpalatable option for the consequentialist at this juncture. Recall that our whole motivation for looking at a hybrid view was to move consequentialism (or something close to it) in the direction of matching our everyday intuitions. That murder is acceptable if we gain enough by it is precisely the sort of thing we therefore should avoid saying if at all possible. In fact, it is worth pointing out that not only is this a claim the non-consequentialist will clearly find unacceptable, but also one the (typical) consequentialist isn't committed to either, if the advancement of one of my personal projects is all that is at stake. The hybrid theory, taken in this

⁴¹ This is too quick, since as I've presented things in the main text, the people who die in the coffee shop case die in the murder case too, in addition to you. To fix this flaw in the example, add the following detail. Since I'm being maximally prudent, after murdering you I might as well take the opportunity to rob you; this gives me some extra resources and also directs attention away from the real motive for the killing, and thus, from me. Using part of the extra money I acquire in this fashion, I make the same charitable donation I presented as an alternative to the coffee shop visit, because this slightly eases my conscience. This keeps the balance in lives correct.

⁴² The option of donating the money while refraining from *both* the coffee shop visit and the murder is of course superior to either of the other options presented in impersonal terms. But (presumably) it is inferior to either one in personal terms. It is thus not clear where it will end up falling on the combined scale, but even if we (reasonably enough) assume it ends up ranking higher than either of the alternatives, I don't think this undermines anything I have said.

direction, therefore has an unattractive consequence that is not shared by *either* of the theories between which we may have hoped it would be an acceptable compromise. This is not progress.

Going in the other direction – claiming that *neither* action is acceptable, that we must make the UNICEF donation or at any rate do *something* with greater impersonal value than the visit to the coffee shop – is far more plausible; even many non-consequentialists will concede that we ought to give more to the likes of UNICEF, and some might be prepared to entertain the possibility that this is obligatory. At the very least, nearly everyone thinks it would at a minimum be *nice* of us to do so. But despite the possibility of these steps toward a compromise, I don't see how this move can be counted as successful at *addressing* the Demandingness Objection; it is more a case of the consequentialist digging in his heels on this issue, or so it appears at this juncture. To put it another way, while not so unpalatable to the dedicated consequentialist, this sort of move will be fully acceptable to very few others (even if they do admit there is *something* to be said for it), particularly given the context.

We might try to claim that, while the hybrid theory fails to address demandingness concerns in this particular instance, it nevertheless does so in *enough* cases to show that morality is not extremely demanding, or at least constitute progress toward that goal. But not only would a case need to be made for this claim, but the prospects for

making such a case seem very dim when one considers the ease with which we can save lives using the sorts of resources used by many of our everyday activities; it is just too easy to find cases like my coffee-shop example where the consequentialist will say that multiple lives hang in the balance.

The other obvious way in which we might try to salvage a hybrid view is to introduce some sort of constraint that would eliminate murder (and other paradigmatically immoral acts) from contention while not doing the same to the sorts of “mere” letting die that seem to be acceptable by everyday standards (or one which relies on the distinction between intending and foreseeing harm which is often resorted to in the face of failures to justify the doing/allowing distinction). Such a move would find moral significance in, perhaps, the lack of a direct causal relationship between my actions and the deaths in question.

It might very well be true that this is the best way to address the problems raised by the example involving murdering the reader. There is a vast literature on these sorts of moves and to do it justice would at a minimum double the length of this essay, while being of only tangential relevance to its main thrust. It must, therefore, suffice to say that few consequentialists find such constraints palatable, and that there is no consensus even among those moral philosophers who accept them on how they could possibly be justified. Even some who reject consequentialism would not

endorse such a move; and given that it is his hybrid theory we are discussing, it is particularly notable that Scheffler himself is sceptical of the possibility of justifying such constraints, despite their intuitive appeal.⁴³

To restate this point from a slightly different perspective, constraints mark such a significant departure from the idea behind consequentialism – that the moral status of some action depends only on the goodness of the results we can expect to arise from it – that it is doubtful that many consequentialists would accept such constructions into their theory while any hope remained of creating a defensible theory any other way; so as long as we are attempting to reply on behalf of the *consequentialist* or close ally thereof, and not some other sort of moral theorist, we should avoid rights and other such constraints.

As a final point against the hybrid theory, it is, with or without constraints, what I shall, in the next chapter, refer to as a “bright line” view; and all the points I will be making there against such views apply to the hybrid theory every bit as much as they do to traditional (and satisficing) utilitarian views. Therefore, even if there is an account of constraints that the pseudo-consequentialist can (no doubt reluctantly) accept, there is independent reason to reject the hybrid theory.

⁴³ See, in particular, chapter 4 of *The Rejection of Consequentialism*. But Scheffler does seem to back down from this position somewhat in his “Prerogatives Without Restrictions” (reprinted as the third appendix of the Revised Edition of the same book).

Beyond the Hybrid Theory

It seems that the hybrid theory is not likely to give us a way to address the Demandingness Objection within the spirit, much less the letter, of consequentialism. Yet perhaps something can be salvaged from it. The notion at the heart of the hybrid suggestion – that there are at least two *independent* sources of reasons, one personal and expressible in terms of individual projects, and one impersonal and expressible in consequentialist terms, each compelling despite often pulling us in incompatible directions – has great appeal, both intuitively and theoretically.⁴⁴ Suppose that this much is true. There is no reason to think Scheffler's hybrid suggestion is the only possible view on how the two sorts of reasons interact. So let us explore some other possibilities.

It would seem clear that any view which appeals to the *sum* of the personal and impersonal values of some state of affairs (however the former is scaled) will carry at least some of the unattractive implications discussed toward the end of the previous section. (This is setting aside my general scepticism over whether such pseudo-mathematical talk is anything more than a useful explanatory shortcut in the first place.) So we should avoid directly appealing to the sum of the two. At any rate, we might think that such an appeal ignores the *independence* of these two sources of reasons. Perhaps, then, we could claim instead that each scale has some threshold

⁴⁴ This idea is by no means original to Scheffler, much less myself. A particularly influential recent expression of a very similar idea is found in Nagel's *The View from Nowhere*, with the titular view corresponding to the impersonal perspective.

below which the best theory will advise us not to go; any act whose impersonal value is less than X is Immoral, while any act whose personal value is less than Y is A Really Bad Idea, for some X and Y.

Before discussing this suggestion in any detail, a somewhat lengthy terminological digression will be helpful, for my use of the word “immoral” just now may be striking and deserves a little explanation – explanation that, some readers might take it, raises further issues that deserve comment. I must say I am far from convinced that the personal scale has anything to do with *morality* as such.⁴⁵ This in no way makes it not a source of reasons, or a lesser one, or is in any other way meant to denigrate it. But very few writers hold that self-interest, even in the fairly enlightened sense which we may assume the personal scale captures, is a *moral* value, and those who do would for the most part not recognize the *impersonal* scale as having any importance. Until the last paragraph, I had been writing as though both scales were a part of moral theory, when this is at least arguably not the case. I tend to think, instead, that the impersonal scale alone corresponds to morality, while the personal scale corresponds to *practical rationality* instead. And so from here on, I shall occasionally use these terms in ways that reflect these views. Anyone who objects to this when read as a substantive claim would be well-served by simply regarding it as a semantic convention I have chosen to adopt. No point essential to

⁴⁵ This also seems to me to be a further point against the pluralist consequentialism that valued autonomy which I briefly discussed earlier.

the present essay hinges on which of these interpretations of my terminology one prefers.

By doing this, I do not mean to suggest it is a settled matter that morality is in fact captured by something like my impersonal scale; readers bothered by this can regard it as an assumption for the sake of argument. (This clarification gives us, *en passant*, another way of stating the demandingness objection – we can now say that morality, at least in the form which I will be discussing here, conflicts with practical rationality in a way that seems objectionable).

At this juncture, both the consequentialist and the demandingness objector are likely to want to respond. The consequentialist might believe that I have just conceded that he is right as far as morality is concerned; there may be a further issue about how morality and rationality interact, but the consequentialist *qua* consequentialist has no horse in *that* race. If it turns out that consequentialism does not always answer the question of what we ought to do, it nevertheless does answer the question of what we ought *morally* to do. Thus the consequentialist may feel entitled to simply declare victory and ignore the rest of this essay. For her own part, the non-consequentialist who makes the demandingness objection may be indignant at the suggestion that, all along, her point was never about what morality can ask of us, but about how morality interacts with something else entirely. I seem, from this standpoint, to be

claiming that there really is no limit to moral demands; rather, morality is constrained, if at all, by some *other* source of reasons.

