# CONSEQUENTIALISM AND MORAL DEMANDS

BRIAN McELWEE

Submitted for the degree of PhD on September 27, 2005

I, Brian McElwee, hereby certify that this thesis, which is approximately 64,000 words in length, has been written by me, that it is the record of work carried out by me and that it has not been submitted in any previous application for a higher degree.

Date 28/2/06 Signature of Candidate

I was admitted as a research student in September 2000 and as a candidate for the degree of PhD in September 2001; the higher study for which this is a record was carried out in the University of St Andrews between 2001 and 2005.

Date 28/2/06 Signature of Candidate

I hereby certify that the candidate has fulfilled the conditions of the Resolution and Regulations appropriate for the degree of PhD in the University of St Andrews and that the candidate is qualified to submit this thesis in application of that degree.

Date 28/2/06 Signature of Supervisor

In submitting this thesis to the University of St Andrews I understand that I am giving permission for it to be made available for use in accordance with the regulations of the University Library for the time being in force, subject to any copyright vested in the work not being affected thereby. I also understand that the title and abstract will be published, and that a copy of the work may be made and supplied to any bona fide library or research worker.

Date 28/2/06 Signature of Candidate

## THESIS ABSTRACT

The aim of this thesis is to outline a form of consequentialism which denies the deeply unintuitive claim that we have a moral obligation to bring about the best consequences we can. Consequentialism should not be understood primarily as a theory of morality narrowly conceived as focusing on obligation, but instead as a theory of the goodness and choiceworthiness of our actions and practices, and of what there is most reason to do.

I begin from the well-rehearsed objection to consequentialism that it is unreasonably demanding, arguing that this constitutes a good objection to the theory in its traditional form, but showing that my favoured form of consequentialism- one which limits itself to claims about value and reasons- is not susceptible to it. I discuss criticisms of consequentialism from influential work in the second half of the twentieth century, showing how the strongest objections outlined therein apply only to consequentialism as a theory of moral obligation, and not to consequentialism as a theory of value.

Finally, I outline what a consequentialist should say positively about moral demands, explaining the limited role which the theory should have in shaping our moral obligations. I conclude that consequentialists should not be preoccupied with developing a distinctively consequentialist theory of moral demands.

## ACKNOWLEDGEMENTS

I would like to thank all those who have given me feedback on the work in this thesis, at graduate seminars, conferences and in less formal environments. I have received excellent advice and provocative discussion on my subject from a number of progress reviewers from the philosophy department at St Andrews, and from my fellow moral philosophy research students, in particular Simon Robertson. I am especially grateful to John Skorupski for his careful supervision, thought-provoking discussion and guidance over the past four years.

Research for this thesis was made possible by an SAAS/ AHRB studentship (2001-4), and a St Andrews Postgraduate Studentship (2004-5).

Finally, I would like to dedicate this thesis to my family and friends, in gratitude for several years of support and encouragement.

# TABLE OF CONTENTS

# CONSEQUENTIALISM AND MORAL DEMANDS

## INTRODUCTION

The aim of this thesis is to outline, and defend, a form of consequentialism which denies the deeply unintuitive claim that we have a moral obligation to bring about the best consequences we can. Consequentialism should not be understood primarily as a theory of morality narrowly conceived as focusing on obligation, but instead as a theory of the goodness and choiceworthiness of actions and practices, and of what there is (most) reason to do. I go on to outline what remains of moral obligation from the consequentialist point of view.

Much of the contemporary debate about consequentialism can be seen as focused on two main types of objection to the theory. The first line of criticism is that consequentialism allows us (perhaps even requires us) to perform actions that are morally wrong- in Shelly Kagan's terms, it ignores the moral *constraints* that feature in common sense moral thought. The second is that consequentialism makes demands on us that are wholly unreasonable- as Kagan puts it, it ignores the moral *options* that we ordinarily recognise. My main project in this thesis is to tackle this second set of objections by endorsing the central thought that lies behind the objections, but by arguing that consequentialists are not committed to the claim- namely, that we are morally obliged to bring about the best consequences we can- which would make them vulnerable to this demandingness objection. Rather than attempt to construct a less demanding consequentialist moral theory, I argue that consequentialists should limit themselves to offering a theory of value and reasons for action.

The structure of my thesis is as follows:

CHAPTER 1: THE POSITIVE CASE FOR CONSEQUENTIALISM

While the main aim of my thesis is to fend off objections to consequentialism, in this chapter I give a brief summary of what makes the position so attractive in the first place. Even opponents of consequentialism generally acknowledge that the theory has a lot to be said in its favour, and has significant prima facie plausibility. My aim here is to outline some of the different sorts of positive

arguments which are thought to give it this plausibility. I first consider what I call structural arguments for consequentialism, then go on to discuss arguments grounded in a particular conception of the good, including what might be termed meta-ethical arguments.

## CHAPTER 2: THE DEMANDINGNESS OBJECTION

I take as my starting-point the demandingness objection to traditional consequentialism. Take the following position:
**Moral Consequentialism (MC)**: We have a moral obligation to bring about the best consequences we can. The right action is that which brings about the best consequences. All other actions are wrong.

Accepting MC leads to some deeply counter-intuitive results. For example, someone who gives up most of his time and money for the good of others, strives copiously to bring about large amounts of good, but does less than the most he can, has, according to MC, failed in his obligations and is acting morally wrongly. This seems totally detached from our common sense moral thought.

I argue however that there are positions that are recognisably consequentialist, but which do not fall foul of this objection. Consider, for example, the following position:
**General Consequentialism (GC)**: The best thing we can do is bring about the best consequences we can. What we have most reason to do is bring about the best consequences we can.

Now such a consequentialism has, as yet, made no demands whatsoever, and so does not, or at least not obviously, demand too much. I attempt to defend GC in later chapters, while arguing against MC. I suggest that a consequentialism which *demands* the best available consequences is counter-intuitive because of the conceptual link between moral demands and blameworthiness. In response to the demandingness objection, Michael Slote proposes a position which he calls 'satisficing consequentialism', which identifies right actions or permissible actions as those which bring about *good enough* consequences. In showing how Slote's account fails, I highlight the importance of taking account of the notions of effort, cost and sacrifice in considering whether we have met the demands of morality.

Given that our moral obligations then are determined in part by the effort, cost or sacrifice involved in performing a particular action, one might think it open to the consequentialist to work these notions directly into an account of moral rightness and obligation, simply as a second set of relevant factors alongside consequences themselves. I argue that such a project is misconceived, by briefly examining one such attempt by Bjorn Eriksson. I point out particular problems with Eriksson's account, before going on to emphasise structural problems which face any such account, and suggest that considerations relevant to objective choiceworthiness (i.e. consequences) should be kept firmly apart from those relevant to judging blameworthiness (e.g. effort, difficulty, cost, sacrifice).

## CHAPTER 3: MORAL REQUIREMENTS

Perhaps the most prominent defence of orthodox moral consequentialism (MC) of recent times is Shelly Kagan's *The Limits of Morality*. In this section, I try to show exactly why this position and Kagan's argument for it are wrong.

Kagan recognises that his 'extremist' position differs from common sense moral thought in two ways: it denies any *constraints* on an agent maximising the good, and it denies the existence of *options* for the agent to do anything other than maximise the good, options to favour himself or those close to him. I argue against Kagan on this latter divergence from ordinary morality.

In his later article, 'Defending Options', Kagan begins with the question, 'What... must be the case for a given act to be morally required?'[1] He argues that there must be some 'morally relevant reason' for performing the act, and that the 'balance of morally relevant reasons' must support doing the act. For Kagan, these two conditions are sufficient for a moral requirement to hold. He considers two suggestions as to a possible third condition on there being a moral requirement: a motivational condition, proposed by Michael Bratman, and an enforcement condition, suggested by Michael Waldron. I outline my own suggestion as to what the third condition of a moral requirement is. I start from the inadequacies of Waldron's enforcement condition, and re-formulate his Mill-influenced suggestion in terms of sanctionability, in particular the sanctions of blame and guilt.

---

[1] Kagan (1994), 334

In the second part of this chapter, I highlight again the distinction between the two different forms of consequentialism outlined earlier, MC and GC. I argue specifically against MC, that is, against traditional consequentialism, firstly by showing that it does not follow from GC. GC provides an answer to such questions as 'What is the best thing I can do?' and 'What do I have most reason to do?' MC, on the other hand, provides an answer to questions like 'What morally *must* I do?' and 'What am I *obliged* to do?' I argue that when we answer the second type of question (whether we are consequentialists or not), and when we (usually retrospectively) justify ourselves to others, we rarely defend our action on the basis that it was the best thing we could do. Rather, we usually point to the fact that our efforts were sufficient. I offer a framework of (explanatory) reasons why we in fact do less than the best, which I divide into 'cognitive' reasons and 'moral' reasons. When we justify our behaviour to others, we generally appeal to a claim that our 'cognitive' and 'moral' efforts were satisfactory, *not* to a claim that what we did was optimal or what we had most reason to do. This is how our talk of morally justifiable or morally acceptable behaviour operates. Just what it is to behave in a morally justifiable way is for one's 'cognitive' and 'moral' efforts to be sufficient.

I then go on to make clearer the relation between a theory of the good in action and an account of moral obligation by focusing on the role of impartiality in our moral theory. Consequentialist morality is generally thought to give pride of place to impartiality. I argue that impartiality is a feature of the consequentialist theory of the good, but that, in the most plausible consequentialist story about moral obligation and wrongness, impartiality plays a quite different role from the one it is usually thought to play. I distinguish shallow partiality (which is allowed even by traditional consequentialism, MC) from deep partiality, a distinction which runs close to that between instrumental and non-instrumental partiality. Complete impartiality can be seen to be a plausible feature of an account of the good in action without the counterintuitive implications of MC, if we take care again to distinguish the two aforementioned types of ethical questions.

I argue that running together these two types of questions, as frequently happens in moral philosophy and in everyday life, results in a tendency towards adopting an implausible theory of what is good or valuable. This erroneous running together of separate questions leads us away from a consequentialist account of value, of what we have most reason to do. This type of error is compounded by the ambiguous use of the term 'irrational'.

## CHAPTER 4: MORALITY AND ETHICS I: ANSCOMBE, STOCKER AND WOLF

My purpose in this chapter, and the next, is to show why consequentialism should explicitly limit itself to being a position specifically about value, choiceworthiness and what we have most reason to do. I point out that it has often been portrayed, by supporters and opponents alike, as being a theory of many other things in ethics, and I argue that each of these portrayals is a mistake. Rather than thinking that consequentialism can give a direct answer to all the major ethical questions, we have to consider first which questions we should be asking. If we modify the questions, I argue, consequentialism in fact becomes a much more plausible ethical theory, giving a direct answer to some of the modified questions, while remaining quite neutral on others.

In pursuing this thesis, I discuss influential work by Elizabeth Anscombe, Michael Stocker, Susan Wolf and Bernard Williams. These writers see themselves as taking a stand against the orthodoxies of twentieth century English-speaking moral philosophy, and in particular against the two most widely held ethical theories, utilitarianism and Kantianism. Their work has provided much of the impetus for the (re-)emergence in recent decades of virtue ethics. However, I want to suggest that they are equally useful for those who wish to hold on to a substantial part of the modern ethical thought which Anscombe *et al* criticise. For my purposes, this involves examining their criticisms as they apply to consequentialism.

Firstly then, I consider Anscombe's thesis that 'The concepts of obligation, and duty- *moral* obligation and *moral* duty, that is to say- and of what is *morally* right and wrong, and of the *moral* sense of 'ought', ought to be jettisoned if this is psychologically possible.'[2] I argue for the softer thesis that we ought to cease treating such concepts as fundamental. There are two sorts of grounds on which we might criticise these moral concepts; firstly, if they are incoherent or self-contradictory, and secondly, if they ground a harmful mindset or ethical practice that we might be better off without.

I then go on to agree with Anscombe, although for partly different reasons, that much of the confusion which results from our central moral concepts arises because they have inherited the legislative framework which structures ethical thought in the Hebrew-Christian tradition. I suggest two features of a crude law

---

[2]Crisp and Slote (1997), 26

conception of ethics [although Western systems of law as they appear today are by no means as crude as the model I criticise] which make it an inappropriate model for our ethical thinking at any fundamental level, especially if we want to be consequentialists. The first of these is something which Anscombe flags: that laws, in the sense of 'laws of the land', are necessarily codified in terms of external action-types, such as stealing, lying, killing and so on. My conclusion here is that neither of the two dominant strands in the concept of wrongness- blameworthiness and choiceworthiness- can be cashed out in terms of external action-types, in the way that laws are. The second feature of a legislative framework for morality that I question is that the dictates of law are generally expressed as 'musts': requirements, demands or obligations. If this 'mustness' comes to dominate all our ethical thought, then it will squeeze out the important features of the ethical realm which cannot take this form, such as ideals or mere recommendations.

I then discuss Michael Stocker's argument that modern ethical theories, such as consequentialism, are 'schizophrenic'; that they hold up as desirable a disharmony 'between one's motives and one's reasons, values, justifications'.[3] Against Stocker, I argue that we should not drag down our values to fit our motives. That is, we should not modify what we take to be valuable, just because we are not continuously (or indeed ever) motivated significantly to act with regard to it- for example, the well-being of those on the other side of the world. It is also not the case that we ought to be motivated by everything we value all of the time. I also suggest that some feeling of fragmentation is inevitable simply in virtue of the empirical state of the world, where so much avoidable suffering exists. I go on to conclude that consequentialism should not be understood as a theory of the appropriate content of our motivation.

In response to further criticisms made by Stocker, and by Susan Wolf, I argue that consequentialism is to be strictly a theory of choiceworthiness or of what we have most reason to do, not an account of the eudaimon life, or the good life for humans. I make a number of remarks to show how such a distinction is possible, and indeed necessary. In response to Wolf, I also distinguish the claims of GC from other, less plausible claims with which it might be confused. I then go on to argue, in the light of my discussion of Wolf's concerns about utilitarianism, that we should not regard 'moral' reasons as *sui generis* reasons, but instead treat all reasons for action as being grounded in the same sorts of considerations- e.g.

---

[3]Ibid., 66

human welfare.

## CHAPTER 5: MORALITY AND ETHICS II: BERNARD WILLIAMS

In the first part of this chapter, I discuss Bernard Williams's article 'The point of view of the universe: Sidgwick and the ambitions of ethics'. My aims are (a) to argue, with Williams, that the ambitions of ethics should, in some sense, be limited, and (b) to argue, against Williams, that ethical theory, in particular my favoured, relatively unambitious form of consequentialism- one that asserts GC, but denies MC- is not incoherent.

I begin by endorsing a worry that Williams has about Sidgwick's ethical ambitions, namely the idea that one can *deduce* a comprehensive set of moral rules from one or more fundamental principles (e.g. the principle of utility). However, after welcoming Williams's rejection of the claim that consequentialists are committed to an especially calculative mindset, I criticise his discussion of the so-called 'cool hour' and what he calls 'deep dispositions', which the consequentialist wishes to endorse in theory but which, according to Williams, are incompatible with 'the utilitarian consciousness'. I argue, against Williams, that the utilitarian or consequentialist (at least when he limits himself to claiming GC) is not committed to any such consciousness. This leads on to Williams's serious concerns about the relationship between ethical theory and practice. I argue that a theory like GC is not, as Williams thinks all ethical theories must be, distinctively a theory *for practice*. Rather, the theory of value (in action), and hence of what we have reason to do, is first and foremost a theory *tout court*, which can unproblematically aim to take up an impartial standpoint, just as other areas of enquiry do. I make some comments which attempt to undermine the disanalogy which Williams thinks holds between ethical theory and science.

Although Bernard Williams has been one of consequentialism's foremost critics in recent decades, I argue that the consequentialist can make his own position far more plausible by paying heed to one of Williams's most central tenets- the distinction between narrow morality, 'The Peculiar Institution', and wider ethics. I develop a distinction between these two standpoints in a way that can make the status of consequentialism clear. I argue, contrary to the majority of consequentialists, that the consequentialist should take up the latter standpoint, rather than the former. While I use Williams's distinction as a starting-point in

drawing the outline of this dual framework, I highlight important differences between my account and his, in order to pinpoint exactly what place consequentialism has in the ethical landscape, as I see it.

## CHAPTER 6: REASONS

Thus far I have suggested that consequentialism should be understood as a theory of two things: (a) the best thing we can do, and (b) what we have most reason to do. In this chapter, I consider views according to which answers to these two questions come apart. In other words, someone might claim that the best thing that I can do (just like the best thing that can happen) is whatever will result in the best consequences, impartially considered, yet deny that what I have most *reason* to do is therefore that which will bring about the best consequences. In particular, I have in mind the sort of position associated with Sidgwick, developed by Roger Crisp, according to which there are not only impartial reasons, but also egoistic, prudential or agent-relative reasons. I argue against such a position, and against Crisp in particular, on two grounds. Firstly, if we postulate these two types of reasons, we seem committed to saying that there is some balance between the two, and hence always a determinate answer as to what we have most overall reason to do. If this is correct however then we will be left giving a very unintuitive account of someone who strikes an apparently 'healthy' balance between the two types of reason, but then goes on to do more good overall (e.g. bring about better consequences, impartially considered) at his own expense. Are we to say that this person is now doing something that he has *less* reason to do? That he is *less* rational?

Secondly, such a quasi-Sidgwickian position seems to undermine a direct explanatory relationship between value and reasons, which I take to be a major advantage of consequentialist theories of reasons to act. I also make some remarks as to precisely what hangs on claiming the term 'reason' for either position. I go on to question whether we have sufficient independent grasp of the question 'What do I have most reason to do?' to settle the issue as to whether it is best answered in the way the impartialist consequentialist proposes, or in the way the rational egoist suggests. I suggest that the two responses are really just answers to two different questions, and that we are misled into running them together by considerations of motivation which are not relevant to the questions in hand.

Finally, I consider a common objection to my position, namely that it seems to imply that many or most of us are acting irrationally almost all of the time. I point up an ambiguity in the notion of irrationality that underpins this worry, and suggest that on either construal, my view draws more plausible conclusions than the alternatives.

## CHAPTER 7: CONSEQUENTIALISM AND MORALITY

In the final chapter, having denied throughout that we should regard consequentialism as a plausible direct theory of moral obligation, I consider exactly what, if anything, the consequentialist *is* committed to saying about moral obligation. It may be argued that I have simply avoided the demandingness objection by failing to give any account of moral demands at all. I argue that we should be cautious in offering a 'theory' of what moral obligations we have, and outline the limited role that consequentialism has in shaping what we take our moral obligations to be. I discuss two positions to which my theory may be thought to bear some resemblance, namely Samuel Scheffler's hybrid theory and Rule Consequentialism, and suggest that, in refraining from offering a theory of our moral obligations, my form of consequentialism avoids the problems which face these two positions. I conclude that consequentialists should not be preoccupied with developing a distinctively consequentialist theory of moral obligations.

## CHAPTER 1: The Positive Case for Consequentialism

Within philosophical discussion of normative ethics in the twentieth century, two systematic theories dominated- consequentialism (often in the form of utilitarianism) and Kantianism. Towards the end of the century, a third family of views emerged centred around the heading of virtue ethics. A fourth alternative position in normative ethics is to deny that the truths of ethics can be brought into a general system at all, that no systematic ethical theory is possible. Although the central aim of this thesis is to defend consequentialism, it will become clear that the form of the theory I develop is one that seeks to accommodate many of the central concerns of those philosophers who are attracted to virtue ethics and those sceptical about the possibility of systematic ethical theory. Nevertheless, the position I aim to defend is fully-fledged consequentialism, which I take to be the view that the value of our actions is determined solely by the value of the consequences they bring about.

The main thrust of this thesis is not to provide positive argument in favour of consequentialism. Rather, my main goal is to make consequentialism much more appealing by showing that many of the most plausible prima facie objections to it are in fact not good objections to the strongest form of consequentialism, i.e. one that denies that we have a moral obligation to bring about the best consequences we can, and instead limits itself to being a theory of choiceworthiness and of what we have most reason to do. This general emphasis on the deflation and deflection of objections to consequentialism, rather than the setting out of positive arguments for it is justified, I think, by general attitudes to consequentialism within moral philosophy. There are few moral philosophers who have not been attracted to consequentialism at some point, and even its most vehement opponents generally regard it as a theory that has a lot going for it, and thus naturally must be argued against before alternative ethical theories are put forward.[1] Samuel Scheffler outlines the huge prima facie plausibility of consequentialism as follows:

---

[1]Philippa Foot, for example, says, 'utilitarianism tends to haunt even those of us who will not believe in it. It is as if we for ever feel that it must be right, although we insist that it is wrong.' Foot (1985), 196. Similarly, D. D. Raphael says 'Most of us find an immediate attraction in utilitarianism. If we are later convinced that it does not account satisfactorily for the concept of justice, we are not quite sure which way to go.' Raphael (1972-3), 87

'One thing that they [consequentialists of different kinds] all share... is a very simple and seductive idea: namely, that so far as morality is concerned, what people ought to do is to minimize evil and maximize good, to try, in other words, to make the world as good a place as possible.

On the face of it, this idea, which lies at the heart of consequentialism, seems hard to resist. For given only the innocent-sounding assumption that good is morally preferable to evil, it seems to embody the principle that we should maximize the desirable and minimize the undesirable, and that principle seems to be one of the main elements of our conception of practical rationality. Anyone who resists consequentialism seems committed to the claim that morality tells us to do less good than we are in a position to do, and to prevent less evil than we are in a position to prevent. And this does not sound nearly as plausible.'[2]

The appeal of consequentialism rests, I believe, on some very widely held intuitions, which I discuss below. Those who reject consequentialism are not generally people who fail to see the attraction of it, but rather are put off by the apparent, often unexpected consequences of it. Again, these consequences fall into two families: (i) consequentialism requires us to act in ways which are unreasonably demanding; and (ii) consequentialism seems to recommend to us types of action we ordinarily think morally forbidden. The main focus of my thesis is to defend consequentialism against the first family of objections.

In the remainder of this first chapter, however, I aim to outline what I take to be good positive reasons for believing consequentialism to be true. These perhaps will only amount to sketches, but I hope that in taking up the negative task, in the rest of the thesis, of defending consequentialism against various objections, a fuller and richer picture of positive reasons for accepting it will emerge. In this chapter, I hope simply to point up some of the features of consequentialism which give it its appeal. The arguments I canvas are intended to be arguments in favour of consequentialism as a theory of the value of our actions, not as a direct theory of moral obligation (a position I argue against in the next two chapters).

---

[2]Scheffler (1988), 1

We can divide positive arguments in favour of consequentialism into two categories- those which favour a consequentialist structure, rather than a non-consequentialist one; and those which are grounded in a particular conception of the good that is to be promoted.

SECTION 1: Structural Arguments

Firstly then, consequentialism may seem to be at an advantage over its competitors simply in virtue of its structure, i.e. prior to fully spelling out what the good consists in. This line of argument takes its starting-point from the following sorts of immediately intuitive thoughts:
(i) How can it ever be better to make the world a worse place?
(ii) How can it ever be bad to make the world a better place?[3]
(iii) How can something be worth doing if it has no good effect on the world?

These sorts of thoughts seem to count in favour of an ethical picture in which the goodness of the consequences of an action are what determines how good an action it is, how much it is worth doing, how much reason there is to perform that action.

There are a variety of ways in which to flesh out this sort of structural argument.[4] Robert Shaver, for example, suggests, 'The deontologist refuses to approve of committing one act, bad by its lights, to prevent the occurrence of more than one of the same bad acts. Consequentialist theories- whether they value only welfare or things in addition to welfare- do not face this paradox, and so they are, in this respect, attractive.'[5]

In contrast to the consequentialist, the deontologist is someone who denies that only consequences determine the value of our actions. He claims that reasons for action are not all ultimately grounded in the value of the consequences they produce, and thus that sometimes the best action is one that is not consequentially optimal. For the deontologist, some act-

[3]For Scheffler, consequentialism incorporates 'the deeply plausible-sounding feature that one may always do what would lead to the best available outcome overall'. Scheffler (1982), 4
[4] For structural arguments besides the one I outline here, see e.g. Pettit (1991), Kupperman (1981)
[5]Shaver (2004), 236

types have a not-to-be-doneness about them that is not grounded in its consequences. Take some such act, perhaps an instance of torture. The consequentialist can agree that it is very bad. But if it is the only means to preventing other bad acts of the same kind, with the same bad consequences, then it is best to perform this 'bad' act.

The deontologist distinctively thinks that this is not always true. The other not-to-be-done aspect of it (i.e. apart from its bad consequences) may outweigh or trump or simply silence the consequential considerations. Philip Pettit distinguishes the deontologist from the consequentialist by characterising the former as someone who believes that some values are not to be promoted, but to be 'honoured'. For example, if someone has a right not to be tortured, my appropriate reaction to this right (the 'value') is not to minimise rights violations, but rather to ensure that I do not myself perform any rights violations.

So, as Shaver says, the deontologist claims that in some circumstances, we are not to directly perform some bad act, even if it is the only means to preventing several other bad acts of the same type. Thus, for example, according to some deontologists, I should allow five people to be killed, rather than carrying out a killing myself. Or it is more choiceworthy to allow five innocent people to be tortured, rather than to torture one innocent person myself.

The next step for the consequentialist is to ask what the rationale is for the deontologist's insistence that we should never commit one wrong in order to prevent five wrongs of the same type and seriousness being committed by someone else. What we can say straight off is that the rationale does not seem to be grounded in any concern for the potential victims of the wrongs. If we had an impartial concern for all the potential victims, surely we would try to minimise the amount of suffering they will jointly undergo. Perhaps then it is out of concern for the potential victim of *my* action? But why should we have greater concern about him than about the other potential victims who will suffer just as badly if I refrain from doing X, the 'bad' action? After all, their suffering is just as real, just as painful.

Of course, my reluctance, the pain it will cause to me to torture someone, for example, is in itself a very significant bad consequence to be taken into account. But this is a consequential consideration, which according to the

consequentialist is just to be weighed up against others. The deontologist is committed to appealing to something else, beyond the mere consequences.

The main objection to the deontologist here is that any rationale for this honouring structure to some values seems to have an unappealing focus upon the agent himself, *at the expense of* the welfare of potential victims. The motivation for the refusal to countenance doing X might naturally seem to be something like, 'I'm not getting involved', a determination to keep one's own hands clean, to avoid a stain on one's own soul or character. It appears to suggest a moral selfishness or squeamishness which seems rather unsavoury in the face of the terrible suffering of the victims, which could be avoided were you to overcome your 'moral' scruples.

A deeper problem is that even if there is some moral cost to the agent who has to perform a bad action-type, even if this is a genuine disvalue to an agent, then still, surely these considerations hardly outweigh or trump the real suffering of real people in a case of torture, for example.

Of course, deontologists do not explicitly offer any such unsavoury rationale or motivation, and are by no means committed to it. But there is a real pressure upon them to offer some alternative story, some *explanation* as to why I should avoid the bad act, even when it results in much better consequences. And it is not easy to see where the room is, between the interests of the victim, and the unappealing concern for moral cleanliness, for such an account. This is what I take to be the problem for the deontologist.

There are further structural arguments for consequentialism, which are grounded simply in the typical problems facing any deontological account. One advantage of consequentialism is its ability (in principle, at least) to offer a comprehensive account of when some actions are better than others. A notorious problem for deontology is its struggle to explain what is to be done in a case where its edicts conflict. So, to take a simplified example, a deontological system may include the rules, 'Do not lie', and 'Do not break a promise'. But what about the situation in which the only way to avoid breaking a promise is by telling a lie?

There are several options open to the deontologist here. Firstly, there may be a hierarchy of principles. More plausibly, there may be an appeal to

consequences, such that one should break the rule whose violation in the particular situation will do less harm. Some deontologists may be happy to take this option, but it does seem to bring a pressure to assert that what ultimately justifies the rules in the first place is their consequential value.

These structural considerations are by no means conclusive against deontology. But the onus certainly seems to be on someone who denies consequentialism to answer what seems to be a problematic question: *Why* is it that some values are to be honoured rather than promoted?

We can also observe that consequentialism seems to have other theoretical advantages over its rivals- namely, simplicity and continuity with the rest of our evaluative thought. Consequentialism does not draw a sharp dividing line between the narrowly moral and the rest of practical reason- all reasons for action of the same type, and have the same source. It postulates no *sui generis* 'moral reasons': the reasons that count in favour of doing something that is morally required are just the same sorts of reasons that count in favour of doing anything else- namely, that the action will produce good consequences. I say more about this in Chapter 6.

## SECTION 2: Conceptions of Goodness

### *The Value of Welfare*

While for most of my thesis I will remain broadly neutral on which conception of the good is correct, I take it that possibly the strongest arguments in favour of a particular form of consequentialism come from a specific conception of the good. So, for example, there are several features of utilitarianism that make it attractive besides its consequentialist structure. While most of what I say later in my thesis will be compatible with many conceptions of the good, and thus favour disparate types of consequentialism, I take it that any plausible account of the good will have human happiness or well-being at its heart. The utilitarian claim, of course, is that the value of our actions is determined only by happiness or well-being (of humans and other sentient creatures, on most accounts).

A two-part argument can be developed along the following lines:
1. Welfare is uncontroversially valuable, and provides reasons for action.

2. All other apparent reason-giving features of actions in fact give reasons in virtue of contribution to welfare.

Robert Shaver argues that what makes utilitarianism so attractive is that 'justifying the value of welfare requires no argument'.[6] And Geoffrey Scarre suggests, 'Even critics who deny that utilitarianism provides an ultimately satisfactory account of morality must concede that it focuses on something of indisputable significance to us, the quality of our lived experiences.'[7]

If we agree that human welfare is valuable and gives us reasons for action, it then of course becomes incumbent on the consequentialist to cash out other apparent sources of reasons for action in terms of their contribution to welfare. So, for example, we often say that the fact that doing X is the only way to keep my promise is a reason for me to do X. The consequentialist must argue that the reason-givingness of this fact is grounded ultimately in the value of welfare (or, naturally, some other good consequence). So, he might argue that keeping promises generally results in good consequences, especially for the promisee. There may be other good consequences besides, for example resulting from supporting, rather than undermining the generally useful practice of promising. If in fact, in unusual circumstances, no one will benefit, either directly or indirectly, by my keeping my promise, it becomes difficult to see what there is still to be said for keeping the promise in any case.