Nevertheless, I am going to stick to my guns on this interpretation of the debate for the remainder of this essay. This issue is one that another work of this length could easily be devoted to, and so I cannot here say everything there is to say on the subject, but I wish to make two remarks about it, one against each of the above two hypothetical objectors. Firstly, it seems to me that if this is what the demandingness objector wants to claim, she may be confused. Surely she is not claiming that it is morally better *not* to contribute optimally. It just appears obvious that this is, in fact, the best we can do morally; if it isn't, then it is not actually contributing optimally after all. The position that it is morally better not to *require* contributing optimally, however, looks very strange side-by-side with this claim, even if they do not actually contradict one another. Certainly attempts to justify the claim, in this form, have not, to date, been terribly persuasive. This internal tension in the position of one who makes the demandingness objection can be avoided by saying, not that there is a limit to what can be demanded of us morally which is somehow required by morality itself, but that while morality may ask a great deal of us, there are other considerations which make some of the things it asks, not demands, but something less forceful – polite requests, as it were. At least, I find this position more plausible, not to mention easier to understand and justify, than any which states the

demandingness objection in purely moral terms, including those early in the present essay.

The other point I wish to make, which bears on both the objections I have mentioned but especially on the consequentialist one, is that ultimately, when we do moral theory and cognate areas of philosophy, surely we are interested in what we ought to do, full stop. Thus the consequentialist declaration of triumph, given above, seems beside the point. If the consequentialist can only account for one of (at least) two equally important aspects of what we ought to do, surely this is an important respect in which his theory is incomplete. Or, if the consequentialist continues to resist this conclusion, he should at least admit that any declaration of victory at this point would be premature. If what we morally ought to do is only part of what determines what we (just plain) ought to do, then the consequentialist can be right about the former without having achieved any sort of *interesting* victory with regard to the latter issue – and that, surely, is the one we are interested in. The victory is, in every sense of the word, solely academic.

To return to the main discussion, let us say a few words about what can be said for the two-threshold view I just sketched. From the standpoint of our usual moral intuitions, a theory with the “dual threshold” structure I have suggested has some clear advantages over, at least, utilitarianism and other structurally similar

consequentialist theories. It allows for supererogation (A Really Bad Idea that ranks very high on the moral scale), tragedy (a situation where no available option meets either threshold), a variety of acceptable actions rather than just one (assuming there are often many actions that meet both thresholds), and, hand in hand with the latter, degrees of moral admirability, without claiming that those who do more poorly on the moral scale are necessarily *immoral* for doing so. All of these concepts are a part of what I take to be our everyday moral views yet seem to be ones utilitarianism has no room for.

At this point, I owe the reader a few clarifications about what this view says about permissibility. Clearly any act which meets both thresholds is acceptable and any which meets neither is not. The reader will note a pleasing symmetry here, and may be inclined to view this as a point in the theory's favour. But this appearance breaks down when we consider acts which meet one threshold but not the other. I have already suggested one way in which this is true; (at least some) acts which are morally permissible but count as Really Bad Ideas are supererogatory, but we certainly don't want to say something analogous about acts which are self-interestedly wise but immoral (such as murdering you, in our running example!). Presumably supererogation is praiseworthy, but actions which have the reverse pattern tend to include the ones we most want to condemn (at least an act which meets *neither* threshold is to a certain degree self-punishing). So the theory should

say that we should try to stay above both thresholds at all times, but if we cannot, the rational threshold is the one to sacrifice, not the moral one – and so the symmetry the view may have appeared to have breaks down. We may find it *understandable* to fall short of the moral threshold for the sake of meeting the personal/rational one, if we do not fall *too* far short; but this is to say that such an act is only a minor violation of the normative law, not that it isn't a violation at all.

Not all utilitarians have taken the absence of the intuitively pleasing elements I have mentioned – multiple acceptable acts, supererogation, etc – from their theory lying down. There have been a few attempts at a *satisficing* version of utilitarianism, which would, in its simplest form, be exactly like the moral “half” of our dual threshold view – an action is morally permissible if and only if the goodness of the resulting state of affairs meets some minimum standard. Satisficing views have not been widely adopted because it turns out that they encounter quite a number of serious problems. Since our dual threshold suggestion has much in common with satisficing views, it makes sense to discuss these problems here and see if the dual threshold view falls to the same difficulties.

Perhaps the most serious such problem is the one illustrated by a number of famous examples, of which my favourite is

Achilles’ Computer: Achilles is sitting at a computer which is networked to an extremely efficient, fully funded food- and medicine-distribution system.

He can type in any number he likes, and that many of the world's neediest people will have sustainable access to sufficient food and medicine to live comfortable lives from then on, with only about as much effort as it takes the typical Westerner to achieve these things. Once Achilles sets the system in motion, however, no-one can override the number he chooses for at least ten years. He knows what the largest number he could usefully enter is (and the optimal number, if for some reason it differs from this). For no particular reason, Achilles enters a number that is about half as large as the optimal number.

Intuitively, Achilles has done far less than he should; he has no reason not to enter the optimal number, yet he does not do this. But on any reasonable satisficing view, Achilles must have acted permissibly; after all, the world that results from his actions is *far* better than any that has ever resulted from any of yours or mine, and presumably we have acted in an intuitively morally permissible way on *some* occasions. So if we have ever met the threshold, then clearly Achilles has too. Thus the satisficing view would seem to entail that Achilles acts permissibly, which seems ridiculous to most of those who see this example; at the very least this seems to show a failure to reconcile our intuitive judgments with consequentialism, which was presumably the whole point of considering the satisficing view in the first place.

Does the dual threshold view fare better? Not without modifications, for there is little likelihood that entering a relatively low number in the computer will count as a Really Bad Idea (i.e. fall short of the threshold for practical rationality). Entering a large number costs Achilles nothing that entering the smaller number doesn't, so it

seems that neither action is *personally* better or worse than the other.⁴⁶ So neither of the forms of criticism of which this theory admits seem to be available; entering the smaller number is neither Immoral nor A Really Bad Idea. But this does seem to be an action we want to be able to criticize. So we must once again take a step back and try to preserve what is attractive about the dual threshold view while avoiding this problem.

Our problem with what Achilles does seems to be, not just that he does less good than he could (for if we are interested in satisficing theories in the first place, including the dual threshold view, we obviously think that is at least sometimes okay), and certainly not that he does less well for himself than he could, but that it would *cost* him nothing to do better. This leads to another principle for integrating the two scales into a single theory – whenever we can do better impersonally without doing worse personally, we must, morally speaking, do so; and perhaps the obverse is true (whenever we can do better personally without doing worse impersonally, we must, prudentially speaking, do so). Or, to expand this in a plausible direction, if action A is *strictly* superior to Action B, in the sense that it is better on at least one of the two scales and worse on neither, we may not choose to do B if A is available.

⁴⁶ Granted, a very large number of people have good reason to be upset with Achilles. But we can suppose the identity of the person who enters the number is a secret, and in any event, most of the people with such reasons – namely, those he failed to help – are not in a position to do very much by way of exacting retribution.

If this principle is the *only* one that governs the areas of life – morality and practical rationality – which we are discussing, then we again seem to rank the cases we have been discussing in this chapter in the intuitively wrong order, a general problem with any theory that does not recognize some form of constraints. (The murder will come out strictly superior to the coffee shop visit.) But suppose we can overcome this difficulty. This new theory is silent on, and therefore could be taken to condone, cases in which we give up something of great moral value for only minor personal gain. This does, at least, make the coffee shop visit acceptable when compared to the donation, but we might worry about whether this is a plausible or attractive view in its own right. Say the result of murdering you, salutary though it is on the personal scale (relative to me), is only slightly better than the result of paying you a huge and sincere compliment, an act which we can safely assume is much better from the impersonal point of view (hold all other relevant differences between these two acts, save the small change in my well-being, constant). This theory is silent on which of these two acts is preferable, but our admissible intuitions are not.

The obvious modification seems to be to add a principle that if act A is *spectacularly* better on the impersonal scale than act B, act B's being only a little better on the personal scale does not justify a preference for B. If both scales admitted of precise measurement, we could precisify this principle too – if the impersonal difference

between A and B is greater than (say) 10x in A's favour, then no (scaled) personal difference of less than x can justify performing B instead of A. In any event, this formulation gives the general idea of what the suggestion is. (This is a stronger principle than the above "strict-superiority" one in the sense that it gives a definite answer in a wider variety of cases. But note that this one is asymmetrical – it does not let us trade off personal for impersonal value at anything like the same rate as we may trade in the other direction.)

Setting aside once again the problem of precise measurement, I wonder if this proposal does not bring us right back to making demands that are, commonsensically, unreasonable. Take a fairly mundane example. The overpowering Houston heat has followed me to my favourite Winnipeg coffee shop, but the even more overpowering Houston air conditioning has not. So I duck across the street to the convenience store for an ice-cream treat. They have none, but they do, as has come up a number of times already, have a UNICEF box. Being in a reflective mood, I ask myself what I should do. It is at least plausible that the above principle says that, simplifying down to the options of donating \$5 or donating nothing (just to make things clearer), I should donate the \$5; I won't miss it much and the expected impersonal value of my doing this is considerable. But repeat this for my next \$5. And my next, and my next, and so on up to (but, mercifully, not including) the point of actually impoverishing myself. There's a bank machine there, should I run short of

cash on hand. I am sure the reader can see where this is going – namely, right back where we started.⁴⁷

⁴⁷ The obvious move is to redescribe this as a single action. However, I see no well-motivated way of doing so. Exactly which such “composite actions” are legitimate and which are *ad hoc*? I know of no answer that I find persuasive.

IV – Scalarity

If we suppose that we must always do the morally optimal action, we have an extremely demanding theory – by many lights, unreasonably so. Such a theory demands that we give up much of what makes human life enjoyable and allows us to express who we are as people. And most of us strongly believe that this is not an acceptable result, even if it is admitted that we can and should do far better than we do by faraway victims of famine and disease. These are the intuitions I am attempting to capture by assimilating demandingness to alienation.

We have sources of reasons other than morality, the non-consequentialist objector should go on to say, and it is sometimes okay to let them take precedence. In the previous two chapters we identified one very important such source, which I called practical rationality. (I believe that it is the most important such source, and that most others that might be proposed are just special cases of it, though I have argued for neither of these claims.) This is the ability to decide our own goals and projects and how to go about achieving them. Chapter 3 was spent assessing various possible views on how these two largely independent sources of reasons interact, but none of the pseudo-mathematical schemes discussed there seemed satisfactory. Yet the intuition remains that these views are onto something – that it is legitimate to make tradeoffs between morality and practical rationality, that making the latter completely subservient to the former is supererogatory, not obligatory.

In this chapter, I will explore what I see as a far better consequentialist response to the demandingness objection, namely the Scalar View defended by Alastair Norcross. Yet, this view may seem unsatisfying to many readers, and so much of this chapter will be taken up by an attempt to address the reason for this (namely, an apparent lack of actual action-guidance). While this chapter fails to find a fully satisfactory answer to this problem, it will make several crucial steps toward this goal; near the end of the chapter, I will give a reason to think such a solution is *possible*, and on the way to doing so, I will present a theory I think has considerable promise as the basis for one. I will close with some thoughts on why, even if this project turns out to be a failure, this would not necessarily be such a bad thing; in other words, why it might be okay for a moral theory to fail in providing unambiguous action-guidance.

Against “Bright Line” Views

In this chapter I wish to discuss a different way of reconciling the two normative scales we have identified. Though the foremost current defender of the general type of view I shall be discussing is Alastair Norcross, let us start with a famous remark by John Stuart Mill:

The creed which accepts as the foundation of morals, Utility, or the Greatest Happiness Principle, holds that actions are right in proportion as they tend to promote happiness, wrong as they tend to produce the reverse of happiness.⁴⁸

⁴⁸ John Stuart Mill, *Utilitarianism*, Ch II, ¶ 2

It is interesting that this passage says, not that the only right action is that which most promotes happiness, but merely that actions are right *in proportion* to their tendency to do so. Mine is not a work of historical scholarship, and so my focus will not be on reconstructing Mill's views, but this idea of rightness – which, as mentioned in chapter 1, I take to be interdefinable with moral obligation – intersecting with some sort of *proportionality* is an intriguing one. It should be clear that, whatever it is, this is *not* the expression of an optimizing view, but of something quite different. Let us consider what that idea might be, with an emphasis, not on the historical facts of Mill's actual view, but on what role such an idea might play in the present inquiry.

We have considered at least one view which it seems could be described in terms similar to those Mill uses here, and that is Scheffler's. That view posited a combined scale for (in the terminology I introduced shortly after that discussion) morality and practical rationality, using a combined scale on which any given act was rated in proportion to how much good it did, measured impersonally (along with a similar scale for practical rationality). An act's overall "rank" on the combined scale was proportional to the amount of good that resulted from it; all else being equal, the more good that resulted, the higher the act would rank.

This clearly can't be *exactly* the picture Mill had in mind. Most obviously, Mill was trying to account for morality only, not normativity more generally, as our present project seems to have become (which, I have suggested, is a significant oversight). But there is another important difference. The statement of Mill's which I have quoted says that actions are right *in proportion* as they tend to promote happiness, but Scheffler suggests a "bright line" above which an action is acceptable and below which it is not. And clearly being above the line is what we mean when we say something is right (if we accept the hybrid view). On the hybrid view, no other distinctions seem to be important; certainly none is *as* important. What I believe the Mill passage suggests, by contrast, is that rightness (moral obligation) itself comes in degrees, rather than being an all-or-nothing matter. As I have hinted at several times, this seems to comport with our intuitive judgments; we tend to think of actions as being better or worse, not as merely being right or wrong.

The other consequentialist or pseudo-consequentialist theories I have considered have all tended to suggest the existence of such a "bright line"; the picture seems to be that, while the good (the measurement of what, quite generally, has value morally) certainly comes in degrees, the right (that which actually guides actions, insofar as anything about morality does so) is an all-or-nothing matter. For example, the satisficing view, though an attempt to accommodate the very intuitions I spoke of above, is in fact every bit as much a "bright line" view as the optimizing version

of consequentialism; it merely sets this line somewhat lower than the optimizing view does. It may be, then, that the reason satisficing fails to satisfy the intuitions it tries to respond to is that it has entirely the wrong structure; there is no way to duplicate, or even approximate, the intuition that right and wrong come in degrees by positing a sharp divide between the two, no matter how clever we might be about where we place that line. (Moreover, there is no obvious reason why the right and the good should differ structurally in this way; if anything, when put that way this dichotomy seems like an oddity that needs explanation of its own.) In short, it seems that to really reflect this intuition, we would have to admit that right and wrong are questions of degree – and do away with the whole idea of having a sharp distinction between the two.

Now, thus far it sounds as though what I am suggesting is to make the right, not be a simple function of the good (as most consequentialists would have it), but simply one and the same as the good. Taking Mill's words literally, then, would mean the degree to which we were morally obligated to perform a given action was directly proportional to the amount of good it produced. But clearly, this can't quite be right. This view still entails that the morally optimal action is the one we have the greatest moral obligation to perform, and that seems to just reduce to a maximizing view, which is precisely what one might assume we were trying to avoid. The view I am trying to capture doesn't so much assimilate the right to the good as omit the right –

and with it the notion of obligation – from our theory entirely. In other words, morality tells us that action A is better than action B, and that’s all. It does not, ever tell us that we must do A or that we cannot do B. Hopefully the information it gives us is relevant to our decision-making, but our moral theory is not to be construed as insisting on this point.⁴⁹ Utilitarianism is, on this view, a theory of the relative value of different actions, and that is all.

To motivate and make plausible such a rejection of the notion of moral obligations, consider what the utilitarian should say about various choices she might find herself making.⁵⁰ Suppose you are trying to choose between donating \$100 and \$200 to UNICEF, while I am choosing between \$200 and \$1000. A utilitarian is in a position to try to convince one or the other of us, but not both, to make the higher donation. Which of us should she talk to? Assuming the negative effect on ourselves is negligible and/or sufficiently similar between the two cases, that we are about equally susceptible to persuasion, and so on (in short, that the dollar amounts are the only morally relevant differences), clearly this utilitarian’s effort would be better spent on me; overall utility will, if she succeeds, improve by about eight times as much as if she successfully persuades you. (And this does not, by the way, seem to change if your choices are instead between \$1000 and \$1100.) Now suppose that our

⁴⁹ Norcross has developed this view in numerous articles, of which perhaps the clearest and most succinct is “The Scalar Approach to Utilitarianism” in Henry R. West (ed.), *The Blackwell Guide to Mill’s Utilitarianism* (Oxford: Blackwell), 2006.

⁵⁰ Here and in the next section, I use modified versions of some of Norcross’ illustrative examples.

moral theory turns out to put the line between right and wrong action, in both our cases at that particular moment, at a donation of \$150 (or \$1050, in the case where you're the more generous donor); everything else I have said about the two cases remains true. Does this change the answer, so that persuading you is suddenly more important? Clearly not. So why are we so concerned to draw such a line in the first place?

One thing that can be said about this **scalar view**, whatever its historical pedigree or philosophical merit, is that it is totally invulnerable to the demandingness objection, for it makes no demands at all. It merely gives us advice. Whether we take it is our own affair.

Scalarity and the Right

There is a persuasive-sounding objection that may be made at this point, which I will spend this section and the next considering (Appendix B discusses, and rejects, an alternative response to it). I certainly feel the force of this objection, but do not think it is ultimately tenable. Nevertheless, for the moment my Devil's advocate hat will be on, even more so than it has been throughout most of this work; I will try to make the most forceful case for it I can.

According to this objection, the above suggestion of going *solely* with a scalar approach to value goes too far. It's true, we might want to say, that we think some

actions are better than others without necessarily condemning either (or condoning either). But it does not follow that there is no use for a sharp distinction between the right and the wrong, or even that such use as the distinction might have is as secondary as it must apparently be on this view. After all, some actions clearly *are* one or the other, and a moral philosophy that denies this seems to be missing what seems like an indispensable element of such a theory. Even if we admit that rightness "comes in degrees" in the sense that the above example seems to illustrate, surely that can't be the most important aspect of our attempts to understand it; it seems to be an adequacy condition of such a theory to identify which acts are *right*, full stop.