Of course most consequentialists think that telling the truth, keeping promises and respecting rights are hugely valuable. But they are so just in respect of their contribution to welfare or happiness or some other good consequence. Again, we can turn to an intuitive line of thought: how could something really be better if it made no difference to anyone's good or welfare? As Robert Goodin puts it, 'Why... should we ever require gestures that are of no earthly use to anyone, anyway? Yet any moral theory... that refused to put considerations of utility firmly at its centre must necessarily run the risk of requiring just such empty gestures'.[8]

---

[6]Ibid., 235
[7]Scarre (1996), 25
[8]Goodin (1991), 242

*Meta-Ethical Arguments*

Another set of arguments which count in favour of a consequentialism which makes human welfare central might by termed meta-ethical arguments. Scarre outlines the variety of such arguments as follows: 'Not only is the possession of a flourishing life of obviously great importance to people: it also provides a basis for moral theory which is remarkably empirical and down to earth. Utilitarianism offers a single rational criterion for appraising actions, practices and institutions: the maximisation of utility. It makes no reference to the law of a divine being (though many utilitarians may be religious believers), nor does it posit, as some secularised descendants of religious-based morality do, a moral law without a law-giver. It likewise avoids the metaphysical and epistemological problems of theories which claim the existence of rights, duties and norms known through an obscure faculty of moral intuition. Utilitarians countenance no experience-transcendent sources of moral knowledge, and represent moral thinking as continuous with other kinds of practical reasoning.'[9]

Human welfare seems to be an unmysterious part of the natural world. If this is what ultimately grounds our ethical truths, then we are spared the need to postulate *sui generis* moral properties or non-natural objects in order to provide truthmakers for our ethical truths.

Some sort of welfarist picture seems to afford us the objectivity we expect of ethics, while avoiding the metaphysical and epistemological difficulties of other ethical views. Human welfare is not 'queer', to use Mackie's term.[10] We need not postulate any peculiarly moral faculty, such as moral intuition or a moral sense in order to be in touch with the parts of the world relevant to ethical truth. Thus, a view on which welfare is the source of ethical truths fits in with a desirable naturalist picture of the world.

It is often asserted by moral philosophers that normative or evaluative conclusions cannot be drawn from descriptive premises. 'You can't get an *ought* from an *is*,' as the slogan goes. However, one might just think that

---

[9]Scarre (1996), 25
[10]J.L. Mackie (1977)

welfare facts are precisely the sort of descriptive facts which are able to bridge the 'ought-is' gap. Arguably, experientialist accounts of human welfare are particularly well-placed in this respect. Experienced human happiness appears to have built into it a 'to-be-doneness' of the kind Mackie says objective ethical truthmakers would require; likewise, human suffering seems to have to-be-avoidedness built in. As Shaver puts it, 'The goodness or badness of pleasures or pains is evident from the point of view of anyone capable of experience. This explains our agreement on their value.'[11]

One can also plausibly portray the objective value of welfare as a good starting-point in combating ethical scepticism and objectionable forms of relativism. A consequentialism of this sort can ground the claims that some things are better than others, some actions are better than others, and that not all sets of moral rules, social codes or social practices are equally good. What makes some actions or codes better than others is precisely their effect on human welfare. If we look at the second strand of Mackie's moral scepticism, his argument from relativity, consequentialism again seems to fare well compared to other ethical theories. While ethical disagreement may, to a greater or lesser degree, be a fact of life, as Mackie asserts, one area of widespread agreement is in the valuing of human welfare. Consequentialism is perhaps better placed in some respects, then, to construct an argument against Mackie's attack than is, say, a form of deontology which postulates the existence of a number of moral truths, which are supposed to be self-evident or intuitable by a special moral faculty.

These remarks are not intended to be conclusive. Rather, I hope simply in this opening chapter to have sketched some strong prima facie reasons why we should find (a welfare-based) consequentialism attractive.

---

[11]Shaver (2004), 248

## CHAPTER 2: The Demandingness Objection

My starting-point is the demandingness objection to traditional consequentialism. In this chapter, I aim to show that consequentialism, as it is often expressed, is unreasonably demanding. I consider two possible consequentialist responses to the objection, suggested by Michael Slote and Bjorn Eriksson respectively, and in showing why these responses are unsatisfactory, I show why consequentialism should not be understood as a direct theory of moral obligation at all, but instead as a theory of the value of our action, and of what there is most reason for us to do.

### SECTION 1: Consequentialism Re-defined

It is generally thought that consequentialists are committed to following claims:
**Moral Consequentialism (MC)**: We have a moral obligation to bring about the best consequences we can. The right action is the one that brings about the best available consequences. All other available actions are wrong.

Indeed, it might be thought that to be a consequentialist just is to hold MC to be correct.[1] However, I wish to argue that in fact consequentialists should reject MC. I shall attempt to do the following:

1. Explain why being a consequentialist does not entail being committed to MC
2. Argue that we should reject MC
3. Argue that consequentialism should not be regarded first and foremost as a theory of moral obligation

Firstly then, it is clear that while claim MC is recognisably a consequentialist position, so too are the following:
**General Consequentialism (GC)**: The best thing we can do is to bring about the best consequences we can.
**GC'**: What there is most reason for us to do is bring about the best consequences we can.

---

[1] In *The Limits of Morality*, Shelly Kagan, calls consequentialism 'the view which holds (roughly) that agents are morally required to perform the act that will lead to the best results overall.' Kagan (1989), xi. If we assume, reasonably, (and as Kagan seems to allow) that being morally required to do X is equivalent to having a moral obligation to do X, then Kagan stands as a typical example of someone who presumes that being a consequentialist commits us to MC.

When looked at like this, we can see that claim MC is only one of many possible starting-points for the consequentialist. GC does not logically entail MC for example. Indeed, it seems plausible to deny that we have a moral obligation to do the best thing we can do. Equally, GC' does not logically entail MC. It is plausible to deny that we have a moral obligation to do what we have most reason to do. If this makes sense, then someone can hold GC or GC' while rejecting MC, and yet still be correctly described as a consequentialist.

Now all three claims can be thought, firstly, to have in common some claim such as the following:
(X) Consequences can be ranked according to goodness.

Each position then goes on to endorse in one way or another the bringing about of the best consequences we can. I will take it that to qualify as a consequentialist, one must (i) hold (X) and (ii) make some endorsement claim like MC, GC or GC'.

## SECTION 2: The Demandingness Objection

Now accepting MC forces us to countenance some deeply counter-intuitive results. Foremost among these is the common complaint that it makes our moral obligations extremely demanding. It may well entail, for example, given empirical facts about the world, that I have a moral obligation to give away almost all my money and worldly possessions to help the poor and the starving in the Third World, drop my own pursuits and interests, and devote virtually all my time to raising further funds for the same cause. (If this seems to some a questionable candidate for what will in fact bring about the best consequences I can bring about, we can fill in some similarly demanding alternative.) If I do any less than this, then according to MC, I have failed in my moral obligation. As Shelly Kagan puts it, this form of consequentialism (which he defends) 'forbids me to do anything less' than the most I possibly can, and so if I continue to go about my normal way of life according to my 'common sense' moral code, then 'most of my actions are *immoral*'... 'in going to the movies I do what is morally *forbidden*'.[2] Even if I am a regular donor to charity, I work in the community two nights a week and spend most of my time being a good teacher, friend and citizen, then I have failed in my moral obligations.

---

[2]Kagan (1989), 1

Thus, Moral Consequentialism seems to imply that every action is either obligatory or forbidden. This flies in the face of our ordinary conception of the moral realm as having three categories, rather than two. The three categories under which actions can fall according to 'common sense' morality are:
(a) obligatory;
(b) permitted, but not obligatory;[3]
(c) forbidden.[4]

For the defender of MC, category (b), what is permitted, but not obligatory, is empty; no actions fall under it.[5] This is in sharp contrast to the thinking of ordinary morality, according to which not only do *some* actions fall into this category, but in fact *the vast majority* do. Only options such as killing, stealing, lying and so on are forbidden by ordinary morality, and only those such as keeping one's promises, saving a nearby drowning child, supporting one's children and perhaps making some contribution to those less well of than ourselves count as obligatory according to our common sense morality. The vast majority of options are thus left as 'permitted but not obligatory'.

Kagan, for example, describes common sense morality as follows: 'the moderate insists on the existence of options: admirable as it may be to sacrifice one's career, time, possessions, or life for the greater good, such behaviour- he assures us- is surely not *required* by morality. As long as one stays within the bounds set by moral constraints, and makes the typically *minor* sacrifices occasionally required, one does all that morality demands of us.'[6]

---

[3]In his 'Saints and Heroes', J. O. Urmson argues that we can more usefully set out at least four or five categories of action, so as to allow for those which go above and beyond the call of duty (or obligation). Urmson (1958) For our purposes, such acts can be understood to fall in the second category, 'permitted, but not obligatory'. It is, of course, a further problem for views such as MC that they do not allow for such supererogatory acts.

[4]I will take it the notion of 'forbiddenness' can be cashed out directly in terms of obligation; a forbidden action is simply one which it is obligatory to refrain from. Likewise, we can regard an obligatory action as one which it is forbidden to refrain from.

[5]This is not perhaps strictly true in every sense of 'permitted'. If two acts P and Q have exactly equally good consequences, and every other possible act has less good consequences, then it might be said that one is 'permitted, but not obliged' to do P, and 'permitted, but not obliged' to do Q. One is of course on this model obliged (not merely permitted) to do (P or Q). Such cases are rare however, and if the only acts which fall into the category of 'permitted, but not obligatory' are those which are optimal and equally good to some other possible act, then there will be far fewer acts in this category than according to common sense morality.

[6]Kagan (1989), 10. By 'typically minor sacrifices' here, Kagan means the sort of thing demanded of us in a nearby emergency, like the obligation to save a child drowning nearby. However, I think we can extend this and allow that ordinary morality, as it is conceived by most of us, has a standing requirement that we make *some* contribution to those less well-off than ourselves. It is certainly true that we are not castigated and shamed if we give nothing to charity the way we would be if we

We can thus see that the demandingness objection to consequentialism highlights the huge gap that exists between common sense moral thinking and consequentialism as it is usually conceived.

Such objections as this are usually taken to be objections to consequentialism *per se*. I have characterised them simply as objections to MC. However, when we turn to claims GC and GC', we can see that the demandingness objection does not automatically apply. Given the state of the world, GC entails that the best thing I can do is to give away most of my money and devote my time to fund-raising. Similarly, given the state of the world, GC' entails that what I have most reason to do is to give away most of my money and devote my time to fund-raising. Neither of these positions has as yet made any moral *demands*, although of course there may be some pressure upon General Consequentialists to give an account of the demands of morality. I return to this question later on.

## SECTION 3: Are GC and GC' Unintuitive Too?

Many will feel that already GC and GC' are just as counter-intuitive as MC. Can it really be that what I have most reason to do is to give away most of my money and devote all my time to fund-raising?

Imagine that, after seeing TV pictures of harrowing suffering in the Third World, David resolves to help out, and decides to make a few changes in his lifestyle in order to improve the lot of a few of these sufferers. He sets up a direct debit at his bank which sends 20 per cent of his earnings to Oxfam, sells his football season ticket and instead spends each Saturday afternoon collecting money for Save the Children outside a shopping centre. He feels that making these changes in his lifestyle amounts to doing something pretty good, as he will be making a real, substantial difference to many people's lives. All in all, he feels quite pleased with himself and at one with the world.

Then he runs into his friend George, who was similarly affected by the same TV pictures of the Third World and, like David, has decided to do something about it. He has siphoned off 75 per cent of his earnings for Third World aid, committed himself to collecting money, working in the community and so on all

---

mugged an old lady, murdered someone or allowed a child to drown, but when we discuss morality openly, a comparatively well-off person who gives nothing to charity and is comfortable with this is certainly at least frowned upon, and can plausibly be held to have violated a moral obligation.

weekend and every weeknight but one. He has left himself enough money to pay all his bills and taxes, afford himself a rare treat once in a while and buy his round at the pub on his one evening off from do-gooding. Likewise, he has left himself enough leisure time so that he will stay in the same typically sturdy frame of mind and physical fitness to allow him to remain happy, keep up important friendships and stay healthy enough in body and mind to continue his charitable efforts indefinitely.

When David hears about George's behaviour, he is deeply impressed. He may not feel guilty or ashamed that he is not doing as much George, that he will not be making a difference to as many people's welfare as George is. However, his over-riding thought is 'Well, George is a better man than me'. David may be proud of his resolutions and efforts, but he can clearly acknowledge that George has done even better, he has chosen an even better way of life than David himself.

Now is there some sense in which David has picked out a better life *overall* than George? David's life seems better balanced, containing as it does not only significant active concern for those less well off than himself, but also leaves time and space for his own personal pursuits and relationships, setting aside a significant part of his life for his pursuit of philosophy, literature, music, relaxing with his friends and family and so on. He affords his pursuits and relationships this time and space not simply to the minimum degree that will keep him sound enough of mind to continue maximising the good, but rather according to the actual (non-instrumental) value which he attaches to these pursuits. He may appear to treat his personal interests and relationships as valuable in themselves in a way that George does not. David's life seems to be a healthier, more wholesome, indeed more distinctively *human* life than George's life. We might imagine that David's life has a structure and balance that makes it a better candidate to be called *eudaimon* than George's life. David's life seems the more *desirable*. If one asked the parents of a new-born baby what sort of life they most hope for for their child[7], they might well pick out a well-rounded, complete life like the one David has chosen, rather than the extreme, one-dimensional lifestyle adopted by George.

Nevertheless, we should still recognise George's life as even better. It is vital to remember that George's self-sacrificing behaviour brings vitally important

---

[7]See Urmson (1988), 12

elements of the good life to many, many others. If we value basic nourishment and freedom from disease even as equally important to well-being as developing personal projects and relationships, then we have no choice but to acclaim George's lifestyle, which brings about so much good in so many different lives at relatively little cost (in the form of George's single less well-rounded life), as the better life.

In judging which lifestyle is the best, *given the actual state of the world*, we should not be distracted by which way of life we would pick out as best, *given ideal circumstances*. We can outline a way of life that we think ideal for humans, the type of well-rounded, full and flourishing life that we wish everyone could enjoy. It would of course be for the best if everyone could enjoy such a lifestyle where all their needs and wants were satisfied, they had the opportunity to spend their days in projects they care about deeply, enjoy relationships of mutual respect and love, and generally feel at home in the world.

Let us imagine some time in the distant future where poverty, suffering and disease were all but eliminated, and all humans had ready access to food, medicine, education and comprehensive, organised health care. How *then* would I go about maximising the good? It would certainly not be by dashing around the world, desperately looking for trouble and suffering, in the hope of making some small difference to someone. Such behaviour would not bring about the best consequences. Rather, bringing about the best consequences would involve pursuing personal projects I care deeply about, developing my character, talents, virtues and intellect, and forging full, healthy relationships with those around me. This 'ideal-circumstances' lifestyle is no doubt more attractive than George's lifestyle, and may not seem to involve as much moral strain, but it is nevertheless the best one to adopt in these vastly preferable circumstances.

It is important to realise that the life of frantic do-gooding adopted by George is not good in itself, i.e. simply because it is a life of frantic do-gooding. *All else being equal*, we can pick out a healthy, flourishing, well-rounded life as a better life than one of frantic do-gooding. But all else is not equal. Given the state of the actual world today, the latter life plausibly brings about far more good than the former, and is therefore the best life we could live, the lifestyle we have most reason to adopt.[8]

---

[8]It is not my aim in this thesis to make any particular empirical claims about what in fact is the best way to tackle the problems of Third World poverty. It should also be remembered that the best

SECTION 4: Why is Only the Best Good Enough?

Now GC and GC', as I have so far formulated them, make claims only about the *best* thing we can do and what we have *most* reason to do. But of course, a General Consequentialist will not rest simply with a claim endorsing the bringing about of the *best* consequences. Rather, they will have a scale of endorsements according to varying levels of goodness of consequences. Thus if we imagine that sets of consequences can be ranked on a scale from 1 to 10, then each type of consequentialist will give a stronger endorsement to the bringing about of consequences scoring 10 than consequences scoring 5, which will in turn receive a stronger endorsement than consequences scoring 1. Thus such a consequentialist makes a claim not just about the best consequences possible, but also about *better and worse* consequences. The strength of the endorsement in each case is in proportion to the goodness of the consequences. The position I endorse is what Michael Slote calls Scalar Consequentialism.[9]

Now, when we turn to our original formulation of MC, we can see that it is slightly different in form. MC does not have the notion of 'best' built in the way GC and GC' do. As Slote points out in his 'Satisficing Consequentialism', such formulations in fact are the combination of two claims, only the first of which is distinctively consequentialist according to our re-definition of consequentialism in Section 1, and the second of which is an unnecessary but traditional appendage to MC. Slote expresses it as follows:

'Roughly, then, consequentialism standardly involves the claim that the rightness of acts depends on whether their consequences are good enough together with the particular view that only the best possible (in certain circumstances) is good enough. And given this way of partitioning standard consequentialism, it is not perhaps immediately obvious why these two theses should naturally or inevitably go together. Could not someone who held that rightness depended solely on how good an act's consequences were also want to hold that less than the best was sometimes good enough, that an act might qualify as morally right through having good enough consequences, even though better consequences could have been produced in the circumstances?'[10]

---

thing we in the developed world can do collectively in this regard may be quite different from the best thing I can do individually.

[9]Slote (1985)

[10]Slote (1984), 140

Re-formulating this in the same terms as what I call MC, the two parts of the traditional model are roughly (i) the claim that our moral obligations are ultimately dependent on whether the consequences are good enough, and (ii) the claim that only the best consequences are good enough. Slote goes on to reject the equivalent of (ii).[11] He outlines a position he calls 'satisficing consequentialism' which holds that, in order to fulfil our moral obligations, we do not have to bring about the best consequences we can, but rather merely 'good enough' consequences, where suboptimal consequences may be good enough. Slote sees this as a possible solution to the demandingness objection I raised in Section 2.

## SECTION 5: Good Enough *What*?

The motivation for taking up Slote's position is clear- it seems *too strong* to say that we are morally obliged to bring about the best consequences we possibly can. Thus it is understandable that a consequentialist might seek to soften the claim by saying that we are morally obliged merely to bring about *good enough* consequences, rather than the best we possibly can.[12] However, Slote's satisficing consequentialism leads us into all sorts of difficulties.[13] To see this, we can turn to an example which Slote himself uses, but to which he clearly comes up with the wrong answer.

'How do we react to fairy tales in which the hero or heroine, offered a single wish, asks for a pot of gold, for a million (1900) dollars, or, simply, for (enough money to enable) his family and himself to be comfortably well off for the rest of their lives? In each case the person asks for less than he might have asked for, but we are not typically struck by the thought that he was *irrational* to ask for less than he could have'.[14]

Slote makes it clear that there is no cost whatsoever to the hero in wishing for more; it would cost the hero no more in terms of effort or sacrifice. Slote's

---

[11] Although Slote's discussion continues in terms of right and wrong, I will continue to speak of obligation, which on most models amounts to the same thing if we accept that what is morally wrong is the same as what is forbidden, and thus the same as what we are obligated to refrain from doing.

[12] Bringing about even better consequences than those that are merely good enough would of course be permitted and would presumably be praiseworthy, admirable an so on. The point is simply that bringing about this level of good consequences would not be obligatory.

[13] For a fuller account of the problems facing Slote's satisficing Consequentialism, see Mulgan (2001), Chapter 5.

[14] Slote (1984), 147

conclusion is that the hero's action was right, rational, satisfactory; this is because he has brought about consequences that are 'good enough'. This seems obviously the wrong answer (at least from a consequentialist point of view). To see that this is so, let us change the example very slightly. Let us imagine that the hero has been told by the genie or God that he is to come along in a day's time and wish for whatever he wants, and his wish will be granted. So the hero has the opportunity to come back and discuss the matter with us, his friends, family and community; there is no question of the hero being caught on the hop, as it were, and simply blurting out, 'I wish for a pot of gold' or 'I wish for a million dollars'. Assuming there is no cost involved, and no limit to the genie's powers (or God's powers) clearly the rational thing to do (again, at least if we are consequentialists) is to wish that the best possible consequences be brought about- perhaps wishing for everyone for the rest of time to be perfectly happy, for example. If, after discussing this with his community and realising the goodness of the consequences he can bring about, he goes along to God or the genie and says, 'I'll settle for a pot of gold', then his action is straightforwardly irrational, wrong, blameworthy; he had a moral obligation to bring about better consequences than he did.

Now, in this fantasy case, we may well conclude that the hero does indeed have a moral obligation to bring about the best consequences he can. And yet, I wish to claim that this does not hold as a general rule, but rather is peculiar to this example. What is it then about the example which makes it the case that there is an obligation is to bring about the best consequences one can, whereas David in Section 3, who fails to bring about the best consequences he possibly can, has fulfilled his moral obligations?

The difference lies precisely in the stipulation in Slote's fairytale that there is no *cost* or *effort* required on the part of the hero if he is to bring about even better consequences than he does, whereas there is significant cost or effort imposed on David if he is to bring about even better consequences than he does. Even though we can see David is not as admirable or praiseworthy as George, still he seems to have done 'enough' in some sense; it seems harsh and unreasonable to *demand* more. There is a stage at which we say his *efforts* are satisfactory, rather than there being, as Slote claims, a stage at which the goodness of the consequences produced are satisfactory.

So the alternative softening claim to Slote's is: *we are permitted to put in less than the greatest effort we possibly can in bringing about the best consequences.* I

don't intend to try to pin down this notion of moral effort or sacrifice with great precision. Indeed I am not sure that much more precision is desirable. The provisional answer to the question posed at the head of this section is therefore 'good enough *effort*'.

## SECTION 6: Good Enough Effort for What?

This notion of a good enough effort seems to be a major strand in our moral thinking. Even though he has done less than George, we have a strong intuition that David's efforts are more than reasonable, that he has surpassed an acceptable minimum of moral effort. He has *done enough*. But enough for what? In answering this we can ask how we feel about someone who we consider has clearly *not* done enough. Let us imagine a third person, Andrew, owner of a small business, who not only refuses, even when asked, to give any money to the Third World after seeing the harrowing TV pictures, but further lies, cheats, swindles and steals from his own customers in order to promote his own economic good. Now what sentiments do we feel towards Andrew when we hear about his behaviour? Well, for a start, we *blame* him. Furthermore, when we confront him about his behaviour, and emphasise the ill effects he is wreaking on others and point out that he has made no effort whatsoever to do the right thing, we expect him to *blame himself*, to feel *guilty*. Of course, it is not just empirically true that we *do* blame someone such as Andrew, but further we consider him blame*worthy*; it is *appropriate* to blame him.[15] Even though both David and Andrew do less than the most they possibly could to bring about the best consequences, only Andrew is blameworthy. We think it appropriate that Andrew be blamed and that he feel guilty, whereas we think it absurd and unjust if David is blamed or feels guilty.

Now perhaps we can understand better why Moral Consequentialism is so unintuitive. MC claims that *everyone* who does less than the most they possibly can to bring about the best consequences has failed in their moral obligation, has acted morally wrongly. But our use of the term 'moral obligation' implies that those who fail to fulfil their moral obligations are blameworthy, it is appropriate to for us to blame them, and for them to feel guilty. Thus, if MC is right, then it is correct for us to blame David, and correct for David to feel guilty for pursuing the

---

[15]Let us assume that, in Andrew's case, there are no extenuating circumstances, whatever these might be.

lifestyle he does. This conclusion flies in the face of all our common sense moral thinking.

Should we bite the bullet? No, because we can have the intuitive advantages of consequentialism without the deeply counterintuitive conclusions of Moral Consequentialism by instead endorsing General Consequentialism. In other words, we should regard consequentialism not primarily as a theory of moral obligation, or moral wrongness, but rather as first and foremost a wider ethical theory of the value of our actions, which at best (and perhaps in combination with contingent empirical facts) might derivatively give an account of moral obligation and moral wrongness. The General Consequentialist, who we might call the 'ethical consequentialist', does not make the same unintuitive claims about moral obligation, moral wrongness and hence about when it is appropriate to feel blame and guilt.

The ethical consequentialist will of course face demands to offer a rough account of the moral obligations that we do in fact have. I return to this sort of demand in Chapter 7. There seems no reason, however, to believe at first sight that the moral obligations outlined will be anywhere as demanding as those suggested by MC.

So, to recap, what is wrong with MC is that the term 'moral obligation', as we ordinarily use it, cannot be understood independently of the circumstances under which it is appropriate for us to *blame* the agent in question and appropriate for him to feel *guilty*.

The appropriateness of blame and guilt is at least an integral part of our thought about moral obligation, wrongness, forbiddenness, and so on, even if it is not the whole story. It is because of the centrality of blame and guilt to our main moral categories that the suggestion- implied by MC- that David has failed in his moral obligations sticks in the throat. I thus hope to have offered a framework for discussing the nature of our moral obligations which is unequivocally consequentialist, but is not committed to the unintuitive suggestion that we have a moral obligation to bring about the best consequences we can.

SECTION 7: A Combined Account of Wrongness

So the conclusion to be drawn from our responses to Slote's fairytale example and my example of David the imperfect charity donor is that our intuitions about the notion of moral obligation are grounded in the idea that the demand for maximisation of consequences must be tempered by considerations of effort and sacrifice required of the agent. This might seem to point us in the direction of an account of moral obligation, moral rightness and moral wrongness which simply combines considerations of goodness of consequences with something like considerations of effort, sacrifice or difficulty in performing actions. However, it will be my main aim in this section to argue against an account of moral wrongness which combines these two types of considerations.[16]

A useful model for a combined account of the kind I mention is provided by Bjorn Eriksson in his paper 'Utilitarianism for Sinners'.[17] Eriksson attempts to characterise wrongness in terms not just of goodness of consequences, but also in terms of difficulty, which, as suggested above, might seem an attractive option in the light of the demandingness objection and our reaction to Slote's example. Although Eriksson's account is put in terms of difficulty, which is less plausible, I think, than one in terms of effort and sacrifice, I'll begin by considering Eriksson's own account, then make some remarks to show that an account directly combining consequences with considerations of effort or sacrifice would be similarly misconceived.

Eriksson in fact accepts the claim that the morally right action is the one which brings about the best available consequences. However, his account does take on board the considerations that I have discussed above, and begins by pointing out that all wrong actions (i.e. all non-best actions) are not equally wrong. Rather, there are degrees of wrongness, ranging from actions that are extremely wrong to actions that are almost right. Eriksson claims that the person who fails to do the most they possibly can for the poor has done something 'wrong to a reduced degree'.[18] Whether this terminology makes enough of a concession to ordinary

---

[16]While so far in this chapter my focus has been on moral obligation, in this section my argument will be expressed in terms of moral wrongness, since this is the central concept discussed by Eriksson in the article I concentrate on. Again, I take the two concepts to be connected as follows. We are morally obliged to do what is right; acting wrongly is equivalent to violating a moral obligation. Corresponding to Eriksson's recognition of degrees of wrongness, we might speak of degrees of seriousness of violations of obligations.

[17]Eriksson (1997)

[18]Ibid., 214

language is something I discuss below. The central thrust of Eriksson's account however is that it portrays wrongness as a function of two things: 'The *difference in value* produced by the action as compared to its alternatives' and 'the *difficulty of the action* performed as compared to the difficulty of its alternatives'.[19] Eriksson's account is thus of interest to those who sympathise with the considerations I have raised because it takes account of the fact that our use of the term 'wrong' is often bound up in the difficulty or effort involved in performing certain actions.

He goes on to flesh out the grounds of his theory as follows:
'These variables seem natural because they connect to two kinds of consideration that go well with the spirit of utilitarianism: (i) the general importance of the values different available actions produce and (ii) the principle that "ought" implies "can". The relevance of this principle may not be evident, but the intuitive idea is this: Just as you do not act wrongly at all in "refraining" from an action that would have had the best consequences but you could not perform, it seems appealing to hold that if you almost could not do it (it was very difficult) you did not act so very wrongly when refraining from it.' He goes on to claim that 'letting increased difficulty (of the right action) lead to decreased wrongness (of the performed action) amounts to taking the considerations behind the principle that "ought" implies "can" seriously. Difficulty in the sense that Eriksson intends can be understood as 'a continuum ranging from the limit "cannot" to "cannot be avoided"'.[20]

Eriksson offers a rough list of reasons why actions may be difficult:[21]

1) great physical effort
2) mental effort
3) acting with inadequate information or competence
4) behaviour the agent finds shameful, ridiculous or disgusting
5) overcoming fear of harm and pain
6) suffering, pain or frustration of preferences
7) sacrificing money or other resources
8) setting cherished projects and commitments aside
9) acting against one's moral convictions

---

[19]Ibid., 215
[20]Ibid., 215-6
[21]Ibid., 216

He suggests, rightly I think, that there is a 'firm position in our thought and practice of the simple idea that particular actions are more or less difficult to perform'.[22]

So Eriksson's formula for determining the wrongness of an action is the following: 'For any action *a* its degree of wrongness *Wa* is a function of how difficult *a* is compared to its alternatives, *a*'s *relative difficulty*, and how inferior *a* is as regards value production, *a*'s *value difference*.'[23]

More precisely:
'An action, *a*'s, degree of wrongness *Wa* = *Vra* x *Dra*.

*Vra* = Value-difference, that is the value produced by the right action, *r* - the value produced by *a*.

*Dra* = The relative difficulty of *a* as compared to the right action *r*; that is, the difficulty of *r*/ the difficulty of *a*'[24]

Later, this is refined to: "An action *a*'s degree of wrongness *Wa* = *Vma* x *Dma*, where "*m*" stands for the alternative action that maximises *Wa*.'[25]

The precise formulation Eriksson puts forward is irrelevant to my criticisms, which will be of a more general nature. The crucial implications of Eriksson's theory for our purposes are that:
(i) the wrongness of an act varies inversely with goodness of its consequences (relative to alternatives)
and
(ii) the wrongness of an act varies directly with the difficulty in performing the action (relative to alternatives).

In particular, it is (ii) that I will take issue with.

Now, Eriksson's account produces the following results:[26]

---

[22]Ibid., 217
[23]Ibid., 218
[24]Ibid., 220
[25]Ibid., 222
[26]Ibid., 219

| | | AGENTS | |
|---|---|---|---|
| | | A | B |
| | a1 | 100 (v. difficult) | 100 (v. easy) |
| ACTIONS | | | |
| | a2 | 10 (medium) | 10 (medium) |

Since, in each case, a1 will bring about better consequences, it is better to perform a1 than a2. But further, Eriksson concludes, 'The judgement that a2A would be much less wrong than a2B seems attractive'. And there does appear to be something in this thought. If we deny Eriksson's account then we will need to explain this attractiveness. So I will return to this below.

## SECTION 8: Why a Combined Account is Unsatisfactory

In spite of the apparent attractiveness of the conclusion in this example, I will argue against Eriksson's position on two counts. The first is the one I mentioned above, that calling every action but the best one 'wrong' does not seem to fit with our moral talk. Or indeed, it seems just false. This is of course a problem for all traditional consequentialists, i.e. defenders of MC. They are forced to conclude that the man who has 'merely' given away 50% of his earnings and devotes 'only' most of his free time to good causes has acted wrongly. This sort of claim undeniably sticks in throat. The reason for this, again, is that condemning an agent's behaviour as *wrong* seems to imply that the agent is blameworthy. And that certainly does not seem to be true of the agent who has put in such substantial (though less than perfect) efforts in his devotion to charity. It seems to me that rather than describe such behaviour as 'wrong', the consequentialist should acclaim it as admirable, while admitting the fact that it is 'less than the best possible'.

However, Eriksson's account suffers from a more serious problem than this. First, I wish to argue that his formula is unsatisfactory as it stands. Then I wish to go on to argue that even if we could make amendments to his theory to overcome the unpalatable results, such a combined theory of wrongness is on the wrong lines. The main problem can be expressed as follows: surely it is always better to perform an action that is less wrong (i.e. nearer the right action)[27]. But

[27] At least in cases which do not involve false belief. For an account of this qualification, see Chapter 7 Section 3.

difficulty (which is part of Eriksson's wrongness) has no bearing on the goodness or betterness of the action, no bearing on its choiceworthiness (beyond its contribution to consequences, which has already been accounted for).

It seems to be a dominant strand in our concept of wrongness that considerations of wrongness are generally over-riding in determining choiceworthiness. If we jettison this strand, we do not seem to have wrongness anymore. Wrongness seems to be bound up in the notion of choiceworthiness in the sense that actions that are less wrong are more choiceworthy than ones that are more wrong.

So Eriksson's theory has the result that an action A can have better consequences than an alternative B, and yet be more wrong, simply because it is more difficult to perform. And an action becomes less wrong just because it was less difficult to perform. This has some bizarre consequences. The essential problem is that relative difficulty (difficulty of an action compared to its alternatives) does *not* in and of itself contribute to making an action less wrong.[28] Using one of Eriksson's examples of reasons for difficulty bears this out. For example:

1. 'Great physical effort': Killing someone by poisoning him is less difficult than killing him by throwing him out of a window because it involves less physical effort (assume they have the same consequences), but does not make it less *wrong*.

Perhaps Eriksson does not realise the extent of the unintuitive consequences that result. But he is of course perfectly aware that his theory has the consequence that an action which was less difficult for the agent to perform and which has bad consequences will be judged less wrong than an act which has not so bad consequences but was more difficult to perform. However, he bites the bullet, sticking by his theory in the face of this seemingly unattractive implication of his theory. He says, 'One conceivable response... is to hold that it is unreasonable that... [the wrongness of X is less than the wrongness of Y]... since... [X]... results in a significantly worse outcome. Nevertheless, the ability to generate results *of this kind* is appealing: one sought-after feature was, as argued above, sensibility to differences in difficulty. That... is a result of that feature in work. Unacceptable or not, this particular consequence of the theory points to the need for careful consideration of the weighting of value-difference and relative difficulty. For

[28]Any disvalue attached to the difficulty of the action has already been taken into account in evaluating the consequences of the action.

instance, it is possible that relative difficulty should be given less weight than value-difference.'[29] But any such weighting still leaves the fundamental problem in tact.

Eriksson's theory would have to be severely modified. But even if it was, there would still be fundamental problems. This is because it is essential to keep apart those considerations relevant to the goodness of actions from those relevant to blameworthiness. Putting them together does *not* give us a concept which fits neatly on top of our concept of wrongness. Whatever Eriksson has created, it is not wrongness. Eriksson has just invented a new, mocked-up property. Not only is it not wrongness, but further it is not a property worth inventing. There is no purpose served by running together these two central moral properties. Keeping the two properties apart, we still have the tools to explain the apparent 'attractiveness' of Eriksson's conclusion in the example in the chart. Two agents who perform actions with equally good consequences, but who put in different amounts of effort or sacrifice, are generally to be judged differently, one is more blameworthy or praiseworthy than the other. But their external actions have the same goodness (consequential value), and hence the same choiceworthiness, just as if they were both mere events.

Eriksson in fact seems to be troubled by what exactly he has created. He wonders, 'This means that it is not clear what the theory is supposed to do; what the normative point of it is. This question is of some importance because it may be difficult to find reason to accept the theory if it cannot be shown to be of any use or interest.'[30]

He goes on to propose 'possible uses for the theory'. These are that it 'may be able to explain some normative judgements and work as a method of decision in some kinds of cases'. I have already suggested why Eriksson's theory is unsuitable as a method of decision. What about the claim that it explains normative judgements? Eriksson's account of wrongness is a construct built out of the two factors that appear to me to be relevant to the notion of wrongness (at least from a consequentialist standpoint)- i.e. goodness of consequences and effort or sacrifice, which Eriksson roughly picks out with his appeal to 'difficulty'. So it is unsurprising that Eriksson's account might go some way in explaining our normative judgements about 'wrongness'. However, it does not explain our

---

[29]Ibid., 220-1
[30]Ibid., 222

judgements as well as when we keep the two notions separate, and indeed Eriksson's tacking-together of the two considerations results in, as I have suggested, some deeply unintuitive results.

We may conclude then that Eriksson's theory is a non-starter. However, it is very useful in highlighting the strains that are inherent in our concept of moral wrongness. There seem to be two main strands to the concept, namely choiceworthiness and blameworthiness, and yet the features which (from a consequentialist point of view, at least) appear relevant to each of these strands cannot be easily put together in the way that a combined account of wrongness along the lines of Eriksson's proposes.

The categories of moral right and wrong are best suited to the notion of particular obligations within a social code. There is, as the term is usually used, only one 'right' thing to do- any other course of action would be wrong. This tradition is not so comfortable describing most situations and most behaviour; it only applies when special, particular obligations (perfect duties, perhaps) crop up. But consequentialism aims at assessing *all* our behaviour, not just special 'moral' situations, and sensibly says that *all* our behaviour is of more or less value. It purports to give a comprehensive ranking of all alternatives in every particular situation. It is unsurprising that if the consequentialist claim that actions are better or worse depending on their consequences is crammed into the mould of the deontological categories of 'right and wrong', as MC attempts, the results we get are very unintuitive.

It may well prove to be a strength of consequentialism as I conceive it that it limits itself to being a theory of choiceworthiness, rather than a theory of moral wrongness. The categories of 'narrow' morality- right, wrong and obligation- may be better suited to a deontological ethics than a consequentialist one. Indeed, in Chapters 5 and 7, I go on to highlight some specific difficulties for the notion of a narrowly moral theory which raise the question as to whether consequentialists should strive to develop a moral theory at all.

## CHAPTER 3: Moral Requirements

In this chapter, I examine the concept of a moral requirement, or moral obligation, with reference to Shelly Kagan's account in *The Limits of Morality*, in which he argues for a position much like the one I have termed Moral Consequentialism (MC), which states that we have a moral obligation to bring about the best consequences we can. I suggest that where Kagan goes wrong is in having too thin a concept of moral requirement. While Kagan proposes two conditions on there being a moral requirement, I argue for a third 'sanctionability' condition. Finally, I flesh out the difference between questions of moral requirement and questions of choiceworthiness by discussing the relation of each to the notions of self-justification and impartiality.

### SECTION 1: Kagan's Consequentialism

Kagan defends an orthodox moral consequentialism, which can be summed up in the following claim:
**MC'**- We are always morally required to perform the act that will best promote the overall good.[1]

This position, Kagan points out, conflicts with two features of common sense moral thought, namely the existence of moral constraints and the existence of moral options. Common sense moral thought often forbids us to perform an act which maximises good consequences in circumstances where that act violates someone's right, or is an instance of harming, or an instance of killing or torturing, and so on. Such prohibitions on performing certain types of actions in spite of their having optimal consequences constitute what Kagan calls *constraints* on promoting the good.

Common sense morality also recognises that there are moral *options*. If the act which will promote the good does not violate any moral constraint, it does not usually follow that we are therefore morally required to perform that action. In other words, there is no general moral requirement to promote the good even within the limits of our moral constraints.

---

[1] Kagan (1994), 333

The consequentialist's rejection of moral constraints on promoting the good is notoriously controversial, provoking many well-rehearsed philosophical battles between deontologists appealing to our common-sense intuitions, and determined consequentialists trying to strike a balance between biting bullets and undermining contrived thought-experiments. 'But torture is *wrong*, no matter how good the consequences that result!' argue the deontologists. 'Yes, torture is wrong', reply the consequentialists, 'but that's exactly *because* of the terrible consequences that torture invariably has. And rare scenarios where torture is permissible are going to have to be situations where all the agent's alternatives will have similarly horrific results.' The result in this sort of debate is generally an impasse, with neither side able to convince the other.

However, the orthodox consequentialist's rejection of moral options is perhaps significantly more counterintuitive than his rejection of constraints. If we are always morally required to perform the action which best promotes the good, then any action which does less than this is *morally wrong*. Thus, the vast majority, if not all, of our actions are morally wrong, since it is almost always the case that we could be doing something more productive of value at any given time. Again, when I go to the cinema or to the pub, for example, I have done something morally wrong, since I could have instead been helping the needy, raising money for Oxfam or whatever. This use of our moral terms seems, as I argued in Chapter 2, to be completely detached from their normal meanings, if it implies that someone who does lots for his family, community and the distant poor, *but could be doing even more*, is in fact acting morally wrongly.

This implication seems to render consequentialism, in the form it is often expressed, as encapsulated by MC or MC', completely unintuitive, even if it is right to reject the existence of moral constraints in promoting the good. Nevertheless, Kagan, fully aware of these counterintuitive implications, attempts to defend MC', and reject the existence of moral options to behave in a suboptimal way.

## SECTION 2: What is a Moral Requirement?

In 'Defending Options', Kagan argues against two attempts to defend moral options put forward by Jeremy Waldron and Michael Bratman.[2] He says, 'The

[2]Waldron (1994); Bratman (1994)

question... is whether options can be defended in a way that seems plausible and attractive even in the light of the rest of our considered moral views.' He begins with the question, 'What... must be the case for a given act to be morally required?'[3]

Firstly, he outlines two conditions which he endorses as necessary for the existence of a moral requirement:
1. There must be some morally relevant reason for performing the act.
2. The balance of morally relevant reasons supports doing the act: if the original reason is outweighed (from the moral point of view), it simply cannot ground a moral requirement.[4]
(Note that the second condition presumably implies the first.)

For Kagan, these two conditions are sufficient for there to be a moral requirement. This is not, Kagan stresses, simply in virtue of the meanings of the words: 'this is certainly not because I believe that "best consequences" somehow entails "required". Rather, I offer the general requirement to promote the good as a substantive truth about morality.'[5]

He goes on to defend his account by rejecting two different possible conditions put forward as supplementary to (1) and (2): a *motivational* condition, favoured by Bratman, and an *enforcement* condition, as defended by Waldron.

3a. Motivational condition: the reasons for performing the act must be such that they are capable of motivating the agent to do it.

3b. Enforcement condition: it must be appropriate to sanction the agent should she fail to do the required act.[6]

Kagan fully appreciates of course that even if he sees off Bratman and Waldron, it is open to opponents of extremist consequentialism to propose alternative supplementary conditions. But Kagan is right to emphasise that the onus is on his opponents to come up with a plausible condition which is part of a complete coherent moral theory.

---

[3]Kagan (1994), 333-4
[4]Ibid., 334
[5]Ibid., 335
[6]Ibid., 334

I think Kagan convincingly refutes both Bratman's and Waldron's suggestions as they stand. However, I outline a more plausible version of Waldron's enforcement condition below, which is not vulnerable to Kagan's criticisms of Waldron. If such a condition must be fulfilled in order for there to be a moral requirement, and if in fact optimal actions rarely meet this condition, then we will, in line with common sense moral thought, have many moral options, and will not be under a constant and unrelenting moral requirement to maximise the good. Kagan's extremist consequentialism would thus be defeated and consequentialists would have to modify their theory accordingly.

Firstly, however, I want to briefly examine the two conditions which Kagan *does* endorse. For each, we need to ask two questions: 'Is it a genuine condition upon an act's being morally required?' and 'Is the condition met by all actions which will best promote the good *just in respect of being* actions which will best promote the good?'

The first condition is uncontroversially a genuine condition of an act's being morally required. Further, as Kagan points out, the fact that a particular act will bring about the best consequences overall seems to imply that there is some reason to perform it, even if that reason were to be outweighed by other considerations.

The second condition also seems to be a genuine one. It seems very unintuitive to say that we might be morally required to perform an action which the balance of morally relevant reasons does *not* support. I will however make further comments about this later on. More importantly, we need to ask the second question, 'Is this condition met by all actions which will best promote the good just in respect of being actions which will best promote the good?' Well, of course, this is precisely what sceptics about consequentialism deny. They point towards moral *constraints* on maximising the good, for example if the optimal action is an instance of harming, killing, torturing and so on. This is mere repetition of the earlier debate- it does no more to take us beyond the deadlock between consequentialism and common sense morality I mentioned before. However, again, as Kagan emphasises, the defender of common sense morality or deontology is still going to have to refute the existence of a moral requirement to maximise the good, *within* his preferred deontological constraints. In other words, he must explain why there is no moral requirement to perform that action out of all those that are permitted, which will bring about the best consequences. As regards the challenge put forward by Kagan's consequentialist

extremism, the rest of us are all in the same boat, whether we are common sense moralists or deontologists, or whether we subscribe to the consequentialist intuition that the balance of morally relevant reasons always counts in favour of doing the act which best promotes the good.[7]

## SECTION 3: The Third Condition: Waldron and Mill

So the question remains: Why are we not always required to perform that action which the balance of reasons supports? I think the most plausible supplementary condition is going to be something like an enforcement condition. Firstly, I will briefly discuss Waldron's formulation before going on to offer what I think is a more plausible version of this condition.

Waldron's main formulation of his enforcement condition is:
3b: It is permissible to impose sanctions on P for failing to do A[8]

Now this formulation is inspired by Mill's account of wrongness, which Waldron quotes:
'For the truth is that the idea of a penal sanction, which is the essence of law, enters not only into our conception of injustice, but into that of any kind of wrong. We do not call anything wrong, unless we mean to imply that a person ought to be punished in some way or other for doing it; if not by law, by the opinion of his fellow-creatures; if not by opinion, by the reproaches of his own conscience. This seems to be the real turning point between morality and simple expediency. It is part of the notion of Duty, in every one of its forms, that a person may rightfully be compelled to fulfil it.... There are other things, on the contrary, which we wish that other people should do, which we like or admire them for doing, perhaps dislike or despise them for not doing but yet admit that they are not bound to do; it is not a case of moral obligation; we do not blame them, that is, we do not think that they are proper objects of punishment.'[9]

---

[7]I won't consider in this chapter the suggestion that the explanation of why the balance of reasons does not always or usually support doing the act that will best promote the good is that there are special *agent-relative* or *agent-protecting* or *prudential* reasons which often outweigh impartial considerations in determining what we have most reason to do. This position might be seen as a plausible refuge for those who are attracted by consequentialism, but are put off by its apparent demandingness. However, this position in fact has, I think, some pretty unattractive implications for anyone who takes consequentialism seriously, as I argue in detail in Chapter 6.

[8]Waldron (1994), 314

[9]J.S. Mill, *Utilitarianism*, Chapter 5 in *Utilitarianism, On Liberty, Essay on Bentham*, ed. M. Warnock; London: Fontana 1962, 303-4, quoted in Waldron (1994), 313-4

Now this dense passage of Mill's probably actually contains a variety of closely related conditions of moral obligations or requirements. Unfortunately, Waldron arguably settles upon perhaps the least plausible strand in this extract, the suggestion that it is part of the notion of duty or obligation that 'a person may rightfully be compelled to fulfil it'. Thus Waldron suggests that for there to be a moral requirement, it must be appropriate for us to *force* the agent to perform the action. Thus Waldron explicitly goes beyond saying merely that someone who fails to do something morally required merits some sort of disapprobation or punishment, but rather merits whatever degree of punishment would be necessary to *force* the agent to perform the action.

Kagan sums up Waldron's apparent position as follows:
'P is required to do A only if it is permissible to impose enough sanction on P for failing to do A (i.e., enough to motivate P's doing A)'[10]

This, hopefully, is not what Mill intends, and in fact leads Waldron to some absurd conclusions which Kagan recites in 'Defending Options'. The moral options resulting from Waldron's condition are much fewer than common sense would dictate, and the pattern of options generated seems morally arbitrary. But even at a purely intuitive level, the idea that an action cannot be morally required unless it is appropriate for us to *force* the agent to do it seems disconcerting at the very least.

So I think a better, more general term to sum up a plausible Millian condition for there to be a requirement might be a 'sanctionability' condition, rather than an 'enforcement' condition. Kagan himself points out the possibility of holding a more modest condition: 'P is required to do A only if it permissible to impose some sanction on P for failing to do A (i.e. only if it is permissible to sanction P at all)'.[11] But even this, I think, is not what we are getting at when we make judgements about moral requirements. A slightly different kind of sanctionability condition is necessary.

---

[10]Kagan (1994), 340
[11]Ibid., 340

SECTION 4: Do We Really Need More?

In this section, I want to argue that remaining where Kagan wants with just the two conditions is not a serious option. And at the same time, I will suggest that it is to the sanctions of the sentiments that we must look in order to discover what we are doing when we moral judgements about moral requirements or obligations.

We can certainly make sense of the idea of something being morally required even though no-one will or can or ought to *enforce* it. Kagan is right insofar as he is rejecting Waldron's version. However, I doubt that we can understand the suggestion that something might be morally *required* without it being the case that an agent who fails to perform the action would be morally *at fault*. We should note both parts of this claim: the 'morally' part and the 'at fault' part. To be at fault is to be criticisable, to have made a mistake, and also probably for the mistake to have been down to you, in some sense. To be *morally* at fault is to be morally criticisable- criticisable for your moral efforts, and this seems to equate with being morally blameworthy.

Anyone who takes consequentialism seriously can subscribe to GC, and can agree that promoting the good is the best thing we could do, and is what we have most reason to do. But for Kagan's claim, MC, we need to go beyond mere 'bestness'. *Requirement* seems to be a wholly different concept from mere bestness. There is rather a 'mustness' which has to be account for. To say that X is required is not just to say that it is optimal, that it is the best thing we could do. Rather, it is to say that it *must* be done. But what does this mustness amount to? 'Must' or else what? Required on pain of what? The answer is *not* 'Or else you could do better'. Rather, it is 'Or else you haven't done *enough*'.

So Kagan clearly has to say more if he wants to talks about requirements. We need an account of what it is to do enough, morally speaking. And this is exactly what, in the spirit of Mill, I am trying to outline. To do enough, morally speaking, is to do enough to not be blameworthy.

And in fact, towards the end of 'Defending Options', Kagan himself seems to become uncomfortable with just the two conditions for requirement which he has said are sufficient. He says, if he is really pushed, what he would add to the mere notion of there being a balance of reasons is the idea of there being

'normative necessity'.[12] This is what it takes for there to be a requirement. But what does he mean by this? All he says on the subject is 'I cannot give you a definition of this concept- but for all that it seems to me a notion that we appear to have. Why think otherwise?'[13] He goes on to say, 'The thought we have in common sense morality is that when there is no requirement we are "off the hook" normatively- there is no normative necessity in our acting that way.'[14]

This imagery of being 'off the hook' is, I think, a nice way of thinking about what a requirement is. However, when we believe we are not subject to a moral requirement, we take ourselves precisely to be off the hook *morally*. The feeling of being morally off the hook has a quite distinct feel to it. If, for example, after going to great lengths to fulfil a promise I have made, I feel that I have fulfilled my moral requirement, it is a feeling of moral immunity in the sense of immunity to *blame*. If someone were to blame me for what I had done, how I had spent my time, they would be wrong to. If, on the other hand, I have not striven to fulfil my promise, then it seems appropriate when others withdraw their respect or esteem for me. This is just what we *mean* by something's being morally required.

## SECTION 5: An Outline of a Sanctionability Account

Most of this chapter has been arguing negatively, in aiming to point out the limitations of accounts of moral requirement offered by Kagan. Without attempting to outline anything like a comprehensive alternative account, I want now to sketch how a plausible sanctionability condition on moral requirements might take shape. Waldron himself in fact suggests a different variation of the enforcement condition later on, which I think is much closer to what is necessary than his main formulation, and which we might to take as a starting point:

3c: 'It is appropriate for P to suffer unpleasant feelings if she does not do A.'[15]

In particular, Waldron has in mind the appropriateness of feeling guilt as a condition of there being a moral requirement. To this, I think we can safely add the appropriateness of blame by others. So the third condition under

---

[12]Ibid., 345
[13]Ibid., 345
[14]Ibid., 346
[15]Waldron (1994), 322

consideration will be something like:
3d: 'It is appropriate for P to feel guilt if she does not do A, and appropriate for others to feel blame towards P if she does not do A.' This my preferred third condition on there being a moral requirement.

There are some issues to consider which might seem problematic for such a condition on there being a moral requirement. My answers to these worries will hopefully provide a positive outline of a plausible account.

1. The first and most obvious question to ask is:
How are we to determine what the 'appropriate' level of guilt (or blame) is in any given case? Waldron's own suggestion is that 'the unpleasantness of guilt should be proportionate... to the difference in overall utility that her failure to act morally made.'[16] This cannot be right however. My action may have disastrous consequences through little fault of my own, through very bad luck or as a result of non-culpable ignorance. If so, then it is *not* appropriate for me to feel guilt or for others to feel blame. Instead, guilt and blame are surely appropriate when they are roughly proportionate to the agent's contribution; proportionate to the quality of his effort, striving, sacrifice, intentions and so on. It is these which primarily determine the circumstances in which we take an agent to be blameworthy.

2. It might be suggested that a consequentialist is not entitled to such an account. After all, is he not committed to the view that we should blame people just if, and just to the extent that, blaming them will have good consequences? But this accusation rests on a misconception. We can see this if we distinguish merely *feeling* blame from blaming 'out loud', so to speak. The question of whether we should actively *express* our feeling of blame, through punishment, or even just through voicing our disapproval or contempt, is of course something about which the consequentialist will have something to say. It is best for us to punish people and express our blame towards them in those situations where this will have the best consequences. I should also emphasise that for the consequentialist who limits himself to GC, and denies MC, *all* that follows from his consequentialism in this regard is that the best thing we can do is express blame when doing so will have the best consequences. The question of whether an agent is blame*worthy* or not is a quite separate one. That is the question of whether it is appropriate to *feel* blame towards the agent, and this is determined

---

[16]Ibid., 323-4

primarily by the agent's contribution in terms of effort, intentions and sacrifice.

3. Might a similar accusation be made of the consequentialist about when it is best that we *feel* blame? Is the consequentialist not committed to a view about when it is best that we feel blame too? Yes, it would, according to consequentialism, be for the best if we feel blame only when our feeling blame will have optimal consequences. We should note that this is to a large extent outwith our control. The practical question concerns whether we have reason to strive to change our natural feelings of blame. It may be the case that sometimes it would be 'for the best' if we feel blame inappropriately, e.g. feel blame towards someone who has done no wrong. There is no contradiction here- the consequentialist can offer plausible accounts both of when an agent is blameworthy[17] and of when it happens to be for the best that we feel blame towards someone who does not deserve it. Even in such a case where A has done X, the act that he had most reason to do, yet it is still expedient for us to blame him, the consequentialist can answer all the relevant questions consistently:
What does A have most reason to do? X
Is it a good thing that people feel blame towards A for doing X? Yes (by the consequentialist claim that the best event that can happen is that which brings about the best consequences)
Is A blameworthy for doing X? No
Is A morally forbidden to do X? No, since he is not blameworthy for doing X.

4. Finally, I want to highlight that if we are to outline a code of external actions which we take to be moral requirements, this can at best be a rough and ready account. After all, it does not follow from the external description of the action, such as 'lying to the man with the axe', whether an agent will be blameworthy or not. It may be the case that for a particular action, externally described, agent X will *not* be blameworthy because his intentions (e.g.) were good, whereas agent Y *will* be blameworthy because his intentions were bad. So we can at best *pick out* required actions (for the purposes of developing some general code of requirements) by considering which actions agents would *under normal circumstances* be blameworthy for failing to do. I say more about this in later chapters.

---

[17]Our account of when an agent is blameworthy need not (and indeed, in my view, should not) be a distinctively consequentialist account. It may well simply rely on our ordinary intuitions about when people are blameworthy. I return to the question of what we should say about apparent cases of conflict between consequentialism and our ordinary moral intuitions in Chapter 7 Section 5.

SECTION 6: Self-Justification

In the final two sections of this chapter, I want to try to flesh out the difference between questions of choiceworthiness on the one hand, and questions of moral requirement on the other, specifically by characterising the difference in terms of the notions of self-justification and impartiality.

It is often said that the fundamental question of ethics is 'What should I do?' or 'What ought I to do?' There are several senses in which we might understand such a question. For example:

1. What should I do, given my aims X, Y and Z?
2. What should I do in order to maximise my own interests?
3. What is the best thing I could do?
4. What is there most reason for me to do?
5. What morally ought I to do?

Consequentialism is often understood as attempting to provide an answer to the last of these questions. I am arguing that instead it provides an answer to questions 3 and 4.

Recall the different sorts of consequentialist claims I outlined in the first chapter:

**Moral Consequentialism (MC)**: We have a moral obligation to bring about the best consequences we can. The right action is the one that brings about the best consequences. All other possible actions are wrong.

**General Consequentialism (GC)**: The best thing we can do is bring about the best consequences we can.

**GC'**: What there is most reason for us to do is bring about the best consequences we can.

GC and GC' answer questions 3 and 4 respectively, while MC, traditional consequentialism, attempts (implausibly, I have argued) to answer question 5. Again, MC is subject to the demandingness objection, since someone who brings about very good consequences, puts in lots of effort, makes many sacrifices, gives away half his earnings to good causes and gives up much of his free time to helping the needy, but does less than the best he possibly can, has, according to

this position, acted wrongly; he has failed in his obligations.

My project is to argue for a position which endorses GC, in other words a thorough-going consequentialist position, but which denies the deeply unintuitive claim about moral obligation. My aim in this section is to outline a general framework of how I see the relation between a theory of the good in action (i.e. that which responds to Questions 3 and 4) and 'morality' more narrowly understood (i.e. that which is the subject matter of Question 5).

So, GC says that the best thing that I could do is whatever brings about the best consequences I can. What does this amount to? Simply that the consequentialist judges the goodness of actions the same way he judges the goodness of mere events- both are a matter solely of the consequences. Likewise, the worst thing we could do is that which would bring about the worst consequences, and there will be a scale of 'betterness' between best and worst.

Now, given empirical circumstances, the best thing I can do may well be to give up most of my time and money to the poor of the Third World, since this will bring about the best consequences. This course of action, this way of life is, uncontroversially, extremely demanding, requiring constant effort and sacrifice. Nevertheless, the consequentialist *is* committed to the claim that the best thing one can do is bring about the best consequences possible, and hence the best possible life, given empirical circumstances, is indeed such a life of effort and sacrifice.

The following response might now seem appropriate: 'Yes, I see that this is the *best* thing that I could do, but what *should* I do? Thus, when we ask 'What should I do?', we do not always intend to ask 'What is the best thing I could do?' Sometimes we want to know 'What do I *have* to do?' Is there not some level of 'acceptableness' which constitutes a further 'should' question? 'What should I do?' meaning 'What *must* I do? What level must I reach?' What particular actions must I perform or refrain from? Our ethical talk often seems to focus as much on Question 5 as on Questions 3 and 4.

Even if we were to agree that the consequentialist has offered a correct account of the fundamental question of ethics in claiming that the *best* thing we could do is to bring about the best consequences we can, *we want to know more.* Our moral talk involves not just better and worse types of action, but also a set of requirements or demands on agents, as well as assessment of a person's worth at

some purportedly deep level. Our moral talk employs centrally this notion of 'doing *enough*', behaving *acceptably*, and a notion of being able or unable to *justify* ourselves to others. All these key notions *go beyond saying something simply about actions being better or worse.* They also say something about the worth of the agent, about how the agent should be assessed, not just how the goodness of his actions should be assessed.

Take then our practice of justifying ourselves to others. It is crucial to realise that when we do this, we do not generally think of ourselves as trying to show that we have done the best thing available to us. Rather, what we try to show that we did enough, that we put in enough effort, took enough care. This, I suggest, is the focus of the difference between Questions 3 and 4 on the one hand, and question 5 on the other.

Consider explanatory accounts of why we in fact act suboptimally, why we often perform some action which was not the best action available to us. We can divide these explanations into two broad categories, which I will call *cognitive* and *moral* respectively. Under cognitive explanations, we might include such things as incomplete knowledge about what the best thing to do is, incomplete knowledge about what the best means is to doing the best thing, or even a false theory of value about what makes an action the most choiceworthy one. Under moral explanations, we can include such things as the fact that performing the optimal act requires great self-sacrifice or great effort, or it requires us to sacrifice the welfare of those close to us.

Importantly, we typically subdivide both cognitive and moral explanations into subcategories as follows:
i. Justifiable cognitive reasons
ii. Unjustifiable cognitive reasons
iii. Justifiable moral reasons
iv. Unjustifiable moral reasons

So a justifiable, or defensible, cognitive reason why I do not act optimally might be that, through no fault of my own, I do not know what is the best means to a good end. The best action may be to give a large donation to charity C, but if I, through no fault of my own, have never heard of charity C, then I will not perform the optimal action. This we can call a justifiable or defensible cognitive explanation.

On the other hand, there are cognitive deficiencies which explain our acting suboptimally which are not defensible. For example, when operating heavy machinery, I may perform an action which endangers someone's life, because I do not know what is the correct button to press when something goes wrong. I may be ignorant of the necessary fact through negligence or laziness in failing to read the manual for the machinery, as I should have done. This is an example where my failure to perform the action has a cognitive explanation, but which is not justifiable or defensible.

Likewise, some moral explanations are justifiable, while others are not. It may be that the optimal action for me to perform is one which would require huge sacrifice to me or a member of my family. For example, I might have to put my child at risk in order to save another child. Ordinary morality allows us to give some priority to the welfare of our own children ahead of that of others.

But not all 'moral' deficiencies are defensible. The explanation of my failure to help the distant poor might be that I simply want a third car, and if I write this cheque to Oxfam, I will not be able to afford it. Or, to use a well-worn philosophical example, I see a child drowning, but decide not to help because that would involve getting my new suit wet. These are cases where the explanation of my failure to act optimally is 'moral' in these terms, but not defensible.

There may also be cases that we are unsure whether to describe as cognitive or moral, involving, for instance, self-deception about facts resulting from a subconscious bias towards our own interests.

When we justify ourselves to others, we are not generally saying 'I did the best thing I could do'. Rather, we usually justify our actions on grounds of imperfect knowledge or of demandingness in terms of the cost, effort or sacrifice required. I am not suggesting, of course, that the common sense account of the best thing we can do is usually a consequentialist one. Even when it is, it would be strange to explain ourselves at such an 'ultimate' level every time we justify ourselves to others. It would be bizarre to give a full, detailed consequentialist account when asked to defend our actions, even if we were consequentialists.

So when we give a justification of our behaviour, we are not always (or perhaps even usually) showing we did the best thing we could have done. We are often not even saying we did what we believed was the best thing we could have done.