The primary problem with incorporating both degrees of right and wrong *and* such a bright line is one of motivating any particular choice of where to draw the latter line (as opposed to the mere idea of having such a line in the first place).⁵¹ Once we admit that it is possible, even in principle, to admit one action is preferable to another without calling either one wrong, it becomes very difficult to justify any particular choice to call one action right and another wrong, among any two that are very close together in overall "rightness". Yet, if a line between the two exists at all, then there must be cases where this is exactly what we are committed to claiming.

⁵¹ This problem also exists for theories which *only* have the bright line, such as those I examined in chapter 3, and in fact my arguments against them could be seen as illustrations of this, but it is easier to see in theories which try to have it both ways.

Consider the different donations I discussed at the end of the previous section. Your options are donating \$100 or \$200 while mine are donating \$200 or \$1000; and we draw the line between right and wrong at \$150. As I said earlier, it nevertheless seems clear that a utilitarian should make more effort to convince me to make the larger donation than to convince you, and already that seems to be something of a case for the moral uninterestingness of the distinction. But now consider a facet of this case that I completely ignored before. What could possibly *make* it the case that the line was drawn at \$150? *Why* would a donation of \$140 be wrong but \$160 right in *any* sense of (even limited) philosophical interest?

To put the point more abstractly and yet, perhaps, more clearly, *wherever* we draw the line between right and wrong, there must be some trio of possible actions with the following features.⁵² Action A is just below the proposed line (say, scoring -1 on some scale), B is just above it (say, scoring +1), and C is *far* above it (say, scoring +1000). Yet, we are committed to calling A wrong and B right, in turn implying that, if the difference between right and wrong is to mean anything philosophically interesting, the difference between A and B is more important than the difference between B and C. And this raises at least two questions – what could possibly make it true in the first place, and what the moral relevance or philosophical interest of the

⁵² Technically, this will only be true in possible worlds where there are, at any given junction, a sufficiently wide range of possible actions. The real world, of course, is typically such a world. I assume ones that aren't are of no, or at best minimal and strictly theoretical, interest for moral theory.

distinction is supposed to be, given that as consequentialists it seems we should in fact care more about the difference between B and C in the only sense of “care more” that makes a *practical* difference.

There are at least two candidates for such a line that *do* have some initial plausibility, though even so, it is not clear to me how they could answer either of the above questions. We could draw the line between the optimal act and the next-best one, or at the point where our actions are exactly neutral with respect to producing good (the latter is actually presupposed by the numbers in the above paragraph). If we are sympathetic at all to the demandingness objection, we will reject the first of these suggestions, which leaves the neutral point as our best candidate “bright line”. (To be completely clear, I mean the point at which the world before we act and the world after we act each have exactly as much moral value – good – as the other.) But it is clear at once that this fails to capture the intuitions it is presumably intended as a concession to. There are many cases – consider Achilles' Computer again – where we think even the minimum we ought to do is considerably above the neutral point; most of us think Achilles acts wrongly in a broad range of cases in which his act is much further above this point than any act we have ever performed, and hopefully we think we have acted rightly on some occasions. And we could equally be faced with a tragic case in which, no matter what we do, the situation will be below the line; do we then act *wrongly* (even – especially! – once it is admitted that it's the best

we could have done) if we act so as to reduce suffering as much as possible? Such a claim is too implausible to seriously consider; it simply makes a hash of our normal use of the concept “right”.⁵³ So despite having a little more initial plausibility, this suggestion really seems no better than any other place we might choose to draw such a line.

One who wanted to press the argument for some specific line between right and wrong would, at this point, do best to acknowledge that we react differently to (real or hypothetical) moral agents depending on the alternatives available to them, and attempt to work that into his next reply. We might first consider saying that the line should be located at some fixed *relative* point on the scale, so that, for example, any act that is in the top 20% (say) of the available options is acceptable. Let us ignore, or more precisely, assume solvable, the obvious (and suddenly more acute than before) problems of enumerating just what our options are, and of measuring their value with precision; my opinion, in all honesty, is that these issues alone present at least two insurmountable difficulties, but I am not (in the present work) going to argue for this. Instead, I will merely point out that this too fails at tracking the intuitions that it is supposed to be trying to capture. While the whole premise of this essay has

⁵³ Of related interest, Norcross’ “Good and Bad Actions” (*The Philosophical Review* 106 (1997) 1-34) questions the usefulness and even the intelligibility of the concepts of “good” and “bad” (as distinct from “better” and “worse”) in consequentialist thinking; and his criticisms can be applied with little modification to “right” and “wrong” in the present context. Some of my reasoning in this section has drawn inspiration from this paper and discussions with Norcross related to it.

been that the optimizing view is not reasonable *in general*, there are cases in which we do think we ought to do as it suggests; Achilles' Computer seems to be one of them. So sometimes, this proposal is too permissive. But it would be wrong to conclude that it can be fixed simply by making the percentage of acceptable options smaller, for there also seem to be cases where it isn't permissive enough. At least if most of our usual intuitions are somewhere close to right, the scope of my *unacceptable* options for how to spend an upcoming Saturday night for which I don't currently have plans seems to be mostly confined to things like robberies and shooting sprees, and I trust we may assume these make up a much smaller percentage of my options than that which any plausible theory of the type I am considering would specify.

The Restricted Delta View

The idea behind our response to Achilles' Computer seems to be that Achilles could do better on the impersonal scale *at no* (or in any event negligible) *personal cost*.⁵⁴ Perhaps we shouldn't have to make substantial sacrifices to produce moral value in the world, but surely it is not okay to fail to produce such value when we can do so effortlessly. So perhaps a plausible bright line could be drawn at the best act we could manage without much personal cost. This is effectively a satisficing view of a very specific sort we have not considered so far – where the satisficing (or allied)

⁵⁴ Linking a limit on demandingness to cost is by no means a novel suggestion; for example, much of Kagan's book is devoted to replying (on behalf of the traditional consequentialist, or other "extremist" in his terms) to many different versions of this very idea.

views we have considered so far have said we must reach some minimal point on the moral scale or the combined scale, this one posits a point on the *personal* scale below which we cannot be obligated to go. (It does not prohibit doing so; it just refuses to demand it. Greater good, accomplished by greater sacrifice, is supererogatory.) More subtly, it does not use our absolute point on that scale, but its delta over a given action, to determine where this point is. It says that we must do at least as well, on the moral scale, as the best act we can without lowering our position on the personal scale by more than a certain amount. It is wrong to do worse than this, since we can so easily do better. (And, while maybe the line will turn out to be such that I should make one \$5 donation in UNICEF box cases like the one at the end of the previous chapter, we can – and should, if intuitive plausibility is our goal – set the acceptable delta low enough that it won't take many such donations to reach it.)

The general idea behind this suggestion is plausible. It is in need of some refinement, however, if it is to be coherent at all, and the way in which my examples have been breaking down required donations in \$5 increments actually exposes the reason for this. Suppose that for me, at the moment, a \$5 donation is below the maximum delta that can be required, but a \$20 donation is not. Despite my conviction that the precision it creates is illusory, let us put numbers on these suggestions; my current state of well-being is 10000 Bentham's, the maximum decrease that can be morally required of me in any one act is 10, and the \$5 donation requires a sacrifice of 6. The

\$20 donation would therefore require a sacrifice of just a little more than 24 Bentham's, let's say 24.5 (more than just four times that of the \$5 donation, as the less money I have available, the less I can afford to sacrifice more of it). However, suppose I do donate the \$5. It is not at all implausible that I would still be required to donate another \$5 at this point. And another, and another, and another, each such donation decreasing my well-being by just a little more than 6 Bentham's. But now I have been required to spend a total of \$25 in this way, when I could not have been required to spend \$20 all at once, and it is clear this process is going to keep going for some time. It continues for longer still if I make the donations 1¢ each rather than \$5. That we could be required to donate \$5 five or more times in a very short time (a few minutes, say) but not \$20 all at once is preposterous verging on incoherent, and should clearly be rejected.

Now, we might try to consolidate these cases, claiming that the difference between contributing \$5 four times and \$20 once is illusory. We could, for example, make the delta limit relative to time rather than specific actions; I need not reduce my position on the personal scale by more than so much per day (or other unit of time, but I will assume a day in what follows), as opposed to merely being exempt from any action which would reduce it by that much in one swoop. But this suggestion has many problems. Firstly, suppose I encounter a UNICEF box one afternoon at 2:30 and make the required donation for that day (say we have precisely calculated this to be

\$6.72). The following day I encounter the same Unicef box at about the same time of day. According to this suggestion of a time-based delta, in order to determine what I need to do, I should first check exactly what time it is. If it is 2:20 PM, I'm off the hook. However, if it's 2:40 PM, I need to donate again. Secondly – as though that weren't silly enough – suppose that, the night in between, I see a baby about to drown in a small pond, and am uniquely situated to rescue it, at the expense of merely getting my pants wet.⁵⁵ This is fairly clearly an example of an action I ought to do by most standards, but under this proposal, whether this is the case seems to depend on what other good deeds I have done recently, and at what cost. But this seems irrelevant – the cost of *this* action is negligible, and that is what seems relevant. Such a result is unacceptable for, broadly, the same reason as the murdering-the-reader cases the hybrid theory produced – a result which would appall consequentialists and their opponents equally is not progress.