SECTION 7: Impartiality

I now want to go on to try and make clearer this view of the relation between a theory of the good in action and an account of moral obligation (or self-justification) by focusing on the type of role impartiality plays in our moral theory. Morality is often thought to be impartial, in some important sense. In particular, consequentialist morality is generally interpreted as giving a central place to impartiality. I want to argue that impartiality is a feature of the consequentialist theory of the good, but that, in the most plausible account of moral rightness, moral wrongness, moral obligation and moral forbiddenness, impartiality does not play the role it is ordinarily thought to do. Rather, it plays a quite different role.

Firstly then, it should be noted that, from a consequentialist point of view, we must make a distinction between what I will call shallow partiality on the one hand and deep partiality on the other. Common sense morality allows for both forms of partiality, while MC, traditional consequentialism, allows only for shallow partiality.

*Shallow Partiality*

Being to some extent partial in our behaviour will have the best consequences, so a consequentialist morality obviously allows for partiality in this sense. For example, if, in deciding what to do, I determine to be scrupulously impartial with every decision I make, the best consequences will not result. If I never favour myself, starving myself until I am as malnourished as the most malnourished person in the world, I will fail to bring about as much good as I can, since I will be too sick to do anything useful. Likewise, if I never favour my family, friends, wife, neighbours and so on above complete strangers, I will be unable to form the healthy relationships necessary for me to retain psychological stability. After losing psychological stability, I will no doubt fail miserably to bring about as much good as I could if I was sound of mind. (Although, we should be careful not to over-estimate the extent to which we need to favour ourselves and our neighbours in order to be effective utility-maximisers.) So some shallow partiality does bring about the best consequences and so is obviously sanctioned by the consequentialist. But shallow partiality is the *only* partiality permitted by orthodox moral consequentialism.

*Deep Partiality*

Deep partiality, of the kind recognised and permitted by common sense morality, is that which does *not* bring about the best consequences. Deep partiality is permitted just to that extent to which we are permitted leeway (on pain of having acted morally wrongly, being deserving of blame and guilt) in not putting in the most effort or sacrifice we can in bringing about the best consequences we can. We have an obligation only to put in so much moral effort or sacrifice. Our permission (in the face of blame, guilt) to do less than the most we can, or sacrifice less than the most we can just is the permission to be partial to ourselves and those related to us. Never favouring ourselves is too difficult. We are blameworthy only for favouring ourselves too much or too often; we are *not* blameworthy every single time we favour ourselves at all. Demanding complete impartiality all the time is *too much to ask*. This is how I propose we should understand the approximate relation between partiality and moral obligation. We are not morally obliged to be perfectly impartial; rather, we are required to be *sufficiently* impartial. The level of permissible (deep) partiality is the level of permissible relaxation from moral effort/sacrifice. *Morality, far from demanding complete impartiality, allows us a certain degree of partiality.* There is nothing in consequentialism *qua* theory of the good in action, i.e. in GC, which contradicts this intuitive picture of how moral requirements operate.

For the consequentialist, deep partiality has *no* legitimate role in determining the best thing we could do, what there is most reason to do, but, I suggest, retains a significant role in determining our moral obligations. The moral obligations we do have are not just determined by *the goodness of our conforming* to the obligatory actions and abstaining from the forbidden actions, but are also determined by *the difficulty (in terms of effort and sacrifice) of our conforming* to them. So, again very roughly, the level of permissible deep partiality allowed by morality corresponds to whatever gap in effort or sacrifice exists between the effort required to do the optimal action and the effort demanded or required, on pain of blameworthiness and guiltworthiness.

Apart from shallow partiality, which is not really partiality at a fundamental level as it is simply adopted as a means to an impartially determined end, an account of what there is most reason to do, what it is best to do, has no room for impartiality. It only allows partiality where this will bring about the most well-being for all people, impartially considered.

Hopefully, once we see the clear gulf that exists between Questions 3 and 4 on the

one hand, and Question 5 on the other, we will see that consequentialism is far more plausible as a theory of the good, and that complete impartiality *can* be a feature of our account of the good in action, without having the counterintuitive implications facing traditional consequentialism which gives the same sort of answer to both sets of questions.

Now it may well be that not all failure to do the very best we can (even setting aside what I called cognitive explanations) is a matter of being partial towards ourselves, but there is certainly a large correspondence.

Just to make explicit the relation between my discussion of self-justification and my account of partiality, what I am suggesting is that justifying ourselves to others involves first and foremost showing that we have not been excessively partial towards ourselves and those close to us. What I am trying to suggest is that impartiality *can* be a feature of our theory of the good, as consequentialism requires. Giving an account of the best thing we could do is *not* a matter of striking a balance between partiality and impartiality. But giving an account of moral obligation *is* a matter of striking such a balance.

The more contentious claim is the first one, the claim that our theory of the best thing we could do, or what we have most reason to do, is not a matter of striking a balance. What we have most reason to do is to be judged impartially across all people who could be affected by our actions. I think many of us are reluctant to bite the bullet here, even though we accept the claim that no one person's welfare is more important than anyone else's. This reluctance emerges, I think from the tendency to run questions like 3 and 4 together with Question 5.

There is certainly a drive, when we consider our own behaviour, to say 'I am leading the rational life, the best possible life available to me, or at least something close to it'. This sort of claim is pretty implausible if consequentialism is true. But I think we are often inclined to deny that devoting our time and resources to the needy is the best thing we could do, rather than admitting that we are acting in a way that is far from perfect. And we think like this because of the mistaken thought that we act wrongly, and are thus blameworthy and ought to feel guilty, if we fail to do the best thing available to us, if we fail to do the 'right' action. Consequentialism in particular suffers from this drive because of its traditional claim that the 'right' action is the one that brings about the best consequences. And so it naturally follows that all other actions are wrong. But we struggle to accept that we are acting 'wrongly' all the time, so we deny that the

right action, the thing we have most reason to do is to live a life of relentless moral effort and sacrifice.

But of course we are not acting wrongly all the time. We are not blameworthy for doing less than the best we possibly could. We are blameworthy only for being *excessively* partial.

In Chapter 6, I argue that this error in our moral thinking is mirrored by a similar error in the way the word 'irrationality' is used, to suggest that doing anything less than what we have most reason to do is 'irrational', with its connotations of moral or cognitive faultiness. How are we to describe someone who knows (or thinks he knows) what he has most reason to do, yet fails to do it? Should we automatically dismiss him as irrational? No, we should not, I argue later, because the cost in terms of effort or sacrifice may be very large, and so this is why he does not do what he knows he has most reason to do. (This is a common phenomenon with regard to prudence as well as altruism, in the form of everyday weakness of will.)

My diagnosis is that because we like to think of ourselves as acting perfectly rationally, we drag down our account of the most rational, best way of life to something closer to our own way of life. We have a tendency to move from the thought 'Surely I can't be required to do that much?' to the thought 'That can't be what there is most reason to do', and so this drive messes up our theory of what is good, of what the best thing to do is. We end up by trying to fit deep partiality into an account of the best possible life. But on my account, it is a mistake to deny the claim that the better the consequences of your actions are, impartially considered, the better your action is overall.

The mistake in our moral thinking that I am highlighting results in unconvincing hybrid theories with arbitrary balancing acts between self-interest and altruism, or between partiality and impartiality, as to what constitutes the best possible life. (I return to this in Chapter 6). What such accounts are really often doing is outlining a life where we are self-sacrificing or impartial to a degree sufficient for blame and guilt to be inappropriate. In other words, they are answering Question 5, rather than Questions 3 and 4, as they claim to do. We should be wary of an uncomfortable, arbitrary hybrid account of 'what I have reason to do' which draws an arbitrary line somewhere between complete self-interest and complete impartiality, and says 'this is what you have reason to do'.

In place of this sort of account, I hope to have outlined a framework where we can have an impartial account if the good in action, and of what we have most reason to do, just as consequentialism asserts, while at the same time giving an account of the shape of moral obligation and self-justification which leaves room for partiality, as our common-sense moral thought demands.

## CHAPTER 4: Morality and Ethics I: Anscombe, Stocker and Wolf

In this chapter and the next, I discuss some influential work in moral philosophy from the second half of the twentieth century, by Elizabeth Anscombe, Michael Stocker, Susan Wolf and Bernard Williams. These writers see themselves as taking a stand against the orthodoxies of English-speaking moral philosophy, and in particular against the two most widely held ethical theories of the last century, utilitarianism and Kantianism.

While the work I focus on here has provided the impetus for what is known as virtue ethics, I want to suggest that these writers' thoughts are equally useful for those of us who wish to hold on to a substantial part of the modern ethical thought that Anscombe, Stocker, Wolf and Williams criticise.[1] For my purposes, this will involve examining their criticisms as they apply to utilitarianism, or consequentialism, and considering which of them the consequentialist must take seriously.[2] In particular, I hope to show that these criticisms can be used to highlight the *limited* role which consequentialism should play in our wider ethical thought. Consequentialism should not be seen as providing a straightforward answer to all of the major moral and ethical questions we tend to ask. Rather than thinking that consequentialism can give a direct answer to all the major ethical questions, we have to consider first which questions we should be asking. These philosophers query whether we are asking the correct questions in the first place. If we modify the questions, I will argue, consequentialism in fact becomes a much more plausible ethical theory, giving a direct answer to some of the modified questions, while remaining quite neutral on the subject of others. In particular, I aim to show that consequentialism, understood simply as a theory of the value of our actions, and of what we have reason to do, that is, GC and GC', rather than as a direct theory of moral obligation along the lines of MC, is not vulnerable to the central objections characteristic of this late twentieth century movement in ethics. Indeed, the limited consequentialism which I recommend is in fact well-placed to take on board the most justified concerns which motivate Anscombe, Stocker, Wolf and Williams.

---

[1] For another influential article in a similar vein to those I discuss here, see Cottingham (1983). Many of Cottingham's worries about consequentialism can be tackled in the same way I discuss in this chapter.

[2] In this chapter and the next, I use consequentialism and utilitarianism interchangeably. This is because the philosophers I discuss usually frame their objections as applying to utilitarianism, but I intend my discussion to apply to any impartialist consequentialist theory. However, I take it that human well-being will be a major constituent of the good, in any plausible form of consequentialism.

SECTION 1: Standing Outside of Morality

First then, I will consider some themes from Anscombe's article, 'Modern Moral Philosophy', which was first published in 1958. Anscombe has three main theses:
(1) It is not profitable for us at present to do moral philosophy; that should be laid aside at any rate until we have an adequate philosophy of psychology.
(2) The concepts of obligation, and duty- *moral* obligation and *moral* duty, that is to say- and of what is *morally* right and wrong, and of the *moral* sense of 'ought', ought to be jettisoned if this is psychologically possible.
(3) The differences between the well-known English writers on moral philosophy from Sidgwick to the present day are of little importance.[3]

All three claims are extremely bold, and I think we would struggle to find many moral philosophers today willing to defend either the first or the third. The second claim however, is one that I think we can take much more seriously. Now at first glance it might seem odd for the consequentialist to even countenance questioning the centrality of our prevailing moral concepts- right, wrong, obligation and duty. After all, consequentialism is usually taken simply to be identical to MC, the view that we are morally obliged to bring about the best consequences we can. Nevertheless, I have suggested that MC is not the most plausible form of consequentialism, thus leaving room for the consequentialist to consider both the theoretical and practical possibilities of 'jettisoning' the concepts of moral obligation, right and wrong, and the moral ought.

Anscombe herself cites Aristotle's ethical writings as evidence that ethics can operate quite coherently without these central concepts, which she understands to be central peculiarly within *modern* moral theories. She argues, 'Anyone who has read Aristotle's *Ethics* and has also read modern moral philosophy must have been struck by the great contrasts between them. The concepts which are prominent among the moderns seem to be lacking, or at any rate buried or far in the background, in Aristotle.' And then, 'If someone professes to be expounding Aristotle and talks in a modern fashion about 'moral' such-and-such, he must be very imperceptive if he does not constantly feel like someone whose jaws have somehow got out of alignment: the teeth don't quite come together in a proper bite.'[4]

---

[3] Anscombe (1958), in Crisp and Slote (eds.) (1997), 26
[4] Ibid., 26-7

I will not presume to argue one way or another as to how plausible an interpretation of Aristotle this is. Aristotle may seem to operate with notions somewhat akin to moral obligation. But it is certainly true that such concepts play a significantly less central role in Aristotle than they do in most modern ethical thinkers. Firstly, we can note that ethics itself is secondary to politics in Aristotle's writings, a framework which is alien to modern ethics, which reverses the order of importance. More important however, for our purposes, is that, even though Aristotle and modern ethicists share a starting point 'How should I live?' or 'What ought I to do?', this first question of ethics admits of a variety of construals. For Aristotle, the first question is more like 'What is the good life for man?', as opposed to the modern philosophers' question 'What *must* I do?'

An ethics perhaps does not necessarily have to treat as central and fundamental our modern moral concepts of 'obligation', 'duty', 'right', 'wrong' and 'blame'. We can still be doing ethics without treating these notions as the fundamental ones. Anscombe proposes that we replace these concepts with thick ethical concepts like ''untruthful', 'unchaste', 'unjust''.[5] Later on, I will suggest what alternative concepts the consequentialist should adopt as fundamental.

For the moment, I want to emphasise that I am fairly sceptical about the practical possibility of jettisoning the modern moral concepts from our everyday moral practices, primarily just because they are so entrenched in our psychology, language and conventions. Perhaps in the long term, concepts which no longer have a coherent or acceptable meaning will drop out of use in our moral discourse, much as the concept of sin has done, or at least is doing. It may also be the case that this disappearance of certain moral terms would be a good thing. However, that is not to say that we should strive to eliminate all use of the moral terminology and conceptual apparatus entrenched in our moral discourse and practice. Nevertheless, I firmly agree with Anscombe, firstly that it is a conceptual possibility that we can do ethics without using these concepts at all, and secondly, that we should cease to regard them fundamental in our ethical theorising.

Now at this point it might be asked, 'Does it really make sense to talk of standing outside of morality and criticising it?' Is morality not the ultimate standpoint of practical evaluation? So I should emphasise, as should be clear from the first three chapters, that I am specifically *not* rejecting *all* ethical evaluation. But there

---

[5]Ibid., 34

are at least two respects in which we might criticise our moral concepts. First, we can show that the concepts themselves are incoherent, and argue that the central moral categories should be replaced by other, coherent ones. And secondly, the central moral concepts may ground a practice which fares badly when viewed from some other, more fundamental evaluative ethical standpoint.[6] This second type of criticism is what Nietzsche has in mind when he 'rejects morality', speaking of a 'revaluation of all values' which seems to presuppose some other, preferred standpoint from which we perform the revaluation.

Now it is quite open to a consequentialist, and indeed quite natural, I think, to propose that the appropriate standpoint from which we should judge our moral practices should be a consequentialist one. In other words, we can criticise a practice if it is harmful, and we can criticise the dominance of certain concepts within an ethical outlook if such dominance in our ethical thinking leads to bad consequences. A practice is valuable on the other hand if it brings about good consequences, relative to the consequences of alternative ethical practices. In a similar vein, the consequentialist is able to evaluate *parts* of our moral practice, endorsing some and rejecting others. Indeed, presumably, any plausible alternative ethical practices would have to emerge from our current practices to some extent.

I will outline my criticisms of 'modern moral concepts' in both of these respects: firstly, modern moral concepts are to some degree incoherent or at least contain significant tensions, and secondly, they ground a harmful mindset that we might be better off without.

## SECTION 2: The Law Conception of Ethics

In order to formulate the first type of criticism, I want to start from some of Anscombe's remarks on the structure of modern moral concepts. Her diagnosis of the problems of modern morality is that they have inherited a particular legislative framework, a 'law conception of ethics' which structures ethical thought in line with the Hebrew-Christian tradition, but that, outwith this religious context, this structure is incoherent. She says, 'The terms 'should' or

---

[6]Bernard Williams, for example, in *Ethics and the Limits of Philosophy*, suggests, that 'The impartial standpoint can be called upon for a different purpose, not to argue someone all the way from bare practical reason to the concerns of justice or benevolence, but to support or demand some ethical conceptions rather than others.' Williams (1985), 71

'ought' or 'needs' relate to good and bad', but 'have now acquired a special so-called 'moral' sense- i.e. a sense in which they imply some absolute verdict (like one of guilty/ not guilty on a man)' - that is, in a quasi-legal sense. In fact, she goes on to say that 'The ordinary (and quite indispensable) terms 'should', 'needs', 'ought', 'must'- acquired this special sense by being equated in the relevant contexts with 'is obliged', or 'is bound', or 'is required to', in the sense which one can be obliged or bound by law, or something can be required by law.'[7] According to Anscombe, 'In consequence of the dominance of Christianity for many centuries, the concepts of being bound, permitted or excused became deeply embedded in our language and thought.'[8] She concludes that the concept 'morally ought' has 'no reasonable sense outside a law conception of ethics'.[9]

I want to go on to suggest two features of this particular Old Testament law conception of ethics which make it an inappropriate model for our ethical thinking at any fundamental level, especially if we want to be consequentialists. But first we should note a couple of things about the particular ethical framework that Anscombe, and I, have in mind. First, Anscombe's phrase, a 'law conception of ethics' is perhaps a little misleading, since not all law conceptions of ethics are vulnerable to the criticisms that she asserts. Likewise, the worries that I have about a 'law conception of ethics' apply primarily to this tradition, and not, as I will remark below, to the law itself, as it is manifest today. So it is important to appreciate that the conception that I will discuss, which Anscombe criticises is specifically that conception rooted in the Hebrew- Christian, Old Testament tradition; a conception of law modelled on divine command. Our everyday ethical concepts have *this* particular history, and still retain the colourings and connotations of religious notions whose metaphysical underpinnings have long since been rejected.

The first of the features I wish to highlight is that laws in this tradition are codified in terms of what I will call external action-types. What I have in mind is what Anscombe is getting at in the following passage:
'The prohibition of certain things simply in virtue of their description as such-and-such identifiable kinds of action, regardless of any further consequences, is certainly not the whole of the Hebrew-Christian ethic; but it is a noteworthy feature of it'.[10]

---

[7] Anscombe (1958), 30
[8] Ibid., 30
[9] Ibid., 33
[10] Ibid., 35

Laws on this model are framed in terms of 'such-and-such identifiable kinds of action', such as 'stealing', 'lying', 'killing' and so on. Of course, we can make these commands more and more specific, for example, framing a law like 'Don't lie unless the person you are speaking to is a madman intent on murder...' We can only go so far in making such specifications explicit, and can by no means frame a code of rules or laws which would accord with *all* our best judgements. Perhaps a certain level of specificness, combined with the application of common-sense, allows us to perform legal judgements which are sufficient for the purposes of law.[11]

Several problems arise however when we transfer to morality this framework of dictates expressed in terms of external action-types. I have suggested that there are two major strands in the notion of moral wrongness:
(1) To say that X is morally wrong and that Y is not morally wrong is, in most contexts, to say that Y is more choiceworthy than X.
(2) An agent is (in general) blameworthy if he acts wrongly.

The problem then is that neither choiceworthiness nor blameworthiness can be picked out via external action-types. The case of blameworthiness is pretty obvious, since whether an agent is blameworthy or not clearly depends not just on what type of action he has performed, descriptively identified, but also, and perhaps primarily, on what is *internal* to the agent himself. Whether we judge an agent to be blameworthy or not depends on his motives, his intentions, on his moral striving or effort, and the level of sacrifice he puts in. These features are clearly not to be identified *simply* by pointing at the fact that the agent lied or stole or killed.

Equally, the notion of choiceworthiness should not be cashed out in terms of external action-types, I think. For example, we can think of some of the worst external action types, identified in descriptive (non-evaluative) terms, such as torture for example. Surely 'Torturing children is wrong' is a prime candidate for a moral judgement. Nevertheless, we can conceive that there could be some circumstances in which torturing someone was the choiceworthy thing to do. To make this point crudely, we can imagine the sort of extreme example often put forward by consequentialists. A madman is standing over the nuclear button,

[11]The law itself (today) is in fact, with its categories of justification and excuse, far more sophisticated than the crude model I criticise here. In many cases, the judgements of the law are specifically *not* made simply with regard to certain external action-types.

and demands that you torture a child or he will destroy the world. By hypothesis, let's say, you have good reason to believe he will carry out his threat if you do not do the torturing, he will not just go on and press the nuclear button in any case whether you torture the child or not, and that there is no alternative means of stopping him. In this case, we want to say that torturing a child is the choiceworthy thing to do. Even the most horrific of external action-types might be choiceworthy if valuable enough consequences are at stake. (We do not need to be consequentialists to admit this point; we just need to acknowledge that consequences are always one morally relevant factor, even if they are not the only relevant factor.)

My conclusion here then is that neither of the two dominant strands of wrongness- blameworthiness and choiceworthiness- can be cashed out in terms of external action-types, in the way that laws on the divine command model are.

But what alternative framework would be better for expressing fundamental ethical claims than this historical legislative model which still, as Anscombe claims, shapes modern moral theory? What sorts of ethical recommendations am I allowing as coherent?
(1) Recommendations still cashed out in terms of external action-types, but which have the status of general rules, and so are not always over-riding in terms of choiceworthiness;
(2) Purely formal recommendations (grounded in a substantive account of choiceworthiness, such as that the most choiceworthy action is that which will bring about the best consequences), which will have limited effectiveness in guiding our actions.
And there is also scope within my account for:
(3) Ethical judgements concerning blameworthiness, which is to be cashed out, at least in part, in terms of what is internal to the agent, i.e. his efforts, striving, sacrifice and so on.

There is a second feature of the divine law framework which makes it an unsuitable model for a consequentialist ethical theory. This is that the dictates of such laws are expressed as 'musts', requirements, demands or obligations, and never as mere recommendations or ideals. This aspect is what Anscombe is flagging when she speaks of the terms 'bound, permitted or excused'. The 'recommendations' of Old Testament law are commands in the sense that they ought to be obeyed on pain of punishment. The parallel in our moral practice is that moral commands are to be obeyed on pain of blame and guilt, or

blameworthiness and guiltworthiness. But of course, actions, lifestyles and behaviour can be valuable without it being the case that we merit blame and guilt for not acting in the recommended ways. Insofar as moral discourse ties itself to the crude legal framework Anscombe criticises then, it struggles to make claims about mere choiceworthiness without implying that we should feel guilty for not doing the choiceworthy thing. But consequentialism precisely should, first and foremost, be a theory of choiceworthiness, and so it is often interpreted as making implausible moral claims which seem to imply blameworthiness where none exists. When consequentialism is crammed into the legislative model, we get very unintuitive answers. For example, someone like David in Chapter 2 who brings about lots of good in his community, in the Third World, makes great sacrifices and so on, yet *does less than the most he possibly can*, is, according to MC, acting morally wrongly. Again, this usage seems totally detached from common-sense meanings.

This notion of 'mustness' (on pain of guilt and blame) starts to spread across all our ethical thinking. This perhaps links up to some extent with what Bernard Williams calls the dominance of moral obligation. What I am gesturing towards, in any case, is a psychological claim that we start to see all ethical recommendations as obligations, violation of which merits guilt and blame. Within the divine command model, we are subject to sanctions if we fail to obey the recommendation. But consequentialist ethics should not take on this feature, since its recommendations are ideals; they concern what it is best to do, what we have most (objective) reason to do. We are not automatically subject to sanctions just because we have done less than the best.[12]

There may still seem something strange about treating the maximisation of value as an ethical ideal. I will consider such worries presently in the light of some comments by Michael Stocker and Susan Wolf. For the moment, my conclusion from the conflict of the two features of an Old Testament law framework that I have mentioned with central consequentialist tenets is that consequentialism is not well understood as a theory of moral rightness and wrongness, and of moral obligation.

---

[12]Similarly, this second feature of the legal model also, when combined with a consequentialist theory of value in action, notoriously results in a theory which fails to account for the phenomenon of supererogatory action.

SECTION 3: Consequentialism and Motivation

Michael Stocker begins his famous paper, 'The Schizophrenia of Modern Ethical Theories' with the following claim:
'Modern ethical theories, with perhaps a few honorable exceptions, deal only with reasons, with values, with what justifies. They fail to examine motives and the motivational structures and constraints of ethical life. They not only fail to do this, they fail as ethical theories by not doing this'.[13]

Again, consequentialism or utilitarianism is one of the major targets to come under the heading of 'modern ethical theories'. In this section, I wish to argue that consequentialism *should* in fact deal primarily with reasons, values, with what justifies- this is precisely what consequentialism is a theory of. Stocker's criticisms arise from a determination, perhaps by some consequentialists or perhaps by their critics, to regard consequentialism as an ethical theory of everything, as giving a direct answer to every interesting ethical question, i.e. by failing to recognise the limits of consequentialism. I will suggest that consequentialism does *not* fail as an ethical theory by not focusing on motives and motivational structures, and further that, contra Stocker, it does not fail by advocating a way of life significantly lacking in certainly valuable features.

Stocker's worries are put most succinctly in the following passages:
'One mark of a good life is a harmony between one's motives and one's reasons, values, justifications. Not to be moved by what one values- what one believes good, nice, right, beautiful and so on- bespeaks a malady of the spirit. Not to value what moves one also bespeaks a malady of the spirit. Such a malady, or such maladies, can properly be called moral schizophrenia- for they are a split between one's motives and one's reasons [and values and justifications].'[14]

'These theories are, thus, doubly defective. As ethical theories, they fail by making it impossible for a person to achieve the good in an integrated way. As theories of the mind, of reasons and motives, of human life and activity, they fail, not only by putting us in a position that is psychologically uncomfortable, difficult, or even untenable, but also by making us and our lives essentially fragmented and incoherent.'[15]

---

[13]Stocker (1976), in Crisp and Slote (eds.) (1997), 66
[14]Ibid., 66
[15]Ibid., 68

There are several points to be noted in response to this kind of criticism. The first is that we should not drag our values down to fit our motives- i.e. we should not modify what we take to be valuable, just because we are not continuously (or indeed ever) motivated significantly to act with regard to it- for example, the well-being of those on the other side of the world. It is also not the case that we ought to be motivated by everything we value all the time. A disparity between what we take to be valuable, and what we are strongly motivated to act upon may indeed be very unfortunate. But Stocker is surely wrong in laying the fault at the door of our theory of value, rather than our levels of motivation.

Consequentialism should not thus be understood as a theory of the content of our motivation, i.e. as saying that we should be motivated by the thought of maximal happiness or well-being or some such value every time we act. Rather, the most the consequentialist has to say about motivation is that those patterns of motivation which are good ones are those which tend to bring about good consequences, and those which bring about bad consequences are bad patterns of motivation. But what is certain is that we should not modify a plausible theory of what is valuable in order to bring it closer to what we are naturally motivated to do.

Nevertheless, there are a couple of respects in which modern moral thinking can lead to a type of schizophrenia, fragmentation or psychological discomfort. Firstly, any attempt to squeeze a consequentialist account of value into the aforementioned legislative framework of morality, where we are guilty just for doing less than the best possible, does indeed result in psychological unsettling. We are left with the thought that we ought always to be doing something more valuable, and if we regard this ought as moral, i.e. as implying blameworthiness, or guiltworthiness for those who fail to comply, then we can never rest easy pursuing any sub-optimal project or spending time with those close to us beyond the level necessary to keep us sound enough of mind to continue maximising the good. This particular form of moral schizophrenia, then, attaches to MC, but not to GC, and so if an objection built on it is a good one, then it is an objection only to the former strain of consequentialism.

Secondly, and more importantly, I think some feeling of fragmentation is inevitable simply in virtue of the empirical state of the world, where so much avoidable suffering exists. There is thus a huge gap between how we are naturally motivated to live our lives, and what it seems best that we do, simply because the

world is in such a bad state. We cannot, to borrow a Hegelian sentiment, feel truly at home in the world, for this reason. It is unsurprising and indeed quite appropriate that we should feel this sort of fragmentation.

What is clear is that we should not cease to regard the well-being of those on the other side of the world as valuable simply because we struggle to be motivated to actively pursue their concerns at every turn. It is probably the case that we should in fact strive to modify our pattern of motivation to some degree to bring it in line with the things we consider valuable; in other words, we ought to make the plight of the distant starving, sick and so on more vivid to ourselves so that we are better motivated to act for their sake. At the same time, I want to deny that it is appropriate for us to feel guilty simply for doing less than the most we possibly can, simply because of the massive amount of sacrifice that would be involved in doing the most we possibly can, perhaps giving up virtually all our personal projects, closest relationships and so on.

But in this case, it is as much the world that puts us in such a psychologically uncomfortable position as our ethical thought. It is not so much a malady of spirit that exists, as a malady of the world. Perhaps Stocker has a fair point that our *moral* theories, what we treat as morally obligatory, should not stretch us too far from what we are naturally motivated to do, what we naturally care about. But, as I have been stressing thus far, consequentialism is not best understood as a theory of moral obligation in any case.

## SECTION 4: Consequentialism and the Good Life

Stocker also claims of course that 'People who let [e.g. consequentialist values] comprise their motives will... have a life seriously lacking in what is valuable.'[16] And Susan Wolf makes similar points in her article 'Moral Saints'. It is here that I wish to point out another limit on consequentialism. In setting up as an ethical ideal a life which maximises the good impartially considered, consequentialism is *not* engaged in the task of outlining the good life for humans, or the eudaimon life.

For the consequentialist, the best life I could lead, or the life I have most reason to lead, is the life which will in fact bring about the best consequences I can bring

[16]Ibid., 68

about. Of course I want to suggest that such a life is not *required* by any plausible form of consequentialism, because of the costs, effort and sacrifice involved. However, we might think, like Stocker and Wolf, that not only is such a life of copious do-gooding not *required*, but also it is not *desirable*. This life, it seems, will be severely deficient in such central human goods as healthy family life, flourishing friendships and other relationships in which time can be willingly invested, and all sorts of personal projects, interests, hobbies and so on, the pursuit of which will not maximise the good.

However, the consequentialist should be able to happily acknowledge all of these things so long as he limits his claims simply to being about the best life we could lead, the lifestyle we have most reason to lead. The consequentialist claims, given plausible empirical assumptions, that the best life we could lead is one which is extremely demanding in terms of cost, effort and sacrifice. What he should emphasise, as I suggested in Chapter 2, is that, in judging which lifestyle is best, *given the actual state of the world*, we should not be distracted by which way of life we would pick out as best, given *ideal* circumstances. We can outline a way of life we think ideal for humans, the type of well-rounded, full and flourishing life that we wish everyone could enjoy. It would of course be for the best if everyone could enjoy such a lifestyle where all their needs and wants were satisfied, they had the opportunity to spend their days in projects they care deeply about, enjoy relationships of mutual love and respect, and generally feel at home in the world. *All else being equal*, we can pick out a healthy, flourishing, well-rounded life as a better life than one of frantic do-gooding. But, again, all else is not equal, and so if consequentialism is correct then plausibly, the best life I could lead, the lifestyle I have most reason to adopt is a very demanding one.