This problem is not quite as bad as it may sound at this point, since the suggestion now on the table restricts, not the absolute drop in our well-being that can be asked of us, but its ratio with the impersonal good we do; for this reason, it is not as though merely giving the required donation earlier in the day can get me off the hook for saving the baby. But it still leads to unintuitive results; doing enough *more* than my share earlier in the day can still justify my ignoring the drowning baby. While this

⁵⁵ This is of course a classic example from Singer's "Famine, Affluence and Morality" (*Philosophy & Public Affairs* V.1, 1972), used again in Unger's *Living High and Letting Die* (9).

idea – that doing more than our share at one point entitles us to do less at another – does have some intuitive plausibility, surely in at least this particular case, *no-one* would endorse letting the baby drown. Indeed, our reaction, on finding out about my various other good deeds, would no doubt be to question the intent behind the latter, not to judge me any less harshly for what I have just done.⁵⁶

I like the delta suggestion very much, and was prepared to defend it here until I noticed the above problems. I remain cautiously optimistic that the impulse behind it can be given coherent expression without falling prey to this problem, but I do not, at the moment, know what such expression would look like. I must therefore pronounce this attempt to combine a principled distinction between right and wrong with scalarity a failure. But it deserves to be said that I do this only with great reluctance.

Scalarity and Specificity

It is difficult to say where this leaves us. The scalar view seems in one respect to be exactly what we have been looking for – a consequentialist (not even pseudo-

⁵⁶ Another problem with the delta view: Suppose things are *really* going down the drain. No matter what we do, the world after our action will be worse than the one before it. The delta view, as I have stated it, implies that we have no moral obligations, and might as well do all we can to save our own skins; we can't make things better in disinterested terms, so while it would be nice of us to mitigate the damage, nothing in the delta theory seems to force us to. Is this plausible? I suspect it depends on the details of the example, but I submit that the answer will, at the very least, not *always* be "yes".

consequentialist, but outright consequentialist) view which is not vulnerable to the Demandingness Objection. Yet it seems oddly unsatisfying in the sense that it does not actually tell us that such-and-such an act is right or wrong. Most of this chapter so far has been taken up by unsatisfactory attempts to address this worry, a goal which, as of the end of the previous section, we seem to have little reason for optimism about achieving (though I chose to cling to some anyway). And on reflection, this seems to be a symptom of a larger lack of action-guidingness on the part of the scalar view.

I am not sure if a more satisfactory approach is possible, and the next (and penultimate) section of this essay will present some thoughts on why the problems with the scalar theory might not, after all, be problems. But before getting there, I would like to sketch out what features a theory would need to have in order to satisfy our desire for solid action-guidingness.

For some time now, I have been analyzing the problem the demandingness objector raises in terms of two (almost)⁵⁷ equally valid, yet often conflicting, sources of reasons, morality and practical rationality – an approach inspired by Scheffler's hybrid theory, though the last few variations I have considered have little else in common with Scheffler. On any theory I have considered that acknowledges both,

⁵⁷ I say “almost” because everything remotely resembling a tie has been broken in favour of morality, at the expense of practical rationality.

the attempt to meet the demandingness objection consists largely of acknowledging that morality need not be supreme in our decision-making; compromises between the two scales are possible. Each suggestion proposed one or more principles that might govern the interaction between the two scales; for example, the Restricted Delta view gave two such principles, that we may not choose one action over another *strictly superior* action (which might seem fairly obvious), and that morality can demand a certain level of sacrifice from us, which is proportional to the good we produce, but no more (pretty clearly a more substantive principle). The details are for the moment less important than the point that neither scale alone tells us what we need to do – such decisions depend on the other scale, and on the alternatives available.

Now, this general approach, it seems to me, need not be hostile to the scalar view. According to scalarity, morality is only in the business of giving us general advice. But the whole approach described above seems to be committed to saying this anyway. The approach of the above paragraph is a little stronger – it will still insist that *all else being equal* we should do the morally better among our available options, while the scalar view need not make even this fairly minimal claim – but nevertheless the two seem to be fundamentally compatible. From the mere fact that *morality* does not tell us what to do, it does not follow that our normative theory as a whole shares this lacuna.

In this way, accepting the scalar view of morality need not stand in the way of, say, a further development of the Restricted Delta view. And this means that it is possible in principle to meet the objection introduced three sections ago. The details, unfortunately, remain to be filled in – though the Restricted Delta seems to me a potentially important step toward doing so.

Non-Coercive Moral Theory

I have just outlined a reason to think that, even on the scalar view, we need not give up on getting specific advice out of our normative theory when taken as a whole.

Yet, suppose this does after all prove impossible. How bad is this? In this second-last section of this essay, I will suggest that the answer is, not very.

What is it for a theory to made demands of us, anyway? Presumably, for it to say we must do something or refrain from doing it. (Various conditions might be added to this, such as the “demands” not consisting of simply doing what we would have done anyway. If, for example, we accept both psychological and ethical egoism, the latter makes very few if any demands, in this sense. But the condition I have given is surely at least *part* of any reasonable definition.) But really, what is the worst that can happen as a result of our failing to abide by such an instruction? Failing to abide by the law carries the possibility of legal sanction; failing to abide by rules of etiquette carries the possibility of incurring the disapproval of others; failing to abide by the

dictates of practical rationality would seem to be, at least in theory, its own punishment. What similar comments can be made about failing to abide by some moral theory, even one that we sincerely believe in? Precious few, save where it coincides with one or more of the other things I have mentioned, or reaches such an extreme as to actually make others fear for their own safety (and that speaks to the self-interest, not the moral sentiment, of others).⁵⁸ Speaking in particular of those failures to abide by the harshest of consequentialist requirements which do not also incur the wrath of ordinary moral views, we might be a little disappointed in ourselves, and serious consequentialist thinkers of our acquaintance might urge us to do better. But we soon get over it, as do most of those doing the urging. The consequences for ourselves are slight compared to those of, say, forgetting an anniversary; and the consequences for the rest of the world are, if unfortunate, psychologically very easy to simply not take responsibility for.

This does not mean that we have no reason to follow a consequentialist moral theory; I still think we're on reasonably solid ground endorsing PTR, and haven't yet

⁵⁸ Perhaps it's a little cynical, but I believe the latter, which is distinct from moral revulsion, explains why people, say, shun murderers to a greater degree than they do those who pick their noses in public. And of course, not everyone *does* shun murderers; many people are fascinated with, and some even seek out, people who have done great moral wrongs, whereas virtually nobody reacts to people who pick their noses in public in such a way.

It should also be borne in mind that I am speaking of *consequentialist* moral standards here, and people certainly *don't* shun people who are, by consequentialist lights, morally equivalent to murderers – after all, virtually all of us are in this category.

found a good reason to supplement it with other moral principles *other than* those which I think are better regarded as belonging to practical rationality. But it does underscore the point that the sanctions that we suffer by *not* following it are, in the end, rather weak. While I hesitate to draw any conclusion from this with any certainty, it does seem suggestive in at least one respect. Having acknowledged this, it does not seem so bad to endorse a theory that claims only to be offering us advice, to do with as we see fit. Moral theories, we might say, are in practice, only in the business of doing this anyway. From this point of view, to object to such a theory, merely on the grounds that it admits this, suddenly seems rather silly. We might even go so far as to say it is the desire to have such a theory lay claim to greater coercive power than this – power it cannot possibly possess – that seems to need explanation and justification.

I do not pretend this is a knockdown argument but there seems to be a case here that we shouldn't, contrary to the objection raised early in this chapter, be too upset by this aspect of the Scalar View. If this is true – or the project outlined in the previous section is viable, or both – then the Scalar View seems after all to have survived as the answer to the Demandingness Objection that the consequentialist has all along been looking for.

Conclusion

About eleven years ago, on learning that I was for the time being leaving academia, a moderately well-known philosopher told me I was a good student, but the biggest factor holding me back from being a great one was my attachment to my own views. Since returning to philosophy several years later, I have made considerable effort to avoid falling into the traps that were pointed out to me back then. In some ways, this thesis represents the culmination (at least for the time being) of that process, for I have spent virtually this entire work arguing views that are not even mine.

The views I believe, or fear, are most likely correct are as follows. I think the most plausible moral theory is almost certainly some form of consequentialism, one that does not stray far, if at all, from utilitarianism. I suspect that the correct response to the demandingness objection is just to accept that such a theory is very demanding and work to do much, much better than most of us in fact do. I think that, if sufficiently many people contributed sufficiently close to optimally, the “dystopia” I sketched out in chapter 1 actually wouldn’t be so bad. Certainly it would be far preferable to the world we now live in.