One further quotation from Stocker might enlighten what I have tried to say in this section.
'One partial defence of these ethical theories would be that they are not intended to supply what can serve as both reasons and motives; that they are intended only to supply indices of goodness and rightness... What is far from clear, however, is whether these theories were advanced only as partial theories, or whether it was believed by their proponents that duty and so on were really the whole, or at least the only important part, of ethics.'[17]

What I have argued here is firstly that consequentialism should not be

---

[17]Ibid., 75-6

understood as a theory of duty or obligation at all, and secondly that it should not be regarded as 'the whole, or at least the only important part, of ethics'. Perhaps the mood of much utilitarian and consequentialist writing in the twentieth century indeed conceived the theory as a direct answer to all of the important questions of ethics. I hope to have outlined an ethical framework in which the claims of consequentialism are more limited, and as a result, more plausible.

## SECTION 5: Moral Saints

I now wish to turn to an influential article by Susan Wolf called 'Moral Saints'.[18] Like the work of Anscombe, Stocker and Williams, Wolf's article has helped to undermine confidence in the utilitarian and Kantian moralities which dominated moral philosophical discussion for most of the 20th century. Wolf's article focuses specifically on the human *ideals* endorsed by common sense morality, by utilitarianism and by Kantianism. She argues that, while these ideals may be plausible *moral* ideals, they are so only from the moral perspective, and that when we see the moral as just one aspect of humanity, moral ideals do not form healthy, rational or desirable *human* ideals.

Wolf is not so concerned with attacking the moral content of the moral theories in question, but instead with criticising their structure and their relation to accounts of non-moral goods and with questioning exactly what a moral theory is supposed to do. These projects are ones with which I have a lot of sympathy. My project, after all, is to try to defend consequentialism by advocating a different ethical structure or framework in which we should understand consequentialist claims. Specifically, I argue that consequentialism should neither be a direct theory of moral obligation, nor should it attempt to be an ethical theory of everything. Instead, it should limit itself to being a theory of value, of the choiceworthiness of action or behaviour and of what we have most reason to do.

I have argued above that the maximisation of good consequences should not be seen as a requirement, that the consequentialist should defend GC and reject MC. Wolf would agree with my rejection of MC, saying explicitly that we are not required to be morally optimal. Her central criticisms in 'Moral Saints', however, seem to apply to GC as much as MC, in that she wants to deny that the utilitarian ideal agent, the person who maximises good consequences, is even a desirable or

---

[18]Wolf (1982) in Crisp and Slote (eds.) (1997)

rational human *ideal* for us to aspire to.

While I sympathise with many of Wolf's conclusions, there are important ways in which I think the ethical framework I defend comes apart from the one Wolf seems to have in mind. I will try to argue that on each count my favoured form of consequentialism fits better with our intuitions, and can better accommodate the central modifications to our moral thought that Wolf rightly wishes to make. I wish to suggest, against Wolf, that the utilitarian (or consequentialist) ideal is the ultimate ideal in an important sense, that it is the *rational* ideal, and not just one ideal amongst others which are supposedly equally rational or desirable. I do not intend here to add any further positive argument to those I laid out in Chapter 1 as to why the consequentialist ideal is the rational ideal, why GC is true, why what we have most reason to do is to bring about the best consequences we can. Rather, what I do in this chapter is to remove some obstacles which might be thought to stand in the way of accepting GC.

Again, I take it that the consequentialist claim has some intuitive plausibility. The fact that an action will bring about some good consequence seems to count in favour of performing that action. In general, it seems plausible to say that our reasons for action vary with the value at stake. It is these sorts of thoughts which motivate the consequentialist. I take it that the problem for the consequentialist is not that his view does not begin with any prima facie intuitive plausibility, but rather that certain apparent implications of his view will somehow make it untenable.

Wolf begins, 'I don't know whether there are any moral saints. But if there are, I am glad that neither I nor those about whom I care most are among them'... 'By *moral saint* I mean a person whose every action is as good as possible, a person, that is, who is as morally worthy as can be.'[19]

First, I wish to suggest that already there are two different notions in play here, even within this brief characterisation. On the one hand, there is the person who acts or behaves optimally from the moral point of view, i.e. the person whose every action is the best action available to him at any given time. But I think this is not necessarily the same as someone who has perfect moral *worth*. To see this, we can imagine someone who strives honestly and continuously to do the right thing, puts in great moral effort and makes great sacrifices, but who, through

[19]Ibid., 79

exceptional bad luck, often ends up doing more harm than good. Such a person seems to have the same moral worth, i.e. is equally morally praiseworthy, as someone who puts in the same effort, sacrifice and striving, and has the same sorts of motivations but, because he is luckier, because the world is kinder to him, brings about far more good.

Wolf does not distinguish these two different interpretations sharply in the article, but for the consequentialist or utilitarian, I think, such a distinction must be crucial. The utilitarian will focus on the first interpretation of the ideal, not the person who has the highest level of moral commitment or is especially morally worthy as such, but instead the person who in fact brings about the best consequences. It is an open question after all whether having an obsessive commitment to 'morality', as Wolf's moral saint seems to have, will in fact maximise good consequences. I am a bit sceptical that it would. Consequentialism is first and foremost a theory of the value of action, and only derivatively, if at all, a theory of an agent's moral worth, since an agent's moral worth seems to be a function not so much of the good outcomes he brings about, but rather a function of his motivations, efforts and striving.

Insofar as we wish to criticise consequentialism on the grounds that it sets up as an ideal an unattractive moral saint, we must take care not to characterise the moral saint as someone with a certain set of concerns, commitments and motivations, but rather simply as someone who performs the optimal actions from the moral point of view. Further, the consequentialist (at least, the defender of GC) recommends whatever pattern of commitment and motivation in fact happens to bring about the best consequences. So it would not necessarily be distinctive of the optimal agent that he thinks about being moral very much.

Wolf does nevertheless intend her criticisms to apply to utilitarianism, in particular any utilitarianism which sets up the utilitarian moral saint as the human or rational *ideal*. Although Wolf concentrates as much on common sense and Kantian conceptions of moral saints, I will focus on those strands in her argument which apply to utilitarianism or consequentialism.

I will begin by picking out what I take to be her main criticisms:
1. SINGLE VALUE: 'the basic problem with any of the models of moral sainthood we have been considering is that they are dominated by a single, all-important

value under which all other possible values must be subsumed'[20]

2. THE SELF: 'The way in which morality, unlike other possible goals, is apt to dominate is particularly disturbing, for it seems to require either the lack or the denial of the existence of an identifiable, personal self'[21]

3. CONDITIONAL COMMITMENT: Closely related to this is the suggestion that the ideal utilitarian agent can at best be 'conditionally committed' to many of the worthwhile things in life.

4. IMPOVERISHED WAY OF LIFE: The utilitarian moral saint will be a character lacking many things that we take to be good. Wolf says, 'the main thrust of the arguments of this paper has been leading to the conclusion that, when such ideals [of moral sainthood] are present, they are not ideals to which it is particularly reasonable or healthy or desirable for human beings to aspire.'[22] And further, 'if the moral saint is devoting all his time to feeding the hungry or healing the sick or raising money for Oxfam, then necessarily he is not reading Victorian novels, playing the oboe, or improving his backhand. Although no one of the interests or tastes in the category containing these latter activities could be claimed to be a necessary element in a life well lived, a life in which *none* of these possible aspects of character is developed may seem to be a life strangely barren.'[23] (81)

In the remainder of this chapter, I hope to show why these worries are not serious objections to GC, and make it plausible that, in spite of Wolf's worries, the consequentialist ideal can be the ultimate or rational ideal. Firstly however, I want to outline some of Wolf's conclusions that I broadly endorse. For example, Wolf asserts, 'Despite my claim that all-consuming moral saintliness is not a particularly healthy and desirable ideal, it seems perverse to insist that, were moral saints to exist, they would not, in their way, be remarkably noble and admirable figures. Despite my conviction that it is rational and as good for a person to take Katharine Hepburn or Jane Austen as her role model instead of Mother Theresa, it would be absurd to deny that Mother Theresa is a morally

---

[20]Ibid., 91

[21]Ibid., 84. This criticism seems close to the arguments from integrity put forward by Bernard Williams. See e.g. Williams (1973) For a consequentialist's defence against Williams's concerns, see Ashford (2000)

[22]Wolf (1982), 92

[23]Ibid., 91

better person.'

Wolf's focus is on the relationship between the moral and other human goods. Her targets are the moral ideals yielded by common sense morality, utilitarianism and Kantianism, which purport to be overall human or rational ideals. I will say more shortly about my doubts about Wolf's analysis here. Wolf suggests three possible responses to the unattractiveness of the moral ideals yielded by the three prominent moral theories:[24]
(1) To offer alternative theories with more palatable ideals
(2) To understand the theories somehow as not yielding ideals at all
(3) To admit that moral ideals do not, and need not, make the best personal ideals.

Wolf endorses the third option. I will partially endorse it, and partly reject it in the comments to come. Wolf rejects a view which considers it 'a test of an adequate moral theory that perfect obedience to its laws and maximal devotion to its interests be something we can wholeheartedly strive for in ourselves and wish for in those around us.'[25] She concludes that she is not against utilitarianism or Kantianism as such- not against the 'intramoral content', but against the way in which the theories are usually framed in their relation to other non-moral human values. Her conclusion then is that 'we must change our conception of what is involved in affirming a moral theory.'

Finally, I want to highlight a quotation from Wolf which nicely articulates the way in which we should understand a position which rejects an excessive moralism of the kind exemplified in MC, as I wish to do.
'In pointing out the regrettable features and the necessary absence of some desirable features in a moral saint, I have not meant to condemn the moral saint or the person who aspires to become one. Rather, I have meant to insist that the ideal of moral sainthood should not be held as a standard against which any other ideal must be judged or justified... It is misleading perhaps to insist that one is *permitted* to live a life in which the goals, relationships, activities, and interests that one pursues are not maximally morally good. For our lives are not so comprehensively subject to the requirement that we apply for permission, and our non-moral reasons for the goals we set ourselves are not excuses, but may rather be positive, good reasons which do not exist *despite* any reasons that

---

[24]Ibid., 94
[25]Ibid., 94

might threaten to outweigh them.'[26]

This seems correct, and important. Nothing in GC implies that this is not all true. We can shed a utilitarian or consequentialist theory of value of any moralistic connotation which suggests that maximisation is somehow the norm, what is expected, and that showing that our behaviour is morally justified, morally permissible, must consist in showing that it was optimal, or providing some *excuse* why it wasn't. The spectre of moral obligation need not dominate our everyday justifications for every act we perform. Rather, we might regard moral judgements and motivations as kicking in (a) when we transgress in some way, and (b) perhaps in some 'cool hour' of contemplation, to use Williams's phrase.

## SECTION 6: Psychological Limitations

Before my main defence of consequentialism, I want to make some brief comments by way of dismissing what might seem an easy escape for the consequentialist from Wolf's criticisms. This might run as follows. The moral saint as described by Wolf, always on the lookout for a chance to promote the good, shedding all personal relationships and personal projects, constantly striving for moral perfection, will not be able to sustain such a way of life, so good consequences will not result, and so this cannot be the way of life that the consequentialist actually recommends. This is certainly true, I think, but is not enough to fend off Wolf, who is quite sensitive to the psychological limitations on our moral efforts. If the obsessive moral outlook is taken to such a degree that it does lots of harm, then of course it will not be recommended by the consequentialist. Whatever life consequentialism recommends as ideal, it will include an outlook which does not undermine one psychologically and socially, so that one is in no psychologically fit state to go on promoting the good.

Nevertheless, I think consequentialism uncontroversially recommends as ideal, to most people at least, given a world at all like ours- where there is so much needless suffering, where we in the developed world can make a significant difference to the welfare of those in the developing world, and even those deprived in our own societies- a life which is seriously self-sacrificial. For most people of course, perhaps for all, a life in which we sever *all* connections with

---

[26]Ibid., 95

our friends and loved ones, and completely discard all our interests and projects, will be a life which brings about more harm than good because it is psychologically unsustainable. This point should be accompanied with the warning that we can be enormously self-deceptive about the degree to which we need to continue doing the things we enjoy in order to retain psychological stability. The relevant challenge for the consequentialist however is this: the consequentialist ideal is, given the state of the world, an extremely demanding, and quite plausibly perhaps also a seriously undesirable one.

I want to try to disarm Wolf's worries insofar as they apply to consequentialism, firstly by making some distinctions which might make consequentialism more plausible, Wolf's concerns notwithstanding, then by making some general comments about value.

## SECTION 7: Six Distinctions

In this section, I want, by means of six distinctions, to defuse some of the apparent criticisms of consequentialism suggested by Wolf's discussion. GC and GC′ are claims about the value of action, and about what we have reason to do. I wish to distinguish these type of claim from six other types of claim with which they might be confused.

### 7.1- *MORALITY AS A COMPREHENSIVE GUIDE TO CONDUCT*

Firstly then, Wolf argues that morality itself should not serve as 'a comprehensive guide to conduct'. This seems correct to me. And consequentialism, at least in the form of GC, should not be regarded as a comprehensive guide to conduct. Rather it is a theory of the value of our conduct. This may of course then guide us to some extent. But we should certainly not consider the defender of GC to be committed to the implausible idea that we ought to consult some consequentialist principle every time we act. It does not seem desirable that we lead our lives according to a 'comprehensive guide to conduct' at all. Perhaps it is for the best if we have our reflective moments, our 'cool hour', where we assess our actions and lifestyle in the light of what we take to be of value. But we should not conceive ourselves, every time we act, as consulting something called 'morality' to see what we should do, or if what we would do is permitted. Doing so would surely not bring about the best consequences.

7.2- *THEORY OF MORAL OBLIGATION*

The second distinction is the one I have already made throughout, between GC and MC; between a theory of the good in action, and a theory of moral obligation. Wolf herself is aware that an ethical theory need not be a theory of moral obligation. In setting up moral ideals (as GC might be regarded as doing), she says 'it seems wrong to insist that everyone try to achieve these things or to blame someone who fails or refuses to conform'.[27] Thus everything we take to be of value in human behaviour and human motivations should not be characterised as something that (a) every single person ought to strive for all the time, and (b) failure to do incurs blame. Thus in giving a theory of value along the lines of GC, we need not be giving a theory with any peculiarly moral colouring in this sense. I say a bit more in the last section about how important it is to see a theory of value shed of connotations peculiar to morality in this narrower sense.

7.3- *IDEALS*

Thirdly, I want to disambiguate the notion of an ideal. We might understand the phrase 'the utilitarian ideal' as meaning simply the optimal agent, i.e. the one who in fact brings about the best consequences he can. Distinct from this, however, is the concept of an ideal as something specifically held up as something to consciously aim at. To what degree we should encourage people to aim specifically and consciously to be someone who maximises utility is, from the utilitarian's own point of view, an open question. Perhaps we should promote, encourage or set up as ideals in *this* sense more concrete and even more obviously attractive ways of life than that of the 'utility-maximiser'. The conscious ideal of being a utility-maximiser might not draw too many recruits. Perhaps at the practical every day level, we ought to focus on more solid aims than this. Likewise, it is an open question as to what it is best that we regard as obligations. If we regard as obligatory moral standards that are well beyond our reach, we will become dispirited and guilt-ridden, and bad consequences will result.

Wolf, again, is aware of these issues. She says, 'More pragmatic considerations also suggest that, if the utilitarian wants to influence more people to achieve more good, then he would do better to encourage them to pursue happiness-producing goals that are more attractive and more within a normal person's reach.'[28] She suggests that perhaps we ought to set up as ideals or heroes, not

---

[27]Ibid., 93-4

[28]Ibid., 87

those who lead a bleak, self-sacrificing life, but those who bring about *a lot of* good, but lead fuller, happier, more attractive lives, rather than treating as ideals or heroes lives which will not attract many aspirants. And I think there is certainly a lot to be said for this thought.

The thought should not, of course, be overplayed. It is surely a good thing that we hold up as heroes and ideals in this sense such people as Mother Theresa, Gandhi, Nelson Mandela and war heroes, all people whose lives were, to some degree, unattractive, often not well-rounded and involved great sacrifices.

The last three distinctions are even more important ones, and ones to which Wolf is perhaps not so sensitive as she should be.

7.4, 7.5- *WHAT WOULD BE IDEAL? WHAT IF WE ALL DID THAT?*
Wolf says, 'A world in which everyone, or even a large number of people, achieved moral sainthood- even a world in which they strove to achieve it- would probably contain less happiness than a world in which people realized a diversity of ideals involving a variety of personal and perfectionist values.'[29]

Utilitarianism, or consequentialism, certainly does not recommend moral sainthood in that sense, in the sense of moral striving for its own sake. It certainly does not hold up as valuable, or ideal, a world in which everyone is boringly identical. It would not be good if everyone had a similar temperament, narrowly focused on morality, and no-one pursued to any great degree their natural talents, the arts, sports, sciences and so on. This world would certainly be an impoverished one compared to our own.

Nevertheless, it is crucial to see that the relevant question is not 'What is the best kind of world we can imagine?', the answer to which is certainly not a world where everyone is dashing around desperately trying to prevent suffering. Instead, the relevant question- i.e. picking out the only sense in which I intend GC to pick out an ideal at all- is the question that GC answers, is: What do we have most reason to do, given the way the world actually is, and given the way others do in fact live their lives? What is the ideal, given that others are *not* moral saints?

If the world were kinder, and if more people made a substantially greater moral

---

[29]Ibid., 87

effort perhaps, then more equal distribution would result, and then my best way of promoting the good would be by focusing on my own concerns and projects and those of the people around me, developing my talents and interests, certainly not by busybodying all over the world, hunting down people in need. But of course, what it is best that we be, what we have most reason to do in the actual world, from the consequentialist point of view, is determined by the way the world is, by what others in fact do. To consider whether GC is a plausible theory, we should not be led astray by imagining a world in which everyone simultaneously and unexpectedly started behaving in the way that GC recommends as ideal.

### 7.6- *THE GOOD LIFE FOR HUMANS*

Finally, the question of what life we have most reason to live should not be confused with any purported account of the good life for humans. Obviously, the consequentialist- or at least the utilitarian- recommends as ideal a life which maximally promotes the elements of some such human good amongst all people, impartially considered. But in doing so, given the state of the world, it may well be that such an ideal utilitarian agent would in fact live a relatively impoverished life, lacking many features of that good himself, both personal projects and relationships, and the chance to develop many non-moral virtues, skills or interests.

As Wolf says, 'the moral virtues, given that they are, by hypothesis, all present in the same individual, and to an extreme degree, are apt to crowd out the non-moral virtues, as well as many of the interests and personal characteristics that we generally think contribute to a healthy, well-rounded, richly developed character.' And later, 'In other words, if the moral saint is devoting all his time to feeding the hungry or healing the sick or raising money for Oxfam, then necessarily he is not reading Victorian novels, playing the oboe, or improving his back hand.'[30] This is all uncontroversial, I take it. The optimal life from the consequentialist point of view, given the state of the world, would indeed for most people be one that was 'strangely barren', just as Wolf suggests.

In spite of this, the consequentialist claims that what most agents have most reason to do is to make such sacrifices in an attempt to relieve the many lives which are even more bleak than this. Given the state of the world, where extreme suffering is probably preventable by means of self-sacrifice on our part,

---

[30]Ibid., 81

this is what we have most reason to do. Again, the consequentialist does not value self-sacrifice in itself, but rather sees it as an evil necessary as a means to promoting much greater value at stake elsewhere. And again, it should be remembered, I am not suggesting, unlike the defender of MC, that such massive levels of self-sacrifice achieved by the utilitarian moral saint, are morally required.

## SECTION 8: 'Moral' Value and 'Moral' Reasons

Now I want to make some more general comments about the consequentialist's notion of value, and of reasons, and suggest that his position is far more in tune with Wolf's real worries when we see his conception of value and reasons stripped of the special moral taint that they are given by MC, moral consequentialism, by consequentialism as it is usually understood.

### *THE INTERESTS OF THE AGENT AND HIS FRIENDS*

Firstly then, it should be clear that the consequentialist can happily acknowledge the following:
1. It is not in my interest to be a moral saint or optimal agent.
2. It is not in the interests of my friends that I be a moral saint or optimal agent.
3. There are many (pro tanto) reasons not to be, or strive to be, a moral saint or optimal agent; i.e. there are many disadvantages or drawbacks to being, and trying to be, a moral saint.

To consider the third point first, I take it that this is precisely what Wolf successfully illustrates in the article, demonstrating all the undesirable features of being a moral saint or an optimal agent, in particular the fact that one will not have a well-rounded personality or lifestyle, and one would miss out on the chance to do a lot of valuable things.

If we concede the first two points, then we can accept Wolf's opening statement about moral saints: 'if there are [moral saints], I am glad that neither I nor those about whom I care most are among them'.[31] This seems pretty intuitive. The consequentialist can acknowledge that there are numerous disadvantages of

---

[31] Ibid., 79

being, and reasons against trying to become, a moral saint. The consequentialist point is just that these are hugely outweighed by the considerations grounded in the welfare of others in determining what we have most reason to do.

## *CONDITIONAL COMMITMENT*

Another important point about value is in response to Wolf's criticism that the utilitarian moral saint can only be 'conditionally committed' to his projects. In other words, even though he may be permitted to pursue some personal projects- enough to keep him psychologically healthy enough to keep promoting the good- he must always in a sense have one eye open for a chance to maximise the good, he must be ready to drop his projects.

There is definitely something in this common criticism against consequentialism, especially where it applies to MC, where one seems obliged to forever be ready to give up a pursuit 'at the drop of the moral hat', in Wolf's phrase, whenever the good could be better promoted elsewhere. GC certainly removes the moral pressure, i.e. the constant, pressing threat of being blameworthy. Nevertheless, it might appear that the *ideal* consequentialist agent, the consequentialist moral saint, might have to be like this. There are two things to be said in response to this, to further temper the criticism. Firstly, in general, being able to entrench yourself in whatever pastimes tend to promote the good will surely be a better bet than living with a mindset where one is always distracted by the thought of the value to be promoted elsewhere. If this is correct, the consequentialist will of course recommend the former mindset, one of entrenchment in what one is doing, since it will tend to produce more good. Secondly, some readiness or willingness to drop a particular commitment, whether it be a pastime, a relationship or whatever, when it becomes harmful or destructive, or when you could be doing something far more worthwhile, is a good thing. I return to these themes in my next chapter, in discussing Bernard Williams's criticisms of utilitarianism.

## *CONSEQUENTIAL VALUE AND MORAL VALUE*

Wolf frames her position by drawing a sharp distinction between moral values and other human goods; between moral reasons for action and reasons grounded elsewhere in 'what is good for humans'. But I think this distinction is a very

misleading one, at least from a consequentialist point of view. For the consequentialist, as I see it, there are not *sui generis* moral reasons or moral values. Moral values and moral reasons are not of a different kind to the rest of practical value and reasons for action. What grounds *any* reason for action according to the consequentialist is the same sort of thing in every case, namely, the value of the consequences at stake (for example, the welfare of all those affected).

We can thus fend off some of Wolf's main criticisms. For example, she says, 'The point is that, for a moral saint, the existence of these ['non-moral'] interests and skills can be given at best the status of happy accidents- they cannot be encouraged for their own sakes as distinct, independent aspects of the realization of the human good.'[32] But this is exactly what they are, for the consequentialist. There is reason to pursue them just because they promote the human good, not because they are a means to 'doing the moral thing'. If the reasons in favour of pursuing personal projects are in fact outweighed, that is only because there is *even more* human good, or human welfare at stake. The consequentialist simply puts in perspective the human good grounded in the agent's interests, by comparing it against the interests of others. It is essential to distinguish consequentialism, at least in the form of GC, from the overly moralistic mindset. Wolf, I think, is criticising the latter- and correctly so- rather than consequentialism itself. My concern here is to show that one is not subject to Wolf's criticisms just in virtue of having a consequentialist theory of value.

## *THE SUM OF ALL VALUABLE THINGS*

Likewise, the consequentialist, or at least the non-moralistic one who only defends GC, can refute Wolf's accusation that the utilitarian 'values these things only because of and in so far as they are part of the general happiness. He values them, as it were, under the description 'a contribution to the general happiness'. This is to be contrasted with the ways in which these aspects of life may be valued by non-utilitarians.'[33] This criticism of utilitarianism or consequentialism is deeply misconceived, I think. It seems tantamount to saying that everyone is committed to valuing each valuable thing under the description 'a contribution to the sum of all valuable things'. The consequentialist must strongly resist the

---

[32]Ibid., 85
[33]Ibid., 88

order of determination that this sort of picture implies he is committed to. It suggests that the consequentialist holds that human goods should be pursued because pursuing them is a means to doing what is morally right or optimal. In fact, the consequentialist should hold exactly the opposite: it is the case that there is moral reason, or sometimes a moral obligation, to act in a certain way just because of the human goods at stake. Again, it is an overly moralistic mindset that is at fault, rather than consequentialism itself.

Wolf says, 'the basic problem with any of the models of moral sainthood we have been considering is that they are dominated by a single, all-important value under which all other possible values must be subsumed'.[34] I think what Wolf should be attacking is not the idea that all values have something in common, but rather the idea that all value is distinctively 'moral' value. For the consequentialist, the valuable things in the world are what grounds all reasons for action. Whether we wish to call some of those reasons especially *moral* reasons is then a secondary question. The value at stake in situations we wish to describe as 'moral' should not be seen as a fundamentally different kind of value to the value that underpins a well-rounded, flourishing human life. What is at stake is not something distinctively 'moral'. The value at stake is not a different kind of value, from the consequentialist point of view. The value at stake in the consequentialist ideal is just the same as in the human ideal: human good, presumably in the form of, at least, welfare, constituted by the fulfilment of basic needs and the opportunity to flourish in just those ways that Wolf, Williams and others cherish so much.

Finally, to put the point another way, we should not strive for a world in which everyone is especially self-sacrificing. Rather, we do, and ought to, collectively strive for a world in which everyone can lead the sorts of fulfilling well-rounded lives Wolf has in mind. The consequentialist does not recommend self-sacrifice, or and moral striving for their own sake. Hence his ideal is a way of life that maximises the human good, or human welfare, not one which maximises the moral striving of the agent. The consequentialist ideal is thus not a 'moral' ideal in the sense that Wolf hints at.

---

[34]Ibid., 91

*WOLF'S POSITIVE VIEW*

In these last few comments, I have equivocated, just a little, between human welfare and what Wolf calls 'human good'. To correct this, I shall just make a brief comment on her positive view, which she outlines very sketchily at the end of her paper, and, in the light of what I have just been saying, suggest that such an account is under-motivated.

Wolf, then, recommends a new non-moral, yet non-egoistic perspective, which she calls 'the point of view of individual perfection'- 'the point of view from which we consider what kinds of lives are good lives, and what kinds of persons it would be good for ourselves and others to be', a point of view which 'provides us with reasons that are independent of moral reasons'.[35] This may just be equivalent to what I have called the good life for humans. If so, then I do not think this provides any reasons for action over and above those grounded in welfare and so on. If this is not what she means, then it seems fairly mysterious what sort of standpoint she has in mind. In fairness, her comments are, by necessity, fairly sketchy at the end of the article.

I think we can find what she is looking for if we shed the idea of a consequentialist perspective, the perspective of value or welfare, being an especially moralistic one. If Wolf is just rejecting a moralistic perspective, then her project perhaps has very similar motivations to mine. However, I think she is in danger of throwing the baby out with the bathwater. What is at stake in morality, at least on a consequentialist model, is value itself, the welfare and various goods that can accrue to people. We can grasp this sort of value, without the baggage of 'morality', and without the need to postulate any mysterious evaluative standpoint which is not grounded in welfare, or at least in valuable consequences. We can, I think, in this sense, have our cake and eat it.

Wolf says, 'The role morality plays in the development of our characters and the shape of our practical deliberations need be neither that of a universal medium into which all other values must be translated nor that of an ever-present filter through which all other values must pass.'[36] It seems absolutely correct that 'morality' is not 'a universal medium into which all other values must be translated'. But a de-moralised consequentialism like GC need not make any

---

[35]Ibid., 96

[36]Ibid., 97

claim like that at all.

My conclusion then is that the best thing we can do, or what we have most reason to do, is to become a moral saint, at least in the sense of being an optimal consequentialist agent. In this sense, and this sense alone, the consequentialist ideal is the ultimate or rational ideal. By making some important distinctions and by clarifying exactly what I take the consequentialist's claims to be, I hope to have shown that the most plausible form of consequentialism, i.e. one detached from the particular connotations of 'morality', is not vulnerable to the sorts of objections put forward by Wolf and Stocker, and indeed can accommodate the central concerns which motivate them.

## CHAPTER 5: Morality And Ethics II: Bernard Williams

Perhaps the philosopher whose criticisms of utilitarianism and consequentialism have been most influential is Bernard Williams. In this chapter, I show that Williams's strongest criticisms, like those of Anscombe, Wolf and Stocker, can in fact be accommodated by a consequentialism which denies MC, and limits itself to GC. In the first main part of the chapter, I focus on Williams's article, 'The point of view of the universe: Sidgwick and the ambitions of ethics'.[1] I will argue, with Williams, that the ambitions of ethics should, in some sense, be limited. And I will argue, against Williams, that my favoured, relatively unambitious form of consequentialism- one that asserts GC but denies MC- is not, as Williams charges against all ethical theory, incoherent. In the second main part of the chapter, I focus on one particular theme from Williams's writings, namely his distinction between morality and ethics, and make explicit how a distinction along these lines should be endorsed by the consequentialist.

Before considering Williams's views, however, it will be useful to clarify how I intend GC to be understood, and to disambiguate the claim it makes.

### SECTION 1: 'Ought', 'Should' and 'The Best Thing We Can Do'

GC says that 'The best thing we can do is bring about the best consequences we can.' Perhaps more has to be said here about what GC is offering a theory of; about what we should understand by 'the best thing we can do'. I have suggested already that to say that X is the best thing we can do is just to say that X is the most choiceworthy thing to do. But it must be emphasised that I intend this in a fully objective sense. The term 'choiceworthy' might itself be deemed ambiguous. For example, if I am in a situation where I lack knowledge of the facts relevant to determining what the objectively best thing I can do is, I still have to, in my present circumstances of less than full knowledge, *make a choice*. It might be thought that the 'choiceworthy' thing to do in such a situation is to do what is the best bet from your epistemic standpoint, to do what you have

[1]Williams (1995b)

most reason, or sufficient reason to believe you have most reason to do. To adapt an example of Allan Gibbard's, suppose I am lost in a forest and want to get out.[2] It may be that the most objectively choiceworthy thing to do is to go north-west, since this is the quickest way out of the forest. However, the question remains as to what I should do, given my ignorance of the fact that the quickest way out of the forest is to go north-west. In this subjective sense, the choiceworthy thing to do is clearly to walk in a straight line in a randomly selected direction. So there are (at least) two interesting questions that we might be asking when we ask, 'What is the most choiceworthy thing to do?' or when we ask, 'What is the best thing to do?' The point I wish to make clear is that in asserting GC, I intend to be offering an answer to what I have called the objective construal of the question, rather than any subjective construal.