Why, then, have I spent virtually this entire work as an advocate for a largely non-consequentialist Devil? Because, unlike the case with some of my views of a decade

ago, I am not fanatical or wholly unwavering in these convictions. I am sympathetic to the view that there must be something wrong with a theory that asks so much of us, even though I think that line of reasoning is probably mistaken. And I think resolving these conflicting impulses is not only important for getting my own views sorted out, but also happens to be the most important problem in moral philosophy right now. And so I have presented selected parts of the best case that I think can be made, without resorting to the witch-doctory of a Kamm, for this other view I would like to be correct.

I am not saying that my arguments have been at any point wholly disingenuous; I do think that pure consequentialism is potentially alienating, I do see something unsatisfying about the scalar view and have legitimately done my best to identify what it is, and I really would like it best if something like the Restricted Delta were correct. It has been my goal to make the case for these aspects of my views as strongly as possible – and in places that has proven to be strongly indeed. If we compare the totality of my various and conflicting moral views to a panel of judges deciding a case, this thesis is the opinion of the dissenting minority. That minority may not win the day, but still needs to be taken seriously.

Still, it is undeniable that the Scalar View does get around the demandingness objection, and plausible that the Scalar View, despite how things might initially

appear, is otherwise acceptable; I have gone some way, in the last two sections of chapter 4, toward talking myself out of the unease it occasions in me. This leaves me uncertain whether to simply accept the scalar view, or continue biting the bullet and accept a harsher version of consequentialism, one which makes equally harsh judgments on many of my own actions. If I have clarified some of the issues surrounding this decision, however, I will have accomplished something significant; and I think it is fair to say I have managed that much.

Appendices

Appendix A: Metaethical and Methodological Assumptions

Here and there, I have made assumptions which may be controversial, and I have also referred obliquely (or, in one or two cases, quite explicitly) to a test by which I believe intuitions may be measured, that we may determine which to admit into our moral theorizing and which to reject, or at least take with a grain of salt. In this appendix, I will discuss the thinking underlying my belief that these assumptions are admissible, and give a brief statement of what the test in question actually is.

Some philosophers think of ethics in terms of explaining our existing moral intuitions. Some such philosophers – let Frances Kamm represent them – proceed in whole or in significant part by examining intuitive responses to imagined cases. If we think that a given act is permissible in a given case, this creates a strong presumption – and perhaps, under some circumstances, even *proves* – that the act *is* permissible; the same is true of acts we judge to be either mandatory or impermissible. Moral theory is then a matter of noting any patterns that these intuitive responses form; hopefully any such patterns help us answer the difficult questions that arise in those cases where our intuitions are not so clear.⁵⁹

⁵⁹ Kamm explains her methodology, which is of the general sort described here, in many of her writings; I have used *Morality, Mortality Volume II*, 10-13 as my guide.

A problem with this approach should be apparent from comparing almost any randomly chosen page of Kamm with some of the remarks I have made in this thesis. A different philosopher – Norcross, say, or myself – could use Kamm’s method and Kamm’s examples, and yet end up with a completely different theory than the one she defends. Though Kamm makes a few remarks about “sufficiently many”⁶⁰ people needing to share the intuitive responses she reports for them to be valid as data – and even then, allows for their fallibility – in the main body of *Morality, Mortality* and other work of hers with which I am familiar, she makes little or no allowance for the possibility of disagreement.⁶¹ Where Kamm and I disagree, it is not clear to me that either of us is in the majority in our moral verdicts, nor – more crucially – is it clear why any weight should be placed on majority opinion in any event. Kamm herself believes she is engaged in an *a priori* enterprise⁶² – why, then, should *empirical* facts about our responses to cases matter?

⁶⁰ *Ibid*, 11

⁶¹ In fact, even in the introductory matter discussed here, she never acknowledges the possibility of *diametrically opposed* intuitive responses to her cases, only of nonexistent or unclear ones.

⁶² Kamm 11, just two short paragraphs from the “sufficiently many” comment. More recently (*Intricate Ethics* (Oxford: OUP, 2007)), Kamm explicitly confines her epistemic basis to *her own* intuitions – which, while possibly getting away from the problems raised in the main text, does so at the cost of calling the general applicability and even interest of Kamm's theorizing into serious question. It thus seems to me that this newly revised (or clarified) version of Kamm's methodology is even more problematic, so (though this remark may itself come across as a swipe at Kamm) it seems only charitable to ignore it.

Kamm's response to the last question would no doubt be that the empirical fact is merely an indication (and a defeasible one at that) that the *a priori* truths of which she speaks are a certain way; we are aware at some level of at least some moral truths, and such intuitions reflect this fact. Let us assume this doctrine is correct.⁶³ Even if we make this assumption, the possibility – indeed, the reality – of radically opposed intuitive responses to cases presents multiple problems. The most serious of these is epistemic – how are we to know *which* of the opposed intuitions, in cases where they exist, is the correct one? Even on Kamm's view, intuitions are at best *indicators* of the underlying moral reality; they are by no means constitutive of it. So presumably even Kamm would have to admit that an opinion poll, for example, would not settle the issue (the "sufficiently many" comment suggests she thinks it might, if the results were *very* lopsided, but even this is not obvious to me). Presumably there is some feature of the *a priori* fact at issue which would distinguish the correct view from incorrect ones. But Kamm has done little to enlighten us on what it is supposed to be, much less given any cause for optimism about gaining epistemic access to this feature. If intuitions do not settle the issue, it is not clear what, if any, other tools we have with which to do so. In cases where our intuitions

⁶³ We leave the reasons why these moral truths are truths, the mechanism by which we become aware of them, and the reasons for the – as even Kamm acknowledges – partial and imprecise nature of this awareness all unspecified; we similarly gloss over the issue of why, given these limitations, we should *trust* this awareness. All of these are important issues, but they would simply take us too far afield, and it is not necessary to explore them in order to describe what I consider the *main* problem with Kamm's method. I will say that I believe all of these points represent *further* problems for Kamm (and others), some of them quite serious.

are merely unclear, she can derive some sort of principles from the cases where they *are* clear and apply them to the unclear cases; but in cases where they actively conflict, it seems to me that her methodology as written can't even get off the ground, for *which* intuition do we choose to proceed from?

Perhaps we might try to respond to these worries by abandoning the use of intuitions as the main unit of epistemic currency in moral philosophy. While intuitions may provide a starting point, we could seek to admit as moral principles only those intuitions which survive rigorous analysis, declaring that mere intuitive appeal carries no, or at best very little, weight in determining the correctness of a particular bit of moral theory. It is not clear how we are to get off the ground without giving our prior views on right and wrong *some* place in moral theory, so intuitions need to play some role. But with this conceded, we could still try to make this role as small as possible.

Shelly Kagan endorses the closest possibly-viable thing to rejecting intuitions entirely in *The Limits of Morality*.⁶⁴ Intuitions, on this view, must be capable of being rendered coherent by theory; those which can't should be rejected. This coherence, someone sympathetic to Kagan (a category in which I find myself) might charge, is precisely what Kamm's endless epicycles lack; her unwillingness to approach intuitions (at least her own) critically is her undoing. But as much as I might prefer it

⁶⁴ *The Limits of Morality*, 11-15, cf (for example) 114, 376;

to Kamm's, it seems to me Kagan's approach immediately runs into a problem; very basic bits of theory that even Kagan presumably wants to allow in end up being shut out by this methodology.

Consider PTR, which at this point should need no introduction; and recall that this principle was explicitly borrowed from Kagan.⁶⁵ Kagan defends a moral view that admits only PTR as a moral principle, in other words, a consequentialist one by my lights. But though I think *The Limits of Morality* largely succeeds in his immediate goals, a conspicuous feature of the book is that it does not even try to give us reason to accept PTR itself. Kagan merely argues that the "moderate" and the "extremist" alike are committed to PTR, but the moderate is also committed to additional principles; these additional principles do not stand up to scrutiny; therefore PTR survives as the whole of morality.

As I have said, I think this is more successful than not as far as it goes. But it is unclear at best what becomes of this argument if we deny PTR as well, and nowhere does Kagan provide grounds for accepting PTR other than the putative fact that it is common ground between the moderate and the extremist.⁶⁶ I have pointed out that

⁶⁵ PTR, and some other features of Kagan's terminology, were introduced late in the first section of chapter 1.

⁶⁶ In fairness to Kagan, this is quite adequate for his immediate purposes in *The Limits of Morality*, where the main point he is trying to make is simply that the extremist's position is more defensible than the moderate's. But in a more general context it is not so adequate, especially given that there are those, like Scanlon, who seem to deny PTR.

the PTR has great intuitive appeal – indeed, it strikes me as self-evident – but it is not clear that this could be *Kagan's* basis for asserting it, for he at least sometimes sounds like he would reject any approach that would allow that intuition to count as data. (He *could* argue that it is, uniquely, an intuition that can be made coherent by theory, in much the same way I am concerned, in the main body of the present work, to make a case for a claim in that neighbourhood. But he does not attempt to actually do this in any systematic way.)