This is important for a number of reasons. One is that, in everyday discussion, we use phrases such as 'the best thing to do is this' or 'you ought to do that' in a way that is relative to our state of knowledge or state of belief. In the forest example, if I meet someone who is similarly lost, and similarly ignorant of the quickest way out, and ask him what I should do, I do not expect him to answer according to the objective construal of the question, saying, for example, 'You ought to walk in the direction that is the quickest way out'. Given our shared ignorance about what that quickest way out actually is, this would be a singularly unhelpful response.

I intend GC as offering an answer to the objective question. So it may be that the best thing you can do on a particular day is take course of action A, which will result in the best consequences. But if you have no reason whatsoever to believe that A will result in the best consequences, then the practical implications of the theory are limited. This should not be thought an objection to GC however. The simple fact is that very often we have no idea what the best thing to do is, and for a theory of value to reflect this is a virtue of the theory. The great difficulty of determining in many situations what the best thing to do is does not change the truth of the issue as to what in fact is the best thing to do. In this sense, GC is not what Williams calls a 'principle for action'; it is not a ready-reckoner or an especially useful decision procedure. As the forest example demonstrates, GC can only be in a very limited way a guide to what to do in practical contexts.

---

[2]See Gibbard (1990), 18-19

Another thing to be noticed is that in everyday discussion, our talk in terms of 'ought' and 'should' is not just relative to our epistemic states, but often also relative to parameters created by prior decisions. If I am swithering between going to the cinema and going to the theatre, and ask you, 'What should I do?', it would strike us as very odd if you replied, 'Well, what you should do is write a cheque to Oxfam' or 'What you should do is go along and help out at a soup kitchen serving homeless people'. It may well be that from the consequentialist point of view, this is indeed the best thing you can do, what you should do. And I do not think that the consequentialist should drop his insistence that your response is a true reply to my question, literally taken. However, the oddness of the reply should show us that not all questions framed in terms of 'ought' or 'should' or 'the best thing to do' are intended in the 'parameter-free' sense to which GC purports to give an answer. This discrepancy between parameter-free ought judgements and parameter-relative ought judgements is not unique to paradigmatically ethical cases either. If I am deciding whether to smoke inside or outside, I might ask you, 'What should I do?' If you reply, 'You should eat an apple, and give up smoking altogether', a similar oddness results. It might be appropriate for a caring friend to say this once in a while when asked the question, but giving such an all-things-considered judgement every time would surely become very grating indeed.

So again, I am trying to emphasise how GC should be taken. It is offering an answer to the question, 'What is the best thing we can do?', taken in the parameter-free sense, rather than offering an especially appropriate answer to every question framed in terms of 'ought', 'should' or 'the best thing to do'. In asserting GC, I am not suggesting that we should somehow jettison our ordinary 'should' talk. Much of it may remain very useful, even from an ethical point of view. Agent A might say, 'Sure, I can see that the best thing I can do is to live the really self-sacrificial life, but realistically, that's not what I'm going to do.' We might ask in response, 'Given that A is not going to give up most of his income to help the world's poor, or even go off to Africa himself to help out, what *should* he do?' And this still seems a sensible enough question.

We can offer an ethical framework of 'second bests', of what is 'realistic'. 'You should, at least, do *this* much to help', we might say. But this sort of claim is not meant to supplant the all-things-considered judgement that

the very best thing you can do is live the extremely self-sacrificial life. As I have emphasised earlier, GC does not answer every interesting ethical question, and it is certainly not the only useful ethical claim. Bearing this in mind will be useful in considering Williams's discussions of 'ethical theory'.

## SECTION 2: The Ambitions of Ethics

Let me now turn to Williams's paper, 'The Point of View of the Universe'. His aim here is to discuss, taking Sidgwick as his starting-point, 'the possibility of ethical theory'[3]. He begins by saying that 'There are difficulties not only about Utilitarianism, but about the very project of a systematic ethical theory, which emerge with a special clarity from the pages of Sidgwick's book'.[4]

I am going to argue that consequentialism need not and should not, as I suggested in the last chapter, be in the business of providing a systematic ethical theory at all, at least an ethical theory of the kind that I think Williams is attacking. Instead utilitarianism or consequentialism should be more limited in scope and in ambition, in part for reasons which emerge in the light of some of Williams' concerns. To a significant extent, then, the position I argue for is not incompatible with Williams's views. Whether Williams would regard my limited consequentialism as an ethical theory proper is not my main concern. Rather, I focus upon showing that in attacking what he calls ethical theory, Williams does not impugn my form of consequentialism.[5]

This part of Williams's work should be of great concern to any consequentialist, not just because it forms a major part of his critique of consequentialism, but also because we ought, I think, to sympathise to a large extent with his view that moral philosophers, and perhaps utilitarians in particular are often overly-ambitious, frequently because of

---

[3]Williams (1995b), 153

[4]Ibid., 153

[5]Again, I will use 'consequentialism' and 'utilitarianism' interchangeably. Nothing in Williams's article or in my discussion hangs on how we wish to cash out the content of the good. I will use both terms because Williams speaks of utilitarianism, but I intend my discussion to apply to any impartialist consequentialist theory.

a narrow-mindedness and a bias towards what is local to them in time and culture.

Williams complains of Sidgwick, for example, 'It is a general feature of the book [Methods of Ethics] that the account of common sense morality is very linear, lacking any sense of the possibility of alternative moral traditions or the idea that certain moral outlooks of contemporary opinion might represent interests less broad than those of society or mankind as a whole.'[6] And also, 'The lack of any non-moral perspective on the morality of his time is a feature of the work in general.'[7] It is more because of these types of considerations that I share Williams's concerns about the ambitions of ethics, rather than because of his central worry about conflicts between points of view, which I will discuss later on.

Williams is especially alert in pointing out the various forms of ethical experience, ethical consciousnesses and the limitations of ethical theory. And it is in a spirit of humility about the ambitions of our ethical philosophy and an awareness of widely varying cultures, experiences and ethical mindsets throughout the world and throughout history, that we should make ethical claims. Too often perhaps, moral philosophers project the ethical consciousness of their own time, culture and even of their own ethical theory onto mankind, and inflate what is parochial and local into the universal and eternal. Given that the dictates of ethical theory generally claim to be universal, these sorts of worries may seem to leave us saying very little indeed. Universal ethical claims must be limited, formal and free of bold empirical assumptions, for example about psychology. The sort of claims often made by ethical theories that I think we should be cautious about are, for example, claims that we all have particular, distinctive moral motivations, such as the motivation to be able to justify ourselves to others, or the motivation not to be blameworthy.

In this paper, I will argue that the most plausible form of consequentialism, one which limits itself to GC, and rejects MC, gets round many of Williams's worries about utilitarianism, and about ethical theory in general. It is precisely because GC is so limited in its claims, and

---

[6]Ibid., 158
[7]Ibid., 159

relatively unambitious, that its claims, which are universal, are nevertheless very plausible, I hope to argue.

Williams sometimes speaks as if he is attacking only Sidgwick, but more often attacks utilitarianism more generally. Correspondingly, I intend in this chapter to defend my (limited and relatively unambitious) theory against Williams only in small part by defending Sidgwick, but more especially by distinguishing my utilitarianism from that of Sidgwick and others. In many respects, I agree with Williams's concerns about ethical theory in general, but I will try to show that these concerns do not undermine my limited consequentialism about value. It is not clear that Williams intends his critique of what he variably calls 'ethics', 'ethical theory' and 'systematic ethical theory' applies to a mere theory of value, even where that includes a theory of the value of *actions*. However, I think GC is an ethical theory at least in the limited sense that it explains how some actions can be better and worse than others. It is a genuinely normative theory in that sense.

GC is simply a theory of value, and a theory of what reasons for action we have. It can assess the value of other things than individual actions, for example motives, dispositions, practices, institutions and so on. In particular, GC can purport to assess the value of our moral practices, and of what Williams calls 'narrow morality' or 'Morality, the Peculiar Institution', which I will say more about presently. However, like Sidgwick's theory, GC does purport to offer some standpoint purporting to be objective, or even scientific, from which we can assess our moral practices. And it is this standpoint that will fall foul of Williams later on.

## SECTION 3: Deduction Versus Correction

Firstly, I want to briefly endorse one worry that Williams has about Sidgwick's ethical ambitions. That concerns the idea of *deducing* a comprehensive theory of moral rules from one (or more) fundamental principles. It is a major part of Sidgwick's project to deduce (more or less) our common sense moral rules from the principle of utility. Williams quotes Sidgwick, as aiming for one or more ethical principles 'more absolutely and undeniably true and evident, from which the current rules

might be deduced, either just as they are commonly received or with slight modifications or rectifications'.[8]

Such a 'deduction', I think, is an over-ambitious aim, certainly if the principle of utility is our fundamental ethical principle. To try to deduce a set of moral rules from the principle of utility is a pretty bold endeavour, especially if, as I think Sidgwick intends, we mean deduction in a strong sense, as opposed to merely using some principle of utility to shape or mould, in a non-systematic way, our principles and practices, endorsing some practices and criticising others.

Of course, it is an open question at this stage whether utilitarianism or consequentialism even recommends the institution of rules as such, whether we might not be better off without following ethical rules, or more plausibly, without following rules as characterised by (what Williams calls) Morality, the Peculiar Institution, rules which take the form of demands, violation of which merits blame and guilt, rules which are categorical, which we are faced with independently of our motivations, concerns and goals, and which are universalistic in form.[9] It might be for the best, for example, if some people did not rationally follow rules at all; if we make virtue or character or ideals the focus of our ethical consciousness.

There are a number of obvious problems with any serious project of deducing a code of ethical rules from the principle of utility. For example, there are the sorts of considerations I mentioned above about how what suits one culture, one time, one society may not suit others. If we are to deduce the ethical rules for every individual society, then the task is huge. Sidgwick's own efforts, of course, are already pretty substantial. Perhaps in principle we could, with every empirical fact on the table, deduce an optimal ethical mindset and optimal set of ethical 'rules' for any given society, but beyond this I do not think we should hold out much hope for the idea of deducing everyday moral rules from the principle of utility.

A further reservation that I have about this project of deduction is that by 'the principle of utility', Sidgwick clearly means something more like MC

---

[8]Sidgwick (1907; 1962), 102; quoted at Williams (1995), 158
[9]See Williams (1985), Chapter10

than GC. MC is clearly a 'principle for action' in a way that GC is not. GC rather explains how some actions or some behaviour can be more valuable than others, and is only a principle for action in the very limited sense that it recommends some action or behaviour as more *choiceworthy* than others. It is, however, entirely silent about any sanctions which might apply upon failure to act in certain ways. It is not, in itself, a theory of what can rightfully be expected of us, for example.

Instead, the most plausible practical application of utilitarianism is, I think, to take parts of our local, current everyday practices, and criticise them or defend them, moulding them, correcting them, by appeal to whether some alternative practices might produce more good, bring about better consequences. This sort of criticism can work on several levels. Let me suggest three, with examples:
1. We might criticise a particular practice, say slavery, on the grounds that it does far more harm than good.
2. At a higher level, we might criticise our moral standards- for example, according to common sense morality, we are not greatly at fault for failing to help the distant needy, or for failing to take measures to preserve the environment. We might criticise this pattern of moral sentiments and sanctions, on the grounds that if we treated negligence towards the distant needy or the environment as blameworthy, we could make a huge difference to the welfare of the distant needy and of future generations.
3. Beyond this, we might criticise our very morality, Williams's 'narrow morality', the 'Peculiar Institution', our whole ethical practice centred around blame, categorical obligations and so on.

We might of course, like Williams (and Anscombe, as mentioned in the last chapter), complain about narrow morality on the grounds of conceptual incoherence. However, it can also be criticised from the standpoint of consequentialism. This criticism might take the form that our system of moral practice and judgement grounds a harmful mindset, perhaps through focusing too much on negative rights rather than positive duties, or alternatively, through focusing too much on duties, rather than ideals, aspirations and so on. Perhaps it encourages and cultivates negative sentiments like guilt and blame, rather than positive ones like admiration. So-called 'Catholic guilt' is perhaps a good example of this sort of thought. It is suggestions like these, and associations with people using morality as a tool of oppression which in many discussions

gives 'morality' a bad name. One only has to sit in a first year moral philosophy class to realise that many see 'moral' as a dirty word, with connotations of authoritarianism, judgementalism and intolerance.

In any case, my conclusion is that if utilitarianism is to guide our practice at all, it will be in this limited way. It must be a piecemeal correction of our current practices, rather than a deduction from scratch of a comprehensive system of ethical rules from the Principle of Utility. Thus consequentialism should be less ambitious than it has been taken to be by its defenders.

## SECTION 4: Calculativeness

Williams introduces the main thrust of his argument by first defending Sidgwick and utilitarianism against a common objection, but by then going on to use it to kick off his own complaints. The bad objection, in Williams's view, is the idea that utilitarianism recommends that we be continually calculating what the optimal action is. Williams appreciates however that Sidgwick and utilitarians are not in fact committed to such a claim.

He speaks of 'Sidgwick's insistence, perfectly correct in itself, that a moral system should not imply that actions must always be taken as the result of conscious rational calculation, whether it be of prudence or of the universal good. It had been an old objection to Utilitarianism that it could lead to a denial of all natural affections and the stifling of impulse and spontaneity in the interests of a calculative spirit directed to universal good... This problem Sidgwick both saw and took resolute steps to avoid. He saw that from the point of view of Utilitarianism it must simply be an empirical question what motivations actually lead to the greatest good; and, in particular, whether the motivation of thinking about the greatest good is likely to lead to the greatest good... As Sidgwick puts it (p.345) the dictates of reason ought always to be obeyed, but it does not follow that the dictation of reason is always good.'[10]

---

[10]Williams (1995b), 162-3

So we must distinguish a relentless, ever-calculating mindset from a merely reflective one. Williams grants that Sidgwick, and even Bentham, were not as bound to the former as, for example, Godwin was. However, Williams is keen to attack not just the crude calculating rationalistic mindset, but also the role of reflectiveness itself in utilitarian theory.

## SECTION 5: The Cool Hour and the Deep Dispositions

Williams worries about a conflict within the theorist, and within anyone who takes utilitarianism to be correct. He starts from the worries, which Sidgwick himself discussed, that it is possible that utilitarians may be obliged to abstain from avowing it, or indeed to teach a quite different moral theory, if doing so will have the best consequences. Williams famously terms this picture 'Government House Utilitarianism', suggesting some powerful intellectual elite who are privy to the ethical truths, but keep them quiet from the masses for fear of the consequences.

Williams frames a slightly different objection. He says, 'Government House Utilitarianism is unlikely, at least in any overt form, to commend itself today. A version more popular now is to identify the required distinction between theory and practice as a distinction between the time of theorizing and the time of practice, and to use the notion, deployed in moral philosophy by Butler, of the 'cool hour' in which the philosophically disposed moralist reflects on his own principles and practice.'[11]

He has two objections to such a picture. Firstly, he attacks the idea that there may be a timetable according to which the two alternate. Secondly, he says such a picture is guilty of a 'naive conception of what is going on in the cool hour itself'.

He says, 'The belief that one can look at all one's dispositions from the outside, from the point of view of the universe, and that so doing is embodied in a cool hour of personal reflection, is a misrepresentation of that cool hour.'[12] Williams worries that our assessments of the value of

---

[11]Ibid., 166-7
[12]Ibid., 167

our dispositions and so on are no more likely to be accurate assessments when conducted during the cool, reflective hour than when conducted at any other time. But the utilitarian should, I think, resist this parody which Williams foists on to him. He certainly should not be seen as directly recommending a particular pattern or timetable; or prescribing a particular balance of reflectiveness with engagedness that is supposedly maximal. Indeed, as I understand the most plausible form of consequentialism, along the lines of GC, the theory is not committed at all to positing any particular pattern of behaviour, certainly not a one-size-fits-all alternation between action and reflection. The discrepancy here between Williams's conception and mine is that Williams sees utilitarianism as a fundamentally practical theory- a 'theory for practice'. Its acknowledgement is there to guide those who accept it as true, as a principle, almost like a motto or ready-reckoner. And perhaps, taken as a theory of moral obligation, as MC proposes, we might sympathise with Williams's view. In taking up MC as true, in adopting a utilitarian theory of moral obligation, one seems to incur some practical commitments. Williams seems to be suggesting that in taking MC to be true, one buys into a particular mindset or consciousness. That mindset is a deeply contradictory one, Williams thinks.

He uses the notion of 'deep dispositions' to illustrate his point. He says, 'From these considerations Sidgwick can offer a Utilitarian account of various things, particularly of certain dispositions of character which are often thought of as having an intrinsic or non-Utilitarian value. They include such things as the disposition to tell the truth, to be loyal to one's friends and to feel a particular affection and concern for one's own children.'[13]

The idea is that in coming to be utilitarians, that is, in coming to hold a practical theory, we bring upon ourselves a particular mindset, one which is divided against itself. Williams's main concern is that the utilitarian hails as valuable certain good-producing dispositions; but that these dispositions necessarily involve not having a utilitarian outlook.

Williams complains that Sidgwick and utilitarianism must treat the dispositions in an instrumental way: 'the dispositions are regarded just as

---

[13]Ibid., 163

devices for generating certain actions, and those actions, in the end, as the means by which practical reason produces certain states of affairs, those that minister to the universal good. That is what those dispositions look like when seen from the outside, from the point of view of the teleological Utilitarian consciousness. But it is not what they necessarily or usually seem like from the inside; and indeed what the Utilitarian argument may very well yield is the conclusion that they should not seem like that from the inside.'[14]

As I see it however, utilitarianism, or GC at least, does not recommend that utilitarians should intermittently shift to some higher-order, more objective state, far less that an impartial, reflective standpoint must be one which suppresses what the dispositions of truth-telling and loyalty to one's friends feels like on the inside. GC could of course endorse an empirical claim (if it were true) that things might be better if people have certain deep dispositions, even that we might be better off if no-one believes GC to be true. But there is nothing contradictory, or even disturbing in that. It cannot be the empirical claim that Williams objects to; rather it must be the perceived call for action.

## SECTION 6: The Utilitarian Consciousness

I think many of Williams's worries can be summed up by this phrase he uses a couple of times, 'the Utilitarian consciousness'. By this he is referring to a mindset which regards things as valuable instrumentally, as a means to an end. Williams sees this in Sidgwick. He says, 'it is at least perfectly possible- and Sidgwick regards it as true- that the Utilitarian consciousness should not, at least in many departments of life, be over-encouraged.'[15]

I think the utilitarian or consequentialist, at least the defender of GC, should resist strongly the suggestion that he is committed to something substantial called the 'Utilitarian consciousness', as Williams suggests. The phrase is suggestive of someone constantly and anxiously keeping an eye out for a chance to maximise utility, someone who can't fully commit to,

---

[14]Ibid., 164
[15]Ibid., 163

for example, a friendship, because he regards it solely as instrumentally valuable, and *feels under an obligation* to drop it when it ceases to be valuable. There are a few things that I think must be said in response to Williams's thought.

Firstly, and briefly, there is the well-rehearsed point that to say that X has value in respect of feature Y is not equivalent to saying that X is valued as a means to Y. So a utilitarian might say that friendship is valuable in respect of the happiness it affords, without claiming that friendship is valued as a means to happiness. Friendship is valued as an end in itself, just because (in part, at least) friendship involves happiness.

Secondly, we can concede much of what Williams wants. We can agree with him that we should not be under the illusion that we have transparent access to our dispositions, motives and so on. The utilitarian, like anyone else, can happily acknowledge the great limits on accessibility to our motivations and dispositions. Likewise, we should not regard reflection as 'authoritatively revealing'.

And of course Williams is right to say that our dispositions are integral to our humanity. Utilitarians have correctly, I think, taken this as a point about the psychological limitations we have- it would be self-defeating for us to try to shed to any great degree the things that we naturally care most deeply about, our friends, family and interests especially. However a bare belief in a consequentialist theory of value, which is all GC commits us to, should in itself do little to undermine our self-identity.

Thirdly, against Williams, there will of course be an empirical question about the valuableness of reflection on one's way of life, and on one's dispositions. It is hardly controversial to suggest that a life of occasional reflection will tend to bring about more good than a wholly unreflective one, which never bothers to question the first impulse, and the effects of one's behaviour and lifestyle on others. To be sure there is a balance to be struck. A tendency to be wholly unquestioning about one's friendships and about the appropriateness of telling the truth can be disastrous. Precisely what goes wrong in many relationships and friendships is a lack of perspective, an inability to get away from what it feels like from the inside. Surely Williams is not advocating blind faith in friendships and truth-telling? It seems in fact a good thing that the deep dispositions are not as

deep as Williams sometimes suggests. GC does not require us to be forever trying consciously and meticulously and relentlessly to perceive the world in a special, impartial way every time we act. In judging the value of various states of affairs, courses of action or institutions, of course we should try to shed our biases towards ourselves and those close to us. But making an impartial judgement of value is not the same thing as living a life. And there is no reason to think that what is desirable in the former pursuit is automatically desirable in every aspect of the latter.

Finally, and crucially for my purposes, I think that GC is far less vulnerable to Williams's sort of thought than MC is; that is, the notion of being caught up in a 'Utilitarian consciousness' seems far more plausible if one has a direct utilitarian theory of moral obligation. The notion of a utilitarian consciousness is of a mindset coming to dominate our every waking hour. But coming to believe GC, a mere theory of value, should entail no such onset of an oppressive consciousness.

Consider MC on the other hand. It claims that I am morally obliged to bring about the best consequences I can. The feeling that we are morally obliged is one that can haunt us in a way that the idea of best action does not. The idea of being continuously morally obligated seems to bring with it, not just a pressingness, but a demand to be continuously conscious of something. To be clear on the distinctive weightiness that *moral* judgements bring with them, consider the sorts of things we tend to say to describe the phenomenology of moral wrongdoing: 'I couldn't live with myself if I did that'; 'I wish more than anything that I hadn't done that'; 'I could never forgive myself', or 'I don't know myself anymore'. The self-condemnation that pends when we take ourselves to be on the verge of acting morally wrongly is distinctively painful and burdensome. Christopher Bennett describes the painfulness of moral guilt as 'partly the painfulness of a kind of self-division, or self-alienation'.[16] So if I feel that I am continuously on the verge of acting morally wrongly, as MC threatens, this is bound to weigh heavily upon me. It is the relentlessness that creates the discomfort, the schizophrenia, the constant threat of guilt and blame, the psychological tension that Williams ascribes to the utilitarian.

---

[16]Bennett (2003), 134. See also Skorupski (1999b)

The corresponding thought that I am very often acting suboptimally has no such tendency to weigh on us or haunt us. The thought that there might be more good in the world were my beliefs about goodness false is not a crippling thought to anything like the same extent.

To return to the discussion of Government House utilitarianism, it is not a remotely disturbing thought to imagine that teaching people that consequences are the ultimate bearers of value might in fact have bad consequences. Compare that to the unsettling thought that those who believe MC to be true seem, given certain empirical conditions, to be under a moral obligation to hide from the masses exactly what their moral obligations are. GC says simply that the best thing I can do is bring about the best consequences I can. Holding a stringent theory of moral obligation automatically weighs on the theorist in a way that a theory of value does not.

Consequentialism need not, and should not, be in the business of prescribing a universal ethical mindset. And if it is limited to being a mere theory of value, as GC is, it should not bring with it a particular mindset either. Of course, believing a consequentialist theory of value to be true will have some effect on our thoughts and behaviour. This is obvious and, I think, fairly trivial. And of course, it can offer us guidance in everyday life. However, this need not be anything disturbing or fragmenting.

The mindset of keeping one eye out, and not entrenching yourself in a friendship, in order to look for possible chances to maximise utility is doubtless an unhealthy one- a mindset, therefore, that the consequentialist will not in fact recommend. Even though I think the defender of MC may not strictly be committed to this mindset of keeping one eye out, never fully committing, one can see where the pressure comes from. It comes from the threat or possibility that I might be acting morally wrongly.

But in response to Williams's central worry about the conflict between deep dispositions and the utilitarian consciousness, I hope to have deflated this apparent conflict, firstly by suggesting that the dispositions need not be quite so deep as Williams thinks desirable, and secondly, that there need be no over-powering utilitarian consciousness at all which might have been on a collision course with such dispositions.

SECTION 7: Theory For Practice

Williams's worries then, about Government House utilitarianism, about the cool hour, about valuable dispositions incompatible with a utilitarian mindset are all grounded in his concerns about the relationship between ethical theory and practice. Williams believes that any ethical theory must be first and foremost a theory 'for practice'. If this is so, then my favoured form of consequentialism may not be an ethical theory at all. GC is not solely, or even primarily, a theory 'for practice'.

Of course when we discuss our everyday moral practices, which rules to respect, which codes to adopt, and what we should regard as expected, we must take into account such things as transparency and compatibility with our psychological limitations and our natural 'deep' dispositions. But GC is not a theory of these things.

GC is not intended to be a moral theory in this loaded sense of 'governing practice' at all. Rather, it is a theory of value. Although again naturally, and trivially, the behaviour of people who believe a utilitarian theory of value to be true will of course be affected by their belief. GC is a theory of reasons, not in the sense that it recommends a particular mindset, but rather, solely in the sense that it explains how something can be a consideration that counts in favour of an action.

Williams objects to any ethical theory which stakes a claim to make judgements from an objective standpoint because the very idea of an ethical outlook is necessarily entrenched in an agent's own point of view, one which is paradigmatically not an objective point of view. He draws what he calls a 'disanalogy with the philosophy of science'.[17] He says 'Even though one cannot reconstruct a scientific body of belief starting totally from outside it, nevertheless it is an aim, in reconstructing it, to free one's view of the world to the maximum degree from perspectives peculiar to one's historical or local situation- to correct for observer's bias, and indeed to try to see the universe from the point of view of the universe, that is to say from no distinctive point of view at all.'[18]

---

[17]Williams (1995b), 170
[18]Ibid., 170

I take it that this is indeed what an impartialist consequentialist theory of value, just like science on Williams's view, aims at. However, Williams also says, 'I agree with Sidgwick, such [an ethical] theory must aim to be a theory for practice, and to be closely related to reasons for action. It cannot be a reasonable aim, with regard to that purpose, that I or any other particular person should take as the ideal view of the world- even if one then returns from it to one's self- a view from no point of view at all.'[19] But if GC is not, as I have argued, primarily a theory for practice, then it is not bound to any practical 'purpose' and thus not bound to a subjective point of view. So if someone thinks that science can rightfully stake a claim to objectivity, but that value theory cannot, then the onus is on him to explain why.

Given Williams's views of ethical theory, he concludes: 'My own view is that no ethical theory can render a coherent account of its own relation to practice: it will always run into some version of the fundamental difficulty that the practice of life, and hence also an adequate theory of that practice, will require the recognition of what I have called deep dispositions; but at the same time the abstract and impersonal view that is required if the theory is to be genuinely a theory cannot be satisfactorily understood in relation to the depth and necessity of those dispositions. Thus the theory will remain, in one way or another, in an incoherent relation to practice. But if ethical theory is anything, then it must stand in close and explicable relation to practice, because that is the kind of theory it would have to be. It thus follows that there is no coherent ethical theory.'[20]

In order to reply to Williams, I wish to return to my earlier discussion distinguishing the purportedly universal claims of moral philosophy from particular practical codes which may suit our time and our culture, but are not universal in form. It is the latter that are theories for practice in Williams's narrow sense. A theory for practice must be a code for our time, our culture, our consciousness, not purporting to be universal and eternal.

When we consider Sidgwick's fundamental question of ethics, 'in what way is it reasonable or rational to live?', we can answer in two distinct ways. The first answer is the simple consequentialist answer: it is most

---

[19]Ibid., 170
[20]Ibid., 171

reasonable to live in that way which will bring about the best consequences. But that in itself is of limited help 'for practice'. GC is not a principle for action like 'Don't steal' or 'Tell the truth'. The second type of answer will be one which makes reference to what is peculiar to the age and culture in which we live, and perhaps even to the specific individual, his talents, temperament and situation. If moral philosophy's claims purport to be universal, they must be limited to formal claims like GC, or else to universal claims with an endless list of conditions specifying the situation of the individual agent.

Of course we can say much more in response to the fundamental question than simply GC. We can make all sorts of generalisations about what action-types, temperaments, character traits and dispositions tend to bring about good consequences in our particular age. Perhaps these constitute something called 'theory for practice'. Take, for example, our talk of rights, in particular, human rights. This talk is so entrenched, especially with regard to international affairs, that many philosophers and theorists concede that they must speak in its terms, to some degree, even though they do not believe there to be anything out there in the world which has the shape of a right; there are no natural human rights, no natural law; and no non-arbitrary level of what we can 'rightfully' expect of each other. This seems to be an example of theory for practice, in Williams's sense. Theory for practice is an area of talk or a framework of concepts whose function or raison d'être is to guide action.

GC purports to be a truth of a stronger kind than claims about human rights regarded in this way. GC is first and foremost a descriptive theory. The crucial point is that it is to be distinguished from Williams's conception of a 'theory for practice'. If an ethical theory is something which must run together scientific, metaphysical or meta-ethical claims with 'theory for practice', then perhaps Williams is right that there can be no coherent ethical theory. However, I hope to have outlined a model on which the two need not be run together.

## SECTION 8: Morality and Ethics

Now I wish to turn to another theme prominent in Williams's work, his distinction between morality and ethics. In my discussion of Anscombe in

the last chapter, I spoke of how consequentialism should stand outside of morality, narrowly understood. In this second part of this chapter, I aim to make explicit precisely what position I see consequentialism occupying in ethical theory. Since I deny that consequentialism is a plausible theory of moral obligation, as it has traditionally been understood, I hope to clarify just how consequentialism can, and should, stand outside of morality.

In this chapter and the last, I have discussed lines of thought from Anscombe, Stocker, Wolf and Williams which express disillusionment with modern moral thought, and even with morality itself. Again, this sort of critique can have two strands. Firstly, there is the questioning of the empirical assumptions and conceptual categories of moral thought. We can criticise morality on the basis of the false presuppositions it makes and possible conceptual incoherence. Secondly, we can criticise, or at least question, morality on the basis that it stifles other things we take to be valuable, or ignores them altogether. This opens up the possibility of evaluating morality by the lights of standards other than its own. As Williams says, 'the virtues [of the means by which morality deals with deviant members of its community] can be seen as such only from outside the system, from a point of view that can assign value to it, whereas the morality system is closed in on itself and must consider it an indecent misunderstanding to apply to the system any values other than those of morality itself.'[21]

Williams of course takes much of his impetus from Nietzsche's scepticism about morality. Although Nietzsche is sometimes misunderstood as a nihilist about all value, he does in fact favour other forms of evaluation over moral evaluation. He not only questions the metaphysical assumptions on which our moral judgements are based, but also criticises our moral practice as undesirable. But this latter kind of criticism cannot take place from an evaluative nowhere; rather, it presupposes some preferred standard of evaluation by which we can judge our current practices as desirable or undesirable. In order to perform a 'revaluation of all values' of the sort Nietzsche discusses, there must of course be some privileged evaluative standpoint from which the revaluation of the old values takes place. For Nietzsche, the privileged evaluative standpoint is that of the advancement of ideals of (super-)human excellence and

[21]Williams (1985), 195

perfection.[22] But an impartialist consequentialist has a ready-made, and very plausible evaluative standpoint of his own, by which to assess the value of morality. We can firstly evaluate our moral practices by considering whether they themselves produce good consequences relative to alternative sets of ethical practices, and we can also evaluate our moral theory, our moral categories by considering whether they underlie and support a healthy set of ethical practices or an unhealthy one.

Williams says, 'Morality is not one determinate set of ethical thoughts. It embraces a range of ethical outlooks; and morality is so much with us that moral philosophy spends much of its time discussing the differences between those outlooks, rather than the difference between all of them and everything else.'[23] Here, of course, Williams intends that utilitarianism, or consequentialism, be understood as one of the possible 'outlooks' internal to morality. My proposal is instead that consequentialism has the tools to criticise morality itself, from the outside. Rather than view consequentialism as one narrowly moral theory competing with others, consequentialism should stand one step back.

In this section, my aims are as follows. I begin by outlining briefly what I take morality to be, the narrow notion of morality which is the object of Williams's sustained critique. Much of his critique, and those of Anscombe, Stocker and Wolf, as I made clear in the previous section, is not just of our moral thought and practice, but of utilitarian moral theory in particular. My recommendation however is not that the utilitarian, or consequentialist, should defend a consequentialist moral theory against these attacks, but rather to join these philosophers in their critical standpoint outside our moral thought and practice, and question it along with them, refusing to take the categories of narrow morality for granted. A consequentialist about value- the defender of GC- is already in a good position to question the value of our morality, since he has a ready-made, and very plausible, preferred standpoint of evaluation. I go on to consider what sorts of consequentialist criticisms one might make of morality. Finally, I consider the relationship between consequentialism and ethics more generally. Is the consequentialist standpoint in some sense *the* ethical standpoint? Is it the ultimate or fundamental ethical standpoint?

---

[22]See Nietzsche (1990) and Nietzsche (1994)
[23]Williams (1985), 174

SECTION 9: What is Morality?

It should be emphasised firstly that the distinction between 'narrow morality' and 'wider ethics' is, to a large extent, stipulative terminology of art, rather than a distinction well-grounded in ordinary language. How then are we to characterise 'morality' in the narrow sense?[24] What is it that Nietzsche and Williams are attacking, and what is it that the consequentialist is seeking to evaluate?

For Williams, 'morality' is the particular mode of ethical thought which is dominant in our society; 'a particular development of the ethical, one that has a special significance in modern Western culture'.[25] Likewise, what I intend to pick out by the term is *our* particular ethical practice, the actual social phenomenon. While John Skorupski describes morality as 'a distinct sphere within the domain of ethics'[26], I think Williams's characterisation of morality as 'a local species of the ethical'[27] is an attempt to emphasise the *contingency* of morality's domination of modern ethical thought.

For Williams, morality in this sense is just one amongst many possible ethical systems or schemes. It is possible for a society to have, instead of a morality like ours, some other 'scheme for regulating the relations between people that works through informal sanctions and internalized dispositions'.[28] He says, 'Human beings have lived, do live, and doubtless will live, by conceptions of the ethical that do not invite these problems [the problems distinctive of morality], or invite them in a much less dramatic form. For these problems to be endemic not just in morality (in the narrow sense) but in our life, it would have to be the case that morality (in the narrow sense) was inevitably the form that ethical life took, and this is not so.'[29]

What then is it that is distinctive about *our* morality? Two things (closely linked) that make narrow morality distinctive are the particular categories

---

[24]I say a bit more about the characterisation of 'wider ethics' in the final section of this chapter.
[25]Ibid., 6
[26]Skorupski (1998), 564
[27]Williams (1993), 242
[28]Ibid., 241
[29]Ibid., 242

in which its judgements are couched, and the sentiments which are characteristic of the moral impulse. The central categories of morality are those of rightness, wrongness, obligation and forbiddenness.[30] As Williams points out, these terms have non-moral senses, which generally fail to have the particular connotations which the moral versions have. Thus, in turn, we can pick out peculiarly moral obligation by considering the sentiments which ground judgements which use these terms, namely the sentiments of blame and guilt.[31]

The central moral category of obligation-grounded-in-blame has certain nuances that are absent from the notion of obligation in play in other ethical schemes. For instance, Williams characterises morality as 'the local system of ideas that particularly emphasises a resistance to luck'.[32] His thought here is that, according to a central strand in moral thought, one can only be justly blamed for what was within one's control, where this notion of what is within one's control is luck-free.

With this rough characterisation, I hope to have made reasonably clear the object in question when one considers how a consequentialist might criticise, or at least evaluate, 'morality itself'.

## SECTION 10: How Might Consequentialism Criticise Morality?

Again, there are two families of criticisms we might make of morality, which might respectively be termed conceptual and evaluative. Firstly then, a consequentialist can of course have the same conceptual worries about morality as anyone else. For example, one might argue that morality

---

[30] Indeed, one of the main focal points of Williams's attack on morality is just its over-concentration on the category of obligation, its tendency to frame all ethical judgements in terms of obligation. 'Many philosophical mistakes are woven into morality', he says, the main one being that it 'misunderstands obligations, not seeing how they form just one type of ethical consideration'. Williams (1985), 196

[31] A further question might be asked as to how we characterise blame itself. Skorupski suggests, for example: 'the sentiment of blame itself disposes to withdrawal of recognition, expulsion from the community'. Skorupski (1998) We might worry that there are clear non-moral cases of blame: e.g. watching a football game, we might say 'the winger was to blame for the move breaking down', where there is no suggestion of the winger meriting *moral* blame. Some sort of story then about what picks out moral blame from non-moral might be offered in terms of the seriousness of the moral, its other-regardingness and the sort of dispositions to social sanction that Skorupski suggests here.

[32] Williams (1993), 241

is built on certain false empirical or metaphysical assumptions, such as the thought that we are all equally autonomous, or equally capable of recognising and acting upon moral considerations. Similarly, one might suggest that different conceptual strands in our moral categories pull in different directions, threatening to make the moral concepts incoherent. This might plausibly be the case for the two main strands in the concept of moral obligation I highlighted in Chapter 2, namely choiceworthiness and blameworthiness.

But beyond these possible conceptual criticisms, I wish to focus in this section upon possible evaluative criticisms of morality that are distinctively consequentialist.

The consequentialist cannot of course evaluate morality in isolation from consideration of alternative ethical practices which could, at least in principle, take its place. An 'ethical anarchy', in which no sanctions of the sentiments are internalised within us, and in which our behaviour goes unregulated by our own feelings and the feelings of others would surely not be likely to result in good consequences. Thus we must consider the possibility and value of alternatives to prevailing morality. A useful way of considering such a comparison between different possible disciplines is to start with the question 'What is morality good for?'. We might answer that it allows humans to live together in relative peace and trust, enabling us to flourish more effectively, in the knowledge that we can rely upon the non-interference of others with our projects, and upon certain limited kinds of co-operation.

Considering morality in such terms need not amount to seeing morality as having some essential function, but recognising the benefits that our moral practices confer upon us might firstly account for the evolution of morality, and explain to some degree why we persist in taking morality seriously. And seeing morality in this way naturally invites the question, 'Might something else meet the same needs, only better?' Perhaps there is some other discipline, other than the prevailing one in which guilt and blame are central, which would be a better means to the desirable end that is served to some degree by our moral practice, or which would afford us similar benefits without the distinctive harms that morality brings.
Now any evaluative comparison between morality and alternative informal systems of mutual discipline will of course be hugely speculative,

since there are so many unknowns. But my point is that evaluation of morality as a whole does seem philosophically desirable. It would be unsatisfactory to allow morality to go unquestioned, and for its categories to be taken for granted. And more specifically, the consequentialist standpoint does seem to be a proper point of view for evaluation of morality. It seems very plausible that it is considerations of human welfare, for instance, that are relevant in determining the value of an ethical system and practice. We might, for example, reject Nietzsche's rejection of morality precisely on the basis that the alternative ethical frameworks that he would like to see replace morality in our consciousness are likely to be significantly worse in impartialist consequentialist terms than morality.

Williams, while of course far from endorsing any impartialist consequentialist evaluation of morality himself, is quite aware of how, from a consequentialist (or, in his words, utilitarian) point of view, the value of different forms of ethical deliberation and ethical mindset are up for grabs. 'For utilitarians, what is important is that there should be as much welfare as possible. The connection with deliberation is a subsequent question, and it is entirely open. We saw when we considered indirect utilitarianism how the question is open of what moral considerations should occur in a utilitarian agent's deliberations. More than that, it is open whether any moral considerations at all should occur in them... for any utilitarian it should always be an empirical question: What are the implications for deliberation of welfare's being important? In this respect, however, there are many utilitarians who belong to the morality system first and are utilitarians second.'[33]

My suggestion is that consequentialists should, if anything, be consequentialists first, and part of the morality system second, and then only if the morality system is consequentially valuable. It could turn out that peculiarly moral deliberation generally fails to result in good consequences. If some other form of ethical deliberation has better consequences, this will be endorsed by consequentialism. The consequentialist could even recommend less ethical deliberation, less reflection in some circumstances.[34]

---

[33]Williams (1985), 184

[34]Obviously, if our morality does somehow conflict with consequentialism by being consequentially harmful, and if our moral intuitions are themselves so strong as to outweigh

Even if we ought to be cautious in our conclusions, there may still be much that we can say about the consequential value of narrow morality as a whole. We know where to look in evaluating it, even if we are more tentative in criticising it outright than Williams, for example.[35] Here are some possible changes of emphasis in our ethical consciousness which might be thought an improvement upon the moral status quo of an ethical experience centred on blame and guilt.

(1) PRAISE, IDEALS

Firstly, we often think of the central third person evaluative sentiments of morality as being praise and blame. But there is no doubt that the second of these is more dominant. And it is quite open to a consequentialist critic of morality to argue that an ethical experience which put greater emphasis on positive sentiments at the expense of the prevailing negative ones would ground a healthier ethical practice. It seems plausible to argue that an ethical education, for example, focused on altruistic aspirations, encouraging praiseworthy traits, rather than focusing on minimising blameworthy ones might be more constructive. While moral praise is surely a part of narrow morality, it often seems to play a secondary role to its negative counterpart, blame. And this has been more true within moral philosophical discussion.[36]

The focus of the moral consciousness does seem to be wrongness and requirement. It could be that an ethical consciousness built around ideals and goals, rather than demands would produce better consequences. An agent's being motivated more by the thought of praiseworthiness, rather than the motivation to avoid being blameworthy, might form part of a more constructive and productive ethical character. Likewise, from the third-person perspective, a readiness to make positive ethical judgements about others, to see the best in people, the praiseworthy, as opposed to a disposition to blame or condemn, might plausibly be considered an ethical

---

our confidence in consequentialism itself, then we might go the other way. In saying that we might be consequentialists first, I am not denying the wisdom of a reflective equilibrium in our methodology, but rather suggesting that when we carry out such a reflective equilibrium, the advantages of consequentialism are significant enough to outweigh the number of moral intuitions which conflict with it.

[35]Williams's conclusions about morality are generally bold and scathing. For example: 'Its philosophical errors are only the most abstract expressions of a deeply rooted and still powerful misconception of life'. Williams (1985), 196

[36]See Urmson (1958) for errors which arise from this misemphasis.

mindset more likely to promote human flourishing and harmony. Talk of the possibility of more positive ethical mindsets might sound somewhat naive or utopian. Nevertheless, Nietzsche, for example, might sympathise with this line of criticism. One of his most vivid attacks on morality focuses on what he takes to be its typical self-pitying but vindictive thought, 'I am suffering; someone must be to blame'.[37] It is this sort of pathological negative mindset which some might think is characteristic of the moral outlook that could form the basis of a consequentialist rejection of morality. We can, of course, question whether this sort of mindset is integral to morality, or whether it is simply an adulteration of a healthier set of social practices.

In a similar vein, it is quite plausible that people's association of morality with judgementalism and intolerance has bad consequences. If morality then is apt to encourage judgementalism or intolerance more than other possible ethical systems, then this can motivate a consequentialist worry about morality. If this disillusionment then drives people from acceptance of objective ethical norms towards subjectivism, this may be another disadvantage of morality.

(2) EQUALITY AND INDIVIDUALISM

A second aspect intrinsic to morality which might be thought a target of consequentialist evaluation is its emphasis on particular kinds of equality. Of course, most consequentialism (and the kind I assert) is avowedly impartialist, and generally claims that no person's welfare is more important than that of anyone else. In this respect, consequentialism is usually considered egalitarian itself. Nevertheless, it is still open to a consequentialist to argue that the moral categories and our moral practices support an egalitarianism which is somehow stifling or harmful.

For example, there is the strand in our moral thought which insists upon the equal capacity of all to respond to certain types of reasons. This may have bad consequences in holding people responsible for behaviour they could see no reason to avoid. This may incline us towards retributive punishment of people for whom education and rehabilitation would be more appropriate and more productive.

---

[37]Nietzsche (1994), 99-100

Conversely, a central strand in morality perhaps claims that two people, X and Y, are equally blameworthy for performing a similar action-type, even though they have different abilities, different background beliefs and so on. One might argue that those capable of great achievement are entitled to take greater risks than others. Williams's semi-fictional 'Gauguin' might be considered someone who violates the standards of morality, but who, by consequentialist standards, is justified in abandoning his family in an attempt to achieve artistic greatness.[38]

One alternative way of fleshing out a criticism of morality grounded in its emphasis on equality might be by reference to a distinction between so-called 'blame cultures' and 'shame cultures'. If one could argue empirically that cultures in which blame judgements were less dominant in the ethical practices of a society than ethical judgements grounded in, say, shame, disdain, honour and admiration, promote human flourishing and achievement to a greater degree than blame cultures, this could form the basis of a consequentialist critique. One major difference which marks the difference between these two supposed types of culture is that a blame culture assumes an equality of status in a way than shame cultures do not.

Blame-morality sometimes seems to insist on an unquestioning equal respect for all, a respect which is not grounded in actual features of individuals. From outside the perspective of morality, it may in fact seem very obvious that all people are not worthy of equal respect; that respect is due to people according to their abilities, achievements, wisdom, experience and so on. Perhaps a consequentialist line of argument could be developed suggesting that harm results from morality's tendency to treat people as equal in respects in which they are manifestly unequal.

(3) INDETERMINACY

A further respect in which we might think that our moral categories are unsatisfactory and could be improved upon concerns their ability to deal with what are arguably the most pressing global ethical issues- the welfare of the Third World poor and the state of the natural environment. Our moral categories and sentiments have evolved to deal with local, small-scale dealings with others. When we look to the global, large-scale dealings with, for example, the distant poor, or future people affected by our

---

[38]Williams (1976)

treatment of the environment, it becomes radically unclear what we should say from within the confines of our ordinary moral categories.[39] How we might improve in this respect is perhaps unclear. But this indeterminacy or incompleteness may have bad consequences in that we are ill-equipped, both individually and collectively, by our current ethical thought, to respond to global ethical problems. I return to this issue in the final chapter.

The extent to which we ought to regard the questioning of these central features of morality as a rejection of morality *per se* is not clear-cut. Perhaps much of our criticism can be accommodated, and what remains still be called 'morality'. In his book *Human Morality*, Samuel Scheffler for example, argues that many of Williams's main concerns can actually be accommodated within morality.[40] It would similarly be open to the defender of morality to attempt to accommodate any consequentialist objections within the framework of the moral categories.

Also, it is a matter of dispute as to the extent to which societies could actually function without something very like blame-morality. And just to say that some societies have given, or do give, blame and guilt a less central role in their ethical thought and practice, is not to say that *we* could easily, or spontaneously, change. It may be humanly possible to play down these sentiments, without it being very simple or without cost for *us* to change deliberately from one set of dispositions and practices to another. So, even if we can pinpoint parts of our moral practice which could be replaced by consequentially better practices, the implications for what we should do are not immediately clear, given that our patterns of sentiments are deeply ingrained and socially entrenched.

## SECTION 11- The Fundamental Ethical Standpoint?

Finally, I wish to consider whether the consequentialist standpoint is in some sense the fundamental ethical standpoint. I think we can answer this question positively, but with a familiar caveat that I emphasised earlier-

---

[39]See Fishkin (1982).
[40]Scheffler (1992)

that consequentialism should not be taken to be an ethical theory of everything.

We must first consider what the boundaries of the notion of the ethical are. An impartialist consequentialist standpoint may indeed purport to be the ultimate ethical standpoint with regard to the questions of the source of reasons for action. However, as I argue below, we should not be tempted by a direct consequentialist account of, for example, reasons to feel, or even of all questions of value or good broadly construed (for example, aesthetic value), all of which might come under a broad heading of 'ethics'.

A further question to consider is how the ethical relates to the practical. It seems wrong to say that all practical questions are ethical questions. For example, the question of whether I should wash my face or brush my teeth first is a certainly a practical question, but hardly an ethical one. This is because we generally consider the ethical as relating specifically to our interaction with others.

We might question, likewise, whether egoism is really an ethical theory proper, since, as Williams points out, egoistic considerations do not seem to be ethical considerations.[41] Perhaps being in the business of doing ethics even presupposes a concern for others, or a willingness to take their interests, desires or welfare into account in deciding what to do. Nevertheless, I think, contra Williams, that it is plausible to consider normative egoism an ethical theory in one sense. It is after all a competitor to normative ethical theories with altruistic elements; it is aiming to take up the same ground as these. Egoism does seem a candidate answer to the central questions of ethics, 'How should I live?' (what Williams calls Socrates' question) or 'What should I do?' Further, while we are reluctant to call egoistic considerations ethical considerations, they are certainly relevant in answering ethical questions.

Perhaps not so much hangs on where exactly we draw the boundary between the ethical and the non-ethical. My point here is that the sense in which we might think an impartialist consequentialist standpoint to be the fundamental ethical standpoint is the sense in which it claims to provide the ultimate answer to the central questions of ethics, 'What

---

[41]Williams (1985), 11

should I do?', 'What is the best thing for me to do?' and 'How should I live?'

However, let me now turn to the caveat I referred to earlier- that consequentialism should not be understood, as it has often been in the past, as an ethical theory of everything. There are many interesting and important ethical questions to which the consequentialist, qua consequentialist, does not give answer. Let me offer three examples, concerning issues that might be thought to fall within the domain of ethics, broadly construed.

Firstly then, consider the question of what makes an agent virtuous. It would be hopeless to say that an agent is virtuous if and only if he brings about the best available (or good enough) consequences. An agent who strives to do what he takes to be the best thing but is repeatedly unlucky should still be considered virtuous. Rather, we tend to judge whether an agent is virtuous according to his motives, his intentions and his efforts in doing what he takes to be right or good. A consequentialist need not offer a specifically consequentialist account of what makes for moral worth in this sense.

It is perhaps a common misconception in moral philosophy that consequentialism is committed to playing down questions such as these. For example, it is often asserted that the difference between consequentialism and Kantianism is that according to the consequentialist, 'what really matters' is the consequences of one's actions, while, according to the Kantian, 'what really matters' is one's motives and intentions. But of course, we may just be answering different sorts of questions:
(i) What makes for best action? What reasons for action are there?
vs.
(ii) What makes for moral worth? What makes someone blameworthy or praiseworthy?

And a consequentialist who was so inclined can offer a theory of what makes for the moral worth of the agent, perhaps in terms of effort or striving or good will. Having a consequentialist account of value and reasons for actions is quite compatible with a variety of accounts of the

moral worth of an agent built on various different metaphysics of the will.[42]

Secondly, consider questions of what there is reason to feel, many of which might be thought to fall within the remit of ethics. A direct consequentialist account of reasons to feel is not plausible. There is not, for instance, reason to feel blame towards an agent who has acted with perfect intentions but whom it would be consequentially pragmatic to blame. It follows from GC that if performing a 'blame-action' towards an innocent person will, in unusual circumstances, have the best consequences, then this is the best thing we can do. However, this is not to say that they are blame*worthy*, that we have any reason to feel blame towards them. Similarly, perhaps in a particular situation, feeling sad, or even despair, will have no good consequences whatsoever. Nonetheless, there may well be reason to be sad, or to feel despair. Just because feeling sad will do no one any good when one is bereaved, this does not mean one has no reason to feel sad when one is bereaved. Thus a direct consequentialist account of reasons to feel is not plausible.

Thirdly, I think that not even all questions of value can be given a direct consequentialist answer. In other words, it is not the case that in all circumstances that we use the term 'good', we must be in the presence of something that is correctly evaluated in terms of consequences. Aesthetic value is an example: we do not consider a work of art to be good just because it will in fact produce good consequences. A great work of art might be tragically lost in a fire before anyone has had the chance to appreciate it. This does not affect (or at least does not solely determine) how good a piece of art it is. Of course there may be some informative story to be given about the relation between aesthetic value and consequences. One might think a good work of art is generally one which has the potential to afford pleasure, insight, self-understanding and so on. But even this surely only plausibly provides a very partial, incomplete account of aesthetic value.

---

[42]Perhaps a consequentialist account of the individual virtues would be more promising. We may think that were it to be discovered and widely known, for example, that benevolence or courage somehow caused great harm, we would cease to regard these as virtues, and instead treat them as vices. For an argument along these lines, see Driver (2001) and against Driver's view, Skorupski (2004)

My point in these three examples is that, in suggesting that the consequentialist standpoint is a central one within ethics, this does not amount to an implausible claim that consequentialism provides a plausible direct answer to all ethical, aesthetic and evaluative questions. Again, I think that the consequentialist can learn from Williams. He says, 'If there is such a thing as the truth about the subject matter of ethics- the truth, we might say, about the ethical- why is there any expectation that it should be simple? In particular, why should it be conceptually simple, using only one or two ethical concepts, such as *duty* or *good state of affairs*, rather than many? Perhaps we need as many concepts to describe it as we find we need, and no fewer.'[43] It is a mistake for the consequentialist, in asserting GC, to deny the importance of other, distinct questions that fall within the domain of ethics.

In conclusion then, the consequentialist thus should (a) endorse the sort of distinction made by Nietzsche and Williams between narrow morality and wider ethics, and (b) take up a standpoint outside of morality, rather than seeing himself as propounding a moral view competing with, say, a Kantian criterion of moral rightness and wrongness. While it may be plausible in one sense to say that the standpoint of consequentialism is the fundamental ethical standpoint, this is not to say that consequentialism is to provide a direct answer to all ethical questions.

---

[43]Williams (1985), 17

## CHAPTER 6: Reasons

Thus far, I have argued that the most plausible form of consequentialism is one which asserts GC and GC', but which denies MC. In this chapter, I consider the possibility that the question of the best thing we can do, and that of what there is most reason for us to do, might come apart. In other words, someone might accept GC, but deny that we thereby have most reason to do whatever will bring about the best consequences. I will defend GC', however. In the opening section, I give an explanation of what I understand by the idea of a reason. I then go on to consider two connected objections to GC'. The first objection is that my account fails to take account of agent-relative, egoistic or partial reasons. The second is that anyone who accepts GC' will almost certainly be systematically irrational. I show why neither of these objections should trouble the defender of GC'.

### SECTION 1: Reasons for Action

Firstly then, I want to say a little about what I intend by the notion of a reason for action that occurs in GC'. I take it that the two questions which GC and GC' respectively purport to answer- 'What is the best thing we can do?' and 'What is there most reason for me to do?'- are very closely linked.

We might be tempted to ask why the defender of GC should complement his theory with GC'. Might not talk in terms of reasons talk be superfluous? Why not limit ourselves to claims about what is best, and about what is better and worse? The answer, I think, is the notion of *explanation*. We use reasons to explain our claims about what is best, better or worse. Some courses of action are better than others in respect of certain considerations- these are the reasons for performing one action rather than another. We talk about the considerations that go towards making it the case that this is the best thing we can do. So, for example, when we judge that we ought not to kill an innocent man, why it is best not to kill him, we explain our judgement in terms of a list of considerations that count against performing the action. Such considerations include the suffering of the innocent man's wife, the orphaning of his children, striking fear into the local community, cutting short a flourishing life and so on- these are each *reasons* against committing the murder, considerations which explain why it is best to refrain from committing the murder. Unless there are very substantial considerations which count in favour of performing the killing, listing these

considerations gives a full explanation as to why it is best not to commit the murder.[1] We should note that this simple notion of a reason for action makes no reference to the motivation of the agent. Hence, by speaking of a reason for action, I intend what Bernard Williams calls an *external* reason.[2]

This simple notion of a reason, as something that contributes to explanations of betterness, is a very natural and intuitive construal of the question, 'What is there reason to do?' Joseph Raz defines a reason for action as a consideration that counts in favour of an action.[3] This is one of several plausible ways of cashing out the concept without making any reference to the agent's motivation.

Let me take an example. The welfare of others is a good. So there is something good about performing a particular action which increases the welfare of others.
1. There is something good about doing the action.
2. There is something to be said for doing the action.
3. There is a consideration that counts in favour of the action.
4. There is reason to perform the action.

It seems to me that these statements all say pretty much the same thing. The first three statements seem all to be natural, intuitive interpretations of the claim that there is a reason to perform the action. Another way might be to say that a reason for an action is a consideration that makes the action more *choiceworthy* than it would be were the consideration absent. Thus I take GC and GC' to be very closely linked claims.

The consequentialist generally claims that consequences are the only relevant considerations in determining the value or choiceworthiness or favourability of an action, and so is claiming that all reasons for action are ultimately grounded in consequences (for example, consequences for human welfare). It is in this spirit then that I assert GC', using the term 'reason' in a way that is fully in touch with common usage.

---

[1]This is not to say however, that reasons are somehow prior to value. In saying that reasons explain our judgements of betterness and worseness, I simply mean that we point to particular values and disvalues (resulting from various courses of action) to explain our overall judgements about the all-things-considered value or disvalue of a particular action. These particular values and disvalues (or, if you prefer, 'bits' of value in the world) just are, or give rise to reasons, on this model. Thus I do not mean to commit the defender of GC' to the view that value somehow emerges out of reasons. In fact, I favour the opposite view, that reasons for action emerge out of value.

[2]See e.g. Williams (1995c)

[3]See e.g. Raz (1999)

SECTION 2: A Dualism of Practical Reason?

I now wish to consider an alternative framework to GC' which might seem attractive to those who take consequentialism seriously, a view according to which the questions of the best thing we can do and what we have most reason to do come apart. In particular, I will be discussing the position of Sidgwick and more recently, Roger Crisp, according to which there are not only impartial reasons, but also what are variously termed egoistic, prudential or agent-relative reasons. Crisp calls this position the dual-source view, since it judges an agent's practical reasons to emanate from two distinct sources: from the impartial good, and from what is good *for that agent*. In Shelly Kagan's language, Crisp's suggestion is that the explanation of why the balance of reasons does not always or usually support doing the act that will best promote the good is that there are special egoistic (or prudential or agent-relative) reasons which often outweigh impartial considerations in determining what we have most reason to do. This position might seem a very plausible refuge for those who are attracted by consequentialism, but are put off by its apparent demandingness. However, I will argue that such a view in fact runs into major problems, and will go on to defend GC' instead.

When I say that all reasons are impartial reasons or agent-neutral reasons, and deny the existence of prudential or agent-relative reasons, I am not of course denying that we have reasons grounded in our own welfare- we *do* have prudential reasons in this sense. It is just that these reasons do not have any weight, any rational weight over and above what weight they have in an impartial calculus. There is not more reason, all else being equal, for *you* to do X just because X promotes *your* welfare.

I will begin then by outlining the Sidgwickian position set out by Roger Crisp in his article, 'The Dualism of Practical Reason'. Crisp sees his dual-source view as riding a middle course between two extremes, egoism and moralism:
Egoism is defined as the view that 'all reasons for action are ultimately self-interested'.[4] Moralism, on the other hand, is the view that 'all reasons for action are moral reasons'.

It should be noted that moralism as Crisp defines it is Kagan's position, which goes naturally with MC, traditional consequentialism. I will raise some questions

---

[4]Crisp (1996), 53

about the appropriateness of this terminology later on. I do endorse however Crisp's definition of a reason, along the lines suggested by Raz, as the following: 'a reason for or against an action is a property of that action that *speaks in favour of or against it'*.[5]

Crisp's main aim is to argue against Kagan's position, 'moralism'. My aim in this section is to defend a position which avoids the extremes of both moralism and egoism, but without having to postulate a dualism of practical reason, which I hope to show is a very problematic view. In discussing whether reasons for action come from one source or two, Sidgwick, Kagan and Crisp deal primarily with consequentialist impartialism as the possible single-source view, and while other impartialist single-source views are of course possible, I shall follow the lead of these three writers[6], since I am attempting myself to defend such a consequentialist single-source view. So my aim is to argue against a position which takes seriously consequences impartially considered as a source of reasons, but which attempts to postulate a second source of reasons, namely the interests of the agent in question.

Crisp constructs what he calls the 'neutral argument' against Kagan- having conceded that the 'negative' and 'positive' arguments considered by Kagan himself are insufficient. This neutral argument 'allows that all values are assessable from the objective standpoint, and so avoids such failure. But *values* are to be sharply distinguished, conceptually speaking, from *reasons*.' Here we can already see a possible gap opening up between GC and GC'. And further, he says, 'even if all *values* are to be assessed from the objective point of view, it may be that there are certain agent-relative *reasons* which cannot be captured from or grounded in that point of view.'[7] Crisp thereby hopes to avoid Kagan's extremism, but without having to deny 'that there are no subjective or personal values'.

Crisp concludes that 'one cannot ground [moral] options by appeal to the cost to the agent understood in terms of agent-relative value', but instead by appeal to agent-relative *reasons*. Against Crisp, I wish to ground options by appeal to cost to the agent, but *not* by saying that the cost *outweighs* what is at stake in either value or in reason-giving force. Rather my suggestion is that the structure of the

---

[5]Ibid., 56

[6]Crisp himself does not explicitly defend a claim like GC, before going on to attack GC', but these are clearly the sorts of positions he has in mind to defend and attack respectively.

[7]Ibid., 61

concept of obligation is such that we are not necessarily *obliged* to do that which there is most reason to do.