Kagan gestures in the direction of justifying the PTR indirectly by appeal to what our motives would be under certain unspecified ideal conditions,⁶⁷ not dissimilar to the way in which Kamm justifies various principles indirectly by appeal to her intuitions (importantly, in both cases, the connection is claimed to be evidential, not constitutive). But by his own admission, what he offers barely qualifies as a sketch of a full, proper defence of the PTR; and it seems to me that even in that form, it raises a question – why are our intuitions about what is right not admissible, but our motivations in hypothetical scenarios fair game? And the question is especially sharp when we consider that whether one acts rightly would seem, a bit counterintuitively, to be largely independent of motivation – as Kamm would be quick to point out.⁶⁸ If

⁶⁷ *The Limits of Morality*, 387-389

⁶⁸ For this point, and many other astute observations, I recommend some of Kamm's very lucid applied ethics articles (e.g. "Is there a Problem with Enhancement?", *American Journal of Bioethics* 5 (2005), 5-14) over her highly arcane normative ethics.

someone finds a cure for breast cancer, would we think he acted wrongly if he was motivated by fame and wealth? Surely not.

But suppose Kagan is correct that moral motivation and moral reasons are intimately connected, such that the existence of the former (in the proper circumstances) is good evidence of the existence of the latter. Aren't our moral intuitions, at least, a good indication of what our moral motivations are? And doesn't this suggest that, even on Kagan's own view, intuitions do after all have a significant role to play in the justification of our theories, even if this is itself justified in a somewhat roundabout manner?

So treating our intuitions as more or less sacrosanct is deeply problematic, for on this approach we can neither resolve nor account for disagreement (to say nothing of the endless Ptolemaic epicycles we need to add to our theory if we take this approach to its logical conclusion, which is to say, use Kamm's method). But to discard intuitions altogether would be even more problematic, and even the closest thing that looks superficially viable – the Kagan-type approach I have discussed – seems to have great difficulty even getting off the ground; it seems impossible to do moral theory without giving intuitions great evidential weight.

Obviously enough, however, there are many ways to find a middle ground between these two views. One such way can be derived from Peter Unger's work⁶⁹ and involves giving such weight to some intuitions and not others. More specifically, we can distinguish between our intuitions about individual cases and our intuitions about general principles. Where the two conflict, there seems to be some reason to favour the latter; Unger cites numerous potential distorting factors which seem to be of no moral relevance, yet appear to affect our intuitions about individual cases, such as the spatial relationships between the people involved in a given case. This seems to be less true of intuitions about more general principles. In other words, the more general intuitions, about broad principles rather than specific cases, seem less likely to be accepted or rejected on morally irrelevant grounds.

For example, suppose that we can pay \$5 to save the lives of two strangers in some faraway country – a supposition that may very well be true. It is obvious, from their behaviour if not so much from their responses to hypothetical cases, that most people in our society do *not* think it is morally obligatory to pay the \$5. Yet the same people who hold this view would, without a doubt, be prepared to pay the same \$5 if the two people were distant acquaintances in the same city. And while many

⁶⁹Namely, *Living High and Letting Die* (New York: OUP), 1996, with the distinction between moral intuitions and moral common sense introduced at 28. Since these seem to me to both be sorts of intuitions, and moreover, to be on a continuum rather than sharply distinguished, I have opted to simply *call* both of them intuitions. More general intuitions of the sort Unger and I approve of correspond to Unger's "moral common sense" while those about specific cases, which both Unger and myself are relatively dismissive of, are the ones *he* calls "moral intuitions".

people do hold that we owe special obligations to those to whom we stand in certain relations (children, spouses, etc), it seems implausible to extend them to distant acquaintances or to people who merely happen to live in the same city. Does anyone really think that, say, a brief conversation in the line-up for tickets to a movie can create a moral obligation with such stark implications?

On the view that is beginning to take shape, in this case, our intuitions about a specific case conflict with the more general intuition that \$5 is a small price to pay to save two people's lives; and the specific intuition should, in general, lose out to the more general principle. In other words, all else being equal, more general principles win in conflicts of intuitions.

One sceptical about this approach might respond that things are not quite so simple, for the intuition that we owe less to those further away – however silly it may sound when expressed as starkly as it was just now – is nevertheless a strong one, as shown in our actual, everyday actions. Indeed, it seems, in this sense at least, to be at least as strong as the conviction about saving lives on the cheap. And I take it as given that the strength of an intuition – the degree of certainty with which we hold it, as manifested by either our words or our deeds – represents a criterion that is relevant in the conflicts we are discussing, in the same way I have been suggesting that

generality is. If so, we seem to have reason to regard this intuition as morally relevant after all.

Though such responses have from time to time been seriously proposed in the literature, I find this particular one, at least, utterly unconvincing, and I think I can say why; and this will give us a third criterion for evaluating intuitions. It is just too easy to explain this intuition about physical and/or cultural distance away in nonmoral terms, in a way that seems far less compelling (though, granted, still possible) in the \$5-for-two-lives case. First, the intuition being proposed here smacks of tribalism – it makes sense to suppose it evolved under prehistoric conditions in which small, fairly primitive groups had to compete against one another for survival. Favours one's own group in this way, once no longer a compelling necessity for survival, hardly seems like a moral value, or even a particularly reasonable self-interested one; at best, it falls into the category of *understandable* wrong actions that I had occasion to mention in chapter 3, or so this line of thinking suggests. Additionally, we might think we have self-interested reasons to limit our charity in this way, and while it is (as I suggested in the main text) plausible that this is a legitimate reason to limit our obligations somewhat, it is not by any stretch of the imagination a *moral* reason to do so, nor does it in any other way imply that our moral obligations to the distant strangers are less than those toward our fellow citizens. In fact, it looks more like a *diagnosis* of our apparent holding of the belief

that physical or cultural distance matters morally than a *justification* of same – we believe this (on the plausible view to which I am alluding) only because it is in our best interests that we do so – indeed, precisely because it avoids the unpleasant implications about our obligations that I described in the main text.

I call the criterion this suggests *explicability*, but by this I mean not so much our ability to explain an intuition as our ability to explain it *away*, or in other words to come up with nonmoral reasons why we might hold the view it describes despite it failing to describe a genuine moral reason.

Now, the three principles I have mentioned – strength, generality and explicability – seem to form a reasonable system for evaluating any given moral intuition for plausibility, yet it might be objected that already I have stacked the deck against the nonconsequentialist advocate of the demandingness objection. After all, the examples of conflicts I have used so far, and indeed most of the obvious ones (and *all* of Unger's, for that matter) are ones in which the nonconsequentialist does poorly. This would be a plausible charge if these principles had been *motivated* by a desire to achieve that result, but all of them seem independently plausible; one certainly does not need to be a consequentialist to accept any of the reasoning I used in their defence (sketchy though that defence has admittedly been). And in any event, I do not in fact believe, as Unger (for example) appears to, that all or nearly all of the

suggested intuitions that survive the process of comparison to their potential rivals using these criteria are consequentialist. Indeed, chapter 2 revolved around a large category of cases in which the non-consequentialist, using the method I have sketched, does very well indeed.

Appendix B: Contextualism?

In chapter 4, I attempted to find a view that accommodated both Norcross' scalar view of utilitarianism and a "bright line" between right and wrong, failing (though able to show that such seemed possible in principle) and ultimately questioning the very desirability of such a project. There is another view on mixing scalarity with a line (of sorts) between right and wrong, known as contextualism. While it comes the closest of anything I've seen to allowing for both, I believe that it raises more questions than it answers. Since Norcross endorses a contextualist approach to this problem, it would be a significant omission if I failed to say a few words about the subject, though the only conclusion they will lead to (in which I am at all confident) is that more work needs to be done in this area.

I mentioned more than once, and even rested weight upon, the observation that our moral expectations are much higher in (say) the Achilles' Computer case than in any real-life case anyone likely to read this essay will ever have faced. Achilles, it seems, can act wrongly even by doing something with far better results than you or I have ever achieved, which, I have argued, is because he could do better still at no cost to himself. I also suggested that this is an instance of a more general principle, that the moral value of a given action depends on the alternatives available to the agent. In contrast to the elaborate theoretical machinery I tried out in an attempt to further

explain this fact, the contextualist does not think any such further explanation is needed or possible – or, at the very least, does not think any *one* such explanation is uniquely correct.

Instead, the contextualist believes, as the name suggests, that our judgments on whether some action is good or bad, right or wrong, will depend on the context in which they are made. Sometimes an act will produce immensely good results and we will still say it was wrong, as in Achilles' Computer; at other times, an act with bad results might still be said to be right. To the contextualist, what can be said, objectively, about right and wrong begins and ends with such observations about how we use the words in different situations. There are, in other words, no objective facts about right and wrong *themselves* (whether there are moral facts at all – say, about value – is something different contextualists have different views on, as I will discuss momentarily). In the main body of this thesis, I tended to assume, at least until chapter 4, the existence of some correct account of right and wrong that explains how the superficially incompatible judgments I have mentioned can all be correct, or else shows us which we should reject. The contextualist says there is no need for this. We can make observations about what judgments we make in what contexts, but such observations are *only* about that – if there is an objective moral reality, this is not the way to discern it (in sharp contrast with, for example, what Kamm must claim).

There are at least two distinct breeds of ethical contextualism, one championed by Norcross, the other by Unger.⁷⁰ On Norcross' view, where we draw the line between right and wrong, or good and bad, actions, is a purely semantic matter, which can legitimately vary with the context in which the question arises. The ranking of actions remains constant; action A can't be better than action B in one context and worse in another. In other words, there are objective facts about (at least relative) value but not about right and wrong. All moral judgments are comparative for Norcross (hence his scalar view, around which chapter 4 revolved), with the relevant comparisons involving other actions that could have been performed in the same situation; *which* alternative action is the most salient, however, is what varies. In Achilles' Computer, we judge actions based on what Achilles could do without significant personal cost; other standards might prevail in, for example, the rescue-mission cases so popular in the literature, in which (judging from a common-sense standpoint) we might be more lenient in the sense of not requiring strict optimization, yet less lenient in point of requiring significant sacrifices. Two people can disagree on whether an action was right, and each be correct by *some* legitimate set of criteria (though obviously not the *same* one – there is bound to be a certain amount of talking past one another going on in such a situation); but if they disagree

⁷⁰ Mark Timmons' contextualism isn't so much ethical as metaethical, and thus will not concern us here.

on whether that action was better than some alternative, at least one of them is simply mistaken.

On Unger's view, even the ranking is a context-dependant matter. Various features – not of the situation in which we make a decision, but seemingly of the *conversational* context – can make entirely different features morally salient. There can be contexts in which some roughly consequentialist view, such as the one he spends the first six of *Living High and Letting Die*'s seven chapters defending, is appropriate – these are contexts in which we emphasize what he calls our Primary Values – and others in which the ordinary, everyday moral views on which we have few if any obligations to far-away unfortunates are perfectly acceptable – where, in his terms, our Secondary Values are at the fore, or even a version of them adjusted further according to social norms and the psychological difficulty of being motivated by those values.⁷¹

Of these two views, I am fairly sure I at least *understand* Norcross', but I find Unger's puzzling. If we are sometimes right, even semantically, to judge things by the ordinary views on which, for example, our lack of support to poorer nations is our entitlement, then I am somewhat lost as to what it is Unger has been doing in the rest of the book. Much of the point of those first six chapters is quite explicitly to influence his readers' decisions in the direction of a more generous attitude toward

⁷¹ *Living High and Letting Die* chapter 7, especially 167-170.

those in distant nations and, to a certain extent, a more flexible attitude toward things like lying and stealing (when other, more important goods are at stake). Yet, it is unclear why we can't just say that, while in a certain context that is a very desirable thing, that is not the context in which we actually make the relevant decisions. It is fairly obvious that Unger is trying to avoid the Demandingness Objection in this chapter; but it seems to me he can do so only by reducing the whole notion of morally better and worse actions to subjective preference, if not wholesale meaninglessness. If Unger wants to say we should prefer a certain context, that in itself seems to be a normative, probably even a moral, judgment, about which we can ask – in what context does he make *that* judgment?⁷² (Or if it somehow transcends context, then what does he defend contextualist semantics for in the first place?)

Thus, of these two versions of contextualism, I find Norcross' much the more plausible. But this should not comfort anyone who thinks a "bright line" is an adequacy condition on a moral theory, that is, anyone who holds the view I spent the first part of chapter 4 discussing responses to. It seems to me that contextualism, even in this form, does not so much give us a way to draw such a line as acknowledge a complete inability to do so. The objection I pretended to endorse in the second section of chapter 4, recall, was precisely that there seem to be clear

⁷² For that matter, so does the judgment about what comparison is most relevant in a given context under Norcross' contextualism. I thank Jacob Kolman for pointing this out.

examples of right and wrong actions, and we want our moral theory to say so. This, it seems to me, contextualism constitutes a principled *refusal* to do. This is not to deny that there are things that will come out right or wrong in any context we are actually likely to encounter (even on Unger's view, for he acknowledges that some things can't plausibly vary much with context, in the same way we would never call the Himalayas "flat" despite considerable contextual variation in what we mean when we say "flat"). But the contextualist *qua* contextualist can say nothing, on pain of the inconsistency implicit in smuggling in non-context-sensitive moral judgments, about which they are or how we can know, other than simply by seeing what judgments we make in what situations. This would seem to be exactly the opposite of what the objector at that point wanted. I do not claim to have a better answer to the objection, but whatever the best response may be, it clearly isn't contextualism.

Bibliography

The following is a list of works that influenced this essay in some manner. Most, but not all, are explicitly called out in the footnotes at some point; others are referred to vaguely and obliquely, so as to not warrant specific citation, but to still deserve acknowledgment here. Still others have inspired the overall direction of some aspect of the work without actually being specifically cited at any point (for example, the shadow of Peter Singer is *very* visible throughout this essay even though he's only explicitly mentioned once). In a few cases (e.g. the Hurley article), a detailed discussion of a work was planned at one point and omitted from the final version, but a paragraph or two that was originally written toward this end has survived; this section should serve to acknowledge the origins of such passages.

My memory being what it is, this list is probably incomplete; if you discern the influence of some philosopher not listed here, there's a good chance you are correct, and I assure the reader that any such slight against some influence of mine is wholly unintentional.

Ordering is strictly alphabetical, by author's surname then title. No publication information is given for anything over a century old, since such works are usually available in many different editions (and with a standard pagination).

Jeremy Bentham, *Introduction to the Principles of Morals and Legislation*

Thomas Hobbes, *Leviathan*

Paul Hurley, "Does Consequentialism Make too Many Demands, or None at All?" in *Ethics* V. 116 (2006)

Shelly Kagan, "Does Consequentialism Demand too Much?: Recent Work on the Limits of Obligation" in *Philosophy and Public Affairs* V. 13 (1984)

Shelly Kagan, *The Limits of Morality* (Oxford: Clarendon), 1989

Shelly Kagan, "The Limits of Well-Being" in *Social Philosophy & Policy* V. 9 (1992)

Shelly Kagan, "The Structure of Normative Ethics" in *Philosophical Perspectives* V. 6 (1992)

Frances Kamm, "Is there a Problem with Enhancement?" in *American Journal of Bioethics* 5 (2005)

Frances Kamm, *Morality, Mortality Volume 2* (New York: OUP), 1996

Frances Kamm, "Non-Consequentialism, the Person as an End-in-Itself, and the Significance of Status" in *Philosophy & Public Affairs* V. 21 (1992)

Immanuel Kant, *Foundations of the Metaphysics of Morals* (a.k.a. *Groundwork of the Metaphysics of Morals* or *Fundamental Principles of Morals*)

Immanuel Kant, "On a Supposed Right to Lie from Altruistic Motives"

John Stuart Mill, *Utilitarianism*

Tim Mulgan, *The Demands of Consequentialism* (Oxford: OUP), 2002

Liam B. Murphy, *Moral Demands in Nonideal Theory* (Oxford: Clarendon), 2000

Thomas Nagel, *The View from Nowhere* (New York: OUP), 1986

Alastair Norcross, "Good and Bad Actions" in *The Philosophical Review* V. 106 (1997)

Alastair Norcross, "Harming in Context" in *Philosophical Studies* V. 123 (2005)

Alastair Norcross, "The Scalar Approach to Utilitarianism", in Henry R. West (ed.) *Blackwell Guide to Mill's Utilitarianism*, (Oxford: Blackwell, 2005)

Alastair Norcross, "Two Dogmas of Deontology", forthcoming in *Social Philosophy & Policy*

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

Derek Parfit, "Equality and Priority" in *Ratio* V. 10 (1997)

Derek Parfit, "Justifiability to Each Person" in *Ratio* V. 16 (2003)

Derek Parfit, *Reasons and Persons* (Oxford: Clarendon), 1984 (revised 1986)

Plato, *Protagoras*

Peter Railton, "Alienation, Consequentialism and the Demands of Morality" in *Philosophy and Public Affairs* V. 13 (1984)

John Rawls, *A Theory of Justice* (Cambridge, MA: Harvard University Press), 1971

W.D. Ross, *The Right and the Good* (Oxford: OUP), 1930

T. M. Scanlon, *What We Owe to Each Other* (Cambridge, MA: Belknap), 1998

Samuel Scheffler, *The Rejection of Consequentialism* (Oxford: Clarendon), 1982 (revised 1994)

Robert Shaver, "The Appeal of Utilitarianism" in *Utilitas* V. 16 (2004)

Robert Shaver, "Sidgwick's False Friends" in *Ethics* V. 107 (1997)

Henry Sidgwick, *The Methods of Ethics*

Peter Singer, "Famine, Affluence and Morality" in *Philosophy and Public Affairs* V. 1 (1972)

Peter Singer, *One World* (New Haven, CT: Yale University Press), 2002

Judith Jarvis Thompson, *The Realm of Rights* (Cambridge, MA: Harvard University Press), 1990

Peter Unger, *Living High and Letting Die* (New York: OUP), 1996