Crisp offers an example, similar to my example involving David and George in Chapter 2, in order to illustrate his position. But I think he draws the wrong conclusions from it. There are three possible lives open to Anna[8]:
A: A life in which she has flourishing personal relationships, plenty of leisure and a lucrative career, and does nothing for charity.
B: A life in which she has all of these above goods, though not to the same extent because of the fairly large amounts of time and money she donates to Oxfam.
C: A life in which she has none of the above apart from the career, which she pursues solely in order to give all her spare money to Oxfam.

From this Crisp concludes, that 'The reason to promote the good... would support Anna's living life C rather than lives A or B. One might plausibly think, however, that she has stronger reason to live life B.'[9] This, Crisp argues, is because of competing agent-relative reasons.

My argument is to the effect that, contrary to some intuitions we might have, Anna has most reason to live life C. I will argue against Crisp on two main grounds. Firstly, such a Sidgwickian position seems to undermine a direct explanatory relationship between value and reasons, which I take to be a major advantage of consequentialist theories of reasons for action. Secondly, if we postulate these two types of reasons, we seem committed to saying that there is some balance between the two, and hence always a determinate answer as to what we have most (overall) reason to do. If this is correct however, then we will be left giving a very unintuitive account of someone who strikes an apparently 'healthy' balance between the two types of reason, but then goes on to do more overall good (for example, brings about better consequences, impartially considered) at her own expense. Are we to say that this person is now doing something which she has less reason to do? That she is less rational than she was before?

*FIRST OBJECTION*

I take it to be a major advantage of impartialist consequentialism that it promises a direct explanatory relationship between value and reasons. Consequentialism

---

[8]Ibid., 62
[9]Ibid., 62

seems to have the tools to offer a less mysterious account of where reasons come from, of how they are generated, and, crucially, the degree to which we have reason to do things.

There are two points to notice about this. Firstly, I think common sense certainly holds that there is a very close connection between value and reasons. Crisp himself acknowledges this: 'Now reasons and values are clearly closely connected. When we are enquiring about the reasons for doing a certain action, we are often satisfied with a reference to a value that grounds it.'[10] Secondly, if Sidgwick and Crisp think a balance must be struck in determining what we have most reason to do, how are we to go about determining, in any way that is not completely *ad hoc*, where this balance is? A closer connection seems more promising for explaining where reasons for action come from, if we think they vary, subject to certain conditions perhaps, according to the amount of value, perhaps to be cashed out in terms of human welfare, at stake. Crisp's view seems to point towards some sort of *ad hoc* weighting whereby the amount of reason we have floats free of the value at stake. It is not clear what sort of thing might determine how the two variables are to be weighed against each other in a dual-source view, and it is certainly incumbent on the dual-source theorist to provide such an account.

For the defender of GC however, betterness in action just is reasonableness in action; the better action is the more reasonable one. So insofar as we accept that value itself is unmysterious, so reasons for action are unmysterious.

*SECOND OBJECTION*

There is, however, an even more serious objection to the dual-source view. In the Anna example, it seems that if Crisp's conclusion is correct, then the agent-relative reasons must *outweigh* the impartial reasons in determining what we have most reason to do.

Now there is a path which might be taken which argues that there is not necessarily a determinate answer as to what we have most reason to do. One strand of practical reason demands that we choose B and another that we do C, and the two simply cannot be reconciled. This line of thought is closer to that of Sidgwick. On this model, practical reason does not just have two sources, but is

---

[10]Ibid., 56

in fact dualistic in nature in a stronger sense- it yields two different conclusions, almost as if there were two versions of Practical Reason. Crisp himself does not take this path, but Sidgwick, at times at least, holds this sort of position. However, he is deeply dissatisfied about having to do so, and in several places expresses concern that his project has failed, since he has been unable to establish what reason bids us to do. He is haunted by the idea that 'Practical Reason is divided against itself'.[11] Schneewind tells us, 'What he [Sidgwick] finds at the end is that because of one such basic and undeniable fact about human life, practical reason inevitably makes contradictory demands on action. If this is not a formal contradiction within reason itself, its bearing on Sidgwick's real hope for philosophical ethics is sufficiently devastating to make it clear why he thinks his endeavour ends in failure.'[12]

If our enquiry into what practical reason commands ends up giving us two incompatible answers, then our enquiry is prima facie unsatisfactory. We should at least ask whether we cannot find a plausible alternative, unified account. If we end up with two incompatible answers, it is surely better to conclude that they are simply answers to two different questions. Perhaps we lack a strong enough independent grasp of what the question 'What is there most reason for me to do?' even means. Instead, we should maybe accept that there are the following two questions:
(i) What will best satisfy my interests? or 'What is the best thing for me to do, *for me*? (which is a question that each of us naturally has enormous interest in, and whose answers we are strongly motivated to act upon), and
(ii) What is the best thing I can do overall, taking everything and everyone's interests into account?

The dual source theorist seems to allow the coherence of (something like) these two questions, but then demands that there must be some third question beyond these, some *substantive* notion of reason for action, above and beyond these two questions. An account of such a substantive conception of reason would be the sort of thing required to motivate a particular weighting between partial reasons, grounded in the interests of the agent, and impartial reasons. But it is not clear where the substance of such a third conception would lie. Thus we should not be tempted at all by Sidgwick's sometime conclusion that there are simply two irreconcilable answers to the one question.

---

[11]Sidgwick (1907), quoted at Schneewind (1977), 370
[12]Schneewind (1977), 374. See also Brink (1988); Frankena (1992); Mackie (1992); Skorupski (2001)

I shall focus now on Crisp's position, the dual-*source* view, which accepts that Practical Reason does yield a determinate answer as to what we have most reason to do. But I will suggest that his view leads to major problems.

Consider the person who does a lot of good for his family, community and the distant needy, and yet brings about less good than he is capable of. Kagan's extremism (MC) condemns his behaviour as morally wrong, simply in virtue of being suboptimal. The dual-source theorist argues that the balance of reasons is not simply determined by the question of which act will best promote the overall good. Rather, the balance of reasons, for any particular agent, is affected also by how a given act will impact on the agent's good, and this provides more or less weighty reasons against performing the optimal act where this will involve considerable sacrifice to the agent. So he can say that our apparently virtuous yet suboptimal agent may after all be doing what the balance of reasons supports. In this case, however, what are we to say about the agent if he then becomes even more committed to good causes, and does in fact start maximising the overall good? Should we say that he has *ceased* to do what the balance of reasons supports, now that he seems to be ignoring the 'due' weight of the agent-relative or egoistic reasons? If not, then where has the weight of the agent-relative reasons disappeared to? It seems very strange to say that now that he is bringing about even more good overall, he is acting in a way that he has *less* reason to. Surely the correct conclusion from this example is that the agent is morally admirable before and after, but that he has even more reason to bring about the most overall good that he can, than he does to bring about merely lots of good overall?

And something very similar can be said in the case of Anna. Crisp urges those unconvinced by his view to ask themselves the following question: 'is Anna being *unreasonable* in choosing life B over life C?'[13]

Perhaps not. But turn the question on its head. Would Anna be unreasonable if, moved by the horrific plight of those in the Third World suffering from disease and starvation, she chose life C over life B? Is she then being unreasonable? Surely not. The lesson to be learned from the answers to Crisp's question and to mine is that we do not condemn someone as unreasonable just for failing to do what they have most reason to do. Perhaps they are 'less than perfectly rational',

---

[13]Crisp (1996), 63

but not 'unreasonable' or 'irrational'. And it is extremely unintuitive to say that, moved by the terrible plight of Third World sufferers, Anna is less reasonable, less rational, is doing what she has less reason to do, if she opts for the more self-sacrificing life, making a huge, genuine difference to many real people's lives in poor countries. Perhaps either life, B or C, is reasonable, understandable and praiseworthy. But life C is even more choiceworthy, and there is even more reason to pursue it.

Crisp says, 'as Sidgwick recognised, it does not strike us as unreasonable that an agent give some priority to her own interests'.[14] However, likewise, it does not strike us as unreasonable if she sacrifices her own interests in favour of promoting the welfare of people living in terrible hardship.

Now Crisp does make a brief, and, I think, unsatisfactory comment in response to this sort of objection in a small footnote towards the end: 'Am I suggesting that living the sacrificial life (C) is *unreasonable*? Not, of course, in the sense that there is *no* reason for living that life: the promotion of the good does ground a reason to live life C. But living that life would be to ignore competing reasons grounded in the separateness of persons.'[15] Of course, choosing life C does not, on my model, mean ignoring any reasons or values, rather it involves correctly and very admirably putting in perspective the strength of those reasons which emerge from the agent's own welfare, in the face of the possibility of relieving great amounts of genuine suffering by real people.

Crisp is pushed by his unintuitive conclusion that life C is less reasonable than life B into a further problematic claim: 'Whether a person living such a life [C] is a candidate for blame or criticism is a separate question. In these cases, being unreasonable may even be praiseworthy.'[16] But surely it should be possible for an agent to be optimally reasonable and optimally praiseworthy?

I think where Crisp, like Kagan, goes wrong is in the structure of the relationship between reasons and obligations or permissions, options. And this is crucial. When we say that someone has a variety of morally permissible options, we cannot structure this as saying that each permissible action is what we have most reason to do. But Crisp does try to structure it like this. He says, 'Ultimately,

---

[14]Ibid., 66

[15]Ibid., 71

[16]Ibid., 71 footnote 26

when the dual-source theorist argues that there are options, she is arguing that there are agent-relative reasons which are stronger than the pro tanto reason to promote the good.'[17] But this model cannot be right. A sphere or situation in which there are options is one in which we are permitted (morally) to do, say, X or Y or Z. Now there will of course be more reason to do X or Y or Z: there will not be the same amount to be said in favour of each available action, in most cases where we have a range of permissible options. We regard ourselves as having morally permissible options almost all the time. However, that is not to say that all the options are equally choiceworthy or reasonable: there is not the same to be said for each permissible action in any given situation.

Again, there is a substantive difference between:
a. The balance of reasons counts in favour of X, and
b. X is morally required

It may be that we tend to sabotage our conception of what is valuable and of what we have most reason to do because it is very demanding and it seems that our language of morals and rationality tends to identify the optimal with the obligatory. It is this identification that leads us astray in judging what is valuable and what we have most reason to do. Crisp refers to the role of our emotions in discussing the difficulty of bringing about the most good, of being self-sacrificing. He says 'We have an evolutionary background and an emotional background and an emotional make-up which cannot be ignored in moral theory.'[18] This is quite right- the extent to which we naturally care about our own welfare and naturally favour ourselves and those close to us cannot be ignored in our conception of what our moral obligations are. However, it does not change either what it is best that we do, or what we have most reason to do.

I just want to end this section by acknowledging that my view might still sound unintuitive in some ways. Do we really not have any reason to favour ourselves? I want to try to deflate that thought by emphasising all the things that the defender of GC' can acknowledge.

1. We usually do favour ourselves.
2. It is natural for us to favour ourselves.
3. Favouring ourselves is reasonable in that sense can be made of it.

---

[17]Ibid., 72
[18]Ibid., 68

4. Favouring ourselves is not 'irrational'.
5. It is extremely difficult for us not to favour ourselves.
6. We cannot be expected not to favour ourselves to a fairly large degree.
7. We are not (in general) blameworthy for favouring ourselves.

However, beyond saying each of these, is there anything that actually counts in favour of favouring myself?[19] That is, something with normative force which does not apply to favouring others? We can *explain* why people do favour themselves. And we can *justify* favouring ourselves (to some degree) in the sense of showing ourselves not to blameworthy. But we *can't* justify it in the sense of showing it to be rationally optimal.

In Crisp's favoured formulation, a reason is something that 'speaks in favour of an action'. Of course the fact that it is I who will benefit rather than someone else will speak to me in the sense that it will in fact hold sway with me, and is hugely likely to motivate me. But to say that is not to say that it speaks *rationally* in favour of being partial to myself.

I want to finish by emphasising again how my position shares Crisp's ambition of riding a middle course between egoism and moralism, even though I reject his mixed account of the source of reasons. His definition of moralism was the claim that 'all reasons are moral reasons'.[20] I reject this claim. For a start, reasons generated by my own welfare are just as much reasons as those generated by the welfare of others. However, we would be very reluctant to call reasons generated by one's own welfare 'moral reasons'.

More importantly, however, a central part of my project is to detach talk of what generates reasons for action from any automatic implications of blame or guilt, i.e. from that which is characteristically moral. I am trying to divorce the questions of, on the one hand, what we have reason to do, and choiceworthiness, from, on the other hand, questions of moral wrongness and moral obligation. And my general claim is that consequentialism ought to be understood as a theory of the former, and not- as it generally has been- of the latter. So instead of

---

[19]In other words, to employ the terminology from Chapter 3, is there anything which counts rationally in favour of showing deep partiality, rather than shallow partiality? Obviously, it will very often be the case that the best thing I can do is to look after my own interests first simply because this is also the best way to maximise the overall good. The question in hand, however, is whether there is any extra reason to be deeply partial.

[20]Ibid., 53

claiming that all reasons are moral reasons, my claim is that all reasons for action are impartial or agent-neutral.

My picture becomes much more intuitive, I think, once we start to see the optimal standards of practical reason as an ideal, a guide, something to be aspired to perhaps, and which, through our cognitive imperfections and imperfections of will, effort and striving, we do not manage to reach, virtually all of the time, instead of complacently seeing optimally rational action as something that most people achieve most of the time.

## SECTION 3: Irrationality

I now turn to an objection to GC' which might be called the objection from irrationality, which claims that someone who accepts GC' will almost certainly be systematically practically irrational. The suggestion is that his normative judgements are so radically out of synch with his behaviour that he is irrational. My aims will be to argue for the following:

1. The objection from irrationality against my favoured form of consequentialism is not a good one.
2. On the conception of practical irrationality presupposed by the objection, many or even most of us are practically irrational much or even most of the time, and so this conception of irrationality is misleading.

## PRACTICAL IRRATIONALITY

My discussion will focus on one particular form of irrationality. I will use, as a model for the argument, T.M. Scanlon's discussion of rationality and irrationality in the first chapter of his book, *What We Owe to Each Other*. There are different forms of irrationality, which we can divide into three main types, corresponding to what might be regarded as three distinct types of reasons that people have- reasons to believe, reasons to feel and reasons to act. My focus is going to be on reasons to act. Scanlon says, 'Irrationality in the clearest sense occurs when a person's attitudes fail to conform to his or her own judgements: when, for example, a person continues to believe something... even though he or she judges there to be good reason for rejecting it, or when a person fails to form and act on an intention to do something even though he or she judges there to

be overwhelmingly good reason to do it.'[21]

So, corresponding to each of the three types of reason for action, we can give examples of this kind of irrationality, along the lines of Scanlon's model:
1. *Reasons to feel*: A judges there is no reason to feel afraid of spiders. Nevertheless, every time he sees one on the floor, he is terrified.
Therefore, A is irrational.
2. *Reasons to believe*: B judges that the reasons against believing in the existence of God hugely outweigh the reasons in favour. Nevertheless, he continues believing in God.
Therefore, B is irrational.
3. *Reasons to act*: C is in the pub, and judges that he has most reason to go home, since he has a tutorial at nine o'clock the next day. Nevertheless, he stays in the pub and has another three pints.
Therefore, C is irrational.

I am going to concentrate on the third of these ways of being irrational, i.e. by failing to do, or intend to do, what you judge there is reason or most reason to do. We should note from the start that there are perhaps other ways even of being *practically* irrational, so it is not a necessary condition of your being practically irrational that you fail to do what you judge you have most reason to do. For example, perhaps you can be practically irrational just in respect of the judgements you make about what you have reason (or most reason) to do.

## THE OBJECTION

I will begin with an intuitive characterisation of the irrationality objection to GC', then go on to delineate exactly what forms an objection from irrationality might take. A first shot at this sort of objection might run as follows. Consequentialists say that what we have most reason to do is bring about the best consequences we can. Given the state of the world, this seems to require a hugely self-sacrificial way of life. With the number of people starving in the world, our best bet for maximising human welfare is probably to give up our money and our spare time to improving the welfare of the world's poorest people, campaigning for trade justice and so on. Most consequentialists are in fact not maximising the good; they are not giving up 80% of their salary and 6 nights a week to good causes.

---

[21] Scanlon (1998), 25

Most consequentialists are systematically and repeatedly failing to maximise the good with their actions; failing to do what they claim they have most reason to do. Now this is perhaps not terribly surprising in itself, since it is very difficult, very demanding and involves great sacrifice to do the most one can for the world's poorest. It involves giving up most of one's hobbies, personal projects, spare time, and probably jeopardising to a large degree one's closest personal relationships.

However, I think that what those who object to my form of consequentialism particularly balk at is precisely the conjunction of GC or GC' with the rejection of MC. It is the thought that such a consequentialist says we have most reason to live like so, fails to live like so, and rests easy, and does not necessarily feel bad about failing to live like so. In rejecting MC, we say that we have no moral obligation to do what GC' says is optimal, that we need not feel guilty about failing to do it, and that others would not be right in blaming us simply for failing to do what we take to be optimal. There seems to be a suggestion, in asserting GC and GC', but denying MC, that we can be quite 'relaxed' about our failure to act in the way that we take to be optimal. But if we are failing to form and act on an intention to do what we judge there to be most reason to do, and making little effort to remedy this gap between our judgements and our motivation or intention or action, are we not just irrational?

The situation might be thought to be analogous to someone who has incompatible beliefs. It is not necessarily a sign of irrationality that a person has incompatible beliefs. What disturbs us is if someone rests easy about having incompatible beliefs, if he does not seek to revise one or other and make his beliefs compatible again. It is this sort of attitude that might lead us to describe someone as irrational. Likewise in the case of my consequentialist, the objectors reject the idea that someone who says he has most reason to do X, but 'settles for less' is rational. So the objection, I take it, is that on my view- in asserting GC' and denying MC- people who believe the theory will almost certainly be irrational.

Now, what I have called the irrationality objection might be better seen as a family of objections of different strengths. The objection can come in different forms:

*Failure to Act*

1. X is irrational if he fails to do what he has most reason to do.
2. X is irrational if he fails to do what he judges he has most reason to do.
3. X is irrational if he systematically and repeatedly fails to do what he judges he has most reason to do.

*Failure to be Motivated*

4. X is irrational if he isn't motivated (at all) to do what he (sincerely) judges he has some reason to do.

Objection 1 is implausible as it stands. In order to show someone to be practically irrational, it is not sufficient to show that they have acted suboptimally. For example, we do not accuse someone of being irrational simply because they have false beliefs about what is the best thing to do. So for example, if I desire a glass of water, I may drink what is in this glass believing it to be water. However, even if it is petrol (so in fact I had no reason to drink the stuff in the glass), I am not irrational for doing so.

As for Objection 2, it cannot be a problem for my theory that someone who believes it intermittently performs irrational acts, since we all perform irrational acts every so often. To use one of Scanlon's examples, we often judge that we have good reason to floss our teeth, but nevertheless fail to do it. Rational people can be irrational, Scanlon says[22], and this is not a problem; everyone is irrational from time to time, no matter what their judgements about what there is most reason to do. So the charge against my consequentialist cannot be based on mere intermittent irrational acts.

Objection 3, however, I think is where the really focus of the attack lies. The apparent problem for my consequentialist is that he does not just intermittently fall short of what he judges he has most reason to do; rather he systematically and continuously falls short of what he takes to be optimal. If 3 is true, then it seems that almost everyone who believes something like GC', and has a remotely plausible view of what good consequences consist in, is going to be irrational, since we almost all fall well short of acting optimally to promote the welfare of all. I will return to this presently.

I just want briefly to mention Objection 4, which might readily be confused with the above objections. One might think that there is a problem for my view about

---

[22]Ibid., 25

motivation, rather than action. This is not a special problem for a consequentialist, I think. Presumably, people who sincerely judge that what they have most reason to do, that the best thing they could do with their lives is sacrifice most of their time and money to helping those less well off than ourselves, are motivated to some degree to act on this judgement- i.e. if it cost them nothing, no time, no money, no effort, no conflict with other pursuits, they would choose that. It would be very odd for someone to assert sincerely that they judged that there was most reason to do X, and still lack even that level of motivation. If they were to appear to be in such a situation, I am not sure whether we would call them irrational or simply question whether they were indeed sincere in asserting that there was most reason to do X. But in reality, of course, the fact is that it is the tremendous level of effort, sacrifice, dropping our leisure, hobbies, relationships, and so on that generally deters us from doing what we take there to be most reason to do. Rather than lacking any motivation to do the altruistic thing, most people simply have many weighty countervailing motivations.

## AGAINST THE IRRATIONALITY OBJECTION

So I take it that the most plausible form of the objection from irrationality is Objection 3. The first thing to point out is that the objection does not apply simply to consequentialism, but to most, or at least many, common sense views. If the construal of practical irrationality asserted by 3 is correct, then the implications that the objectors claim apply to consequentialism, will also apply to anyone who holds the following set of common sense views:

a. Human welfare plays a large part in determining what we have most reason to do.

b. We do not have most reason to favour ourselves over others (to the great degree required).

c. We can make a significant difference to the welfare of those much worse-off than ourselves.

Furthermore, common sense tells us that we are not *morally obliged* to do what we take to be optimal, and so would reject a non-consequentialist equivalent of MC which claimed that we are morally obliged to do what we have most reason to do, where human welfare plays a large part in determining what we have most reason to do. Most people, I think, hold, or are committed to holding, these views, and yet fail to live a hugely self-sacrificial life promoting the welfare of the

worst-off. If this is correct, then the so-called objection from irrationality applies to mainstream views apart from GC'.

Needless to say, the second view that I have suggested might be a common sense view is somewhat more controversial than the other two, at least at first glance. But I hope to have shown in the last section, in discussing Crisp's view, that the alternative is not plausible. If we think that in determining what we have most reason to do, we need to strike a balance between favouring our own interests and promoting the interests of others, then the question emerges as to what we are to say if someone who has supposedly hit upon the 'balanced' life now steps up her altruistic efforts, sacrificing more of her personal interests, but making even more genuine difference to the lives of real very badly off people. Are we to say that now that she has strayed from the balanced way of life, she is being less rational than she was before, that she is now leading a life that she has less reason to lead- even though she is doing far more good overall than she was before? Again, this seems deeply implausible. Both lifestyles are extremely admirable, but the second one is even more admirable, even more choiceworthy, and there is even more reason to pursue that lifestyle. And I think this view genuinely is in keeping with common sense- we often hear people say things like, 'Of course it would be better to live like Mother Theresa, but that's just too hard'. If this is correct, then it appears that indeed it is natural and common sense to say that we do not have most *reason* to favour ourselves over others, even though it is of course quite natural for us to do so, and by no means blameworthy to do so, in most circumstances.

So what then are our options if we are to avoid the irrationality objection? What should be the response of the consequentialist and- if what I have just said is correct- of common sense to this apparent objection?

Our first option is to reject one of the three tenets of common sense that I have listed. The first seems difficult to deny. The second tenet is the one I have just defended- dropping this does not look promising.

What about the third tenet? We could try to argue that empirically, looking after our own interests is a better bet in promoting overall human welfare. But surely it is implausible to say that we cannot do anything much to help anyone anywhere; even for people sceptical about the results of the work of charities in very poor countries, there are surely many other 'good bets' for making a big difference to the lives of those less well-off than ourselves- volunteer work in

the community, helping the disabled, visiting the elderly, campaigning for trade justice or for stronger governmental action on preserving the environment. It shows a real lack of imagination to say that our best bet at contributing to human welfare is by sitting in the pub, or watching TV, most nights of the week.

These tenets commit us to saying that we have most reason to live an extremely self-sacrificial life that in fact almost all of us systematically fail to achieve, and unless human nature changes dramatically, will continue to fail to achieve. So if we do not want to drop these tenets, it seems that we only have two remaining options in responding to the irrationality objection- either we reject this notion of irrationality or we just accept that we are irrational most of the time, but just say that this does not constitute an objection to our account of what we have most reason to do. Now my main concern here is simply to show that the irrationality objection is not a good objection against my view, so we should not drop any of the three tenets. So I am not committing myself to whether we would be better off rejecting the notion of irrationality that the objection relies upon or just accepting that we are irrational most of the time. What I will do is just make a few very brief comments about the merits of each option.

Firstly then, should we reject the notion of irrationality that is in play here? There is a lot of rhetorical force attached to the term 'irrationality' that seems totally out of place in a charge against someone who believes that he has most reason to live a very self-sacrificial life, but lives only a mildly altruistic and self-sacrificial life. It seems bizarre to say that someone who believes GC', but who 'only' devotes three nights a week and 50% of her income to good causes (when she could be doing more) is *irrational*. We can certainly make sense of her thoughts and behaviour. She is certainly not an especially deficient rational agent in any way. The charge of irrationality seems wholly out of place in such a case. The worst we can say of her is that she is less than perfectly rational, which is a charge we should all be able to live with quite happily. So perhaps we should just deny that someone is irrational if she systematically and repeatedly fails to do what she judges she has most reason to do.

Even if someone were to insist that yes, you simply are irrational if you fail to do what you judge you have most reason to do, surely we should simply accept our irrationality, rather than abandon our conception of what we have most reason to do. Simply because we are not sufficiently motivated to do what we take to be optimal, there is no reason to abandon our view of what *is* optimal. For example, perhaps a majority of British people think they should live more healthily. Yet

they systematically and repeatedly fail to do what they take themselves to have most reason to do, presumably through simple weakness of will. What they should not do, in the light of this gap between judgement and motivation is conclude that in fact they do not have most reason to live more healthily. And yet that seems to be an analogous move to the one that the objectors to consequentialism are suggesting- they propose that because there is a huge gap between the judgement of the consequentialist and what he is motivated to do, he should abandon his theory. Now, of course, the ideal in the two situations is that the agent should be more strong-willed, that he should strive to bring his actions in line with his judgement. We should all live more healthily, and we should all be more altruistic. But given that we systematically and continuously fail to do what we take to be optimal, we should certainly not abandon what we take to be worth doing. We should not deceive ourselves about any of the three tenets. This is a mistake of the kind that, as I suggested in Chapter 4, Michael Stocker makes in 'The Schizophrenia of Modern Ethical Theories'. When faced with a gap between our judgement of what is optimal and our motivation, we should not revise our judgement. Rather we should just accept that we are less than perfectly rational. Acting suboptimally is simply human nature, and it is no black mark against a theory of what we have reason to do that people who assert it usually act suboptimally by the lights of that theory.

# CHAPTER 7: Consequentialism and Morality

## SECTION 1: The Challenge

Thus far, I have defended a form of consequentialism which denies that we have a moral obligation to bring about the best consequences we can, but which asserts that the best thing we can do, and what we have most reason to do, is bring about the best consequences we can- i.e. which asserts GC and GC', but denies MC[1]. Of course, in defending GC and GC', I do not simply intend my consequentialism to be a theory about the best available act, but about the comparative value of all available acts. So in any given situation, if action A, B and C bring about decreasing levels of good consequences, then, by the lights of my form of consequentialism, A is better than B, which is better than C. Again, when I say 'better than', I mean more choiceworthy; there is more reason to do A than to do B, and more reason to do B than to do C.[2]

My main aim in this final chapter is to defend this Scalar Consequentialism against a family of common objections, along the lines that it does not seem to be, or to imply, a moral theory, and is therefore unsatisfactory or deficient in some respect. Even if GC and GC' are plausible, and if I have succeeded in showing that MC is implausible (on the basis that it is unrealistically demanding), it seems that a number of questions about morality still remain. It is open for someone to question my favoured form of consequentialism by issuing one of the following challenges:

(1) 'Haven't you just avoided the demandingness objection by failing to give *any* account of moral demands at all?'
(2) 'Yes, GC is pretty plausible, but it doesn't amount to a moral theory, does it?'
(3) 'When faced with a terrible act, e.g. a case of torture, it isn't enough simply to say, 'well, he could have done better in that situation'. Limiting

[1]Further, I reject the softer claims that we always have a moral obligation to do what we believe will bring about the best consequences, or what we have good reason to believe will bring about the best consequences.
[2]Also, I intend GC as an evaluation not just of our acts, but also of our behaviour, practices and so on. So the best practices we can adopt are those which bring about the best available consequences.

ourselves to saying that seems to leave out so much of our ethical views. Specifically, we want to say that not only could the agent have done better, but he acted *wrongly*.'[3]

My aim thus far has not been to outline a moral theory, but rather to argue simply that consequentialism is not a plausible theory of moral obligation, but is a plausible theory of value, choiceworthiness and what we have reason to do, and that the most plausible objections to consequentialism only apply to MC, not to GC and GC'. Nevertheless it is incumbent on a scalar consequentialist to say a bit about how we might complement our theory of value or choiceworthiness. I have argued that consequentialism should not be regarded, as it often has been in the past, by both defenders and objectors, as an 'ethical theory of everything'. It seems fair to ask then how GC relates to the rest of the ethical realm.

My project might be accused of being unambitious if it fails to offer what is often grandly termed 'a moral theory'. This sort of accusation is a mistake however. The sorts of worries suggested by the three challenges above rather seem to force a scalar consequentialist into doing one of the following: (a) offer an account of what moral obligations we do have (a moral theory); (b) argue for an error theory about moral obligations; or (c) explain why a scalar consequentialist need not offer a 'moral theory'. My main strategy will be to take up the third option, and explain the reservations I have about offering something that might be called a moral theory. So what I hope to do is give a rough outline of how I think my favoured form of consequentialism- one that limits itself to being a theory of value and of what reasons for action we have- relates to other types of ethical claims, but without drawing up a systematic moral theory. In the course of doing this, I will explain just why I think drawing up such a systematic moral theory is unnecessary and, in fact, ill-advised.

---

[3]Tim Mulgan, for example, makes this kind of objection to Slote's Scalar Consequentialism. Mulgan (2001), 143. Mulgan's criticism is of a consequentialism which affirms something like my GC, but then remains silent on questions of moral wrongness. As will become clear in what follows, my position is that the onus is on a scalar consequentialist to at least 'cover the ground' of morality, even if he doesn't offer a fully-fledged, distinctively consequentialist, moral theory.

SECTION 2: The Consequentialist Can Say What the Rest of Us Can Say

My first point in fending off the sort of challenge I have outlined is that there is no reason why the consequentialist cannot make all the same sorts of ethical judgements that common sense makes, and further, no reason why he even cannot say what the rest of us want to say about cases of moral wrongness and moral obligation. This means that the third part of the challenge misses its mark. The challenge here is that the scalar consequentialist, when faced with a case of torture, for example, is reduced to saying simply that the agent could have acted better than he did. While the consequentialist will say this, it certainly is not the only thing he will say. He can firstly obviously say that this was a very, very bad act to perform, one of the worst available, he could have acted much, much better. But more than this, there is no reason why, just because GC is couched in terms of good and bad, better and worse, he cannot avail himself of the more expressive evaluative language that the rest of us use; there is no reason why he cannot, for example, describe an act of torture as horrific, cruel, repulsive or heinous.

It is a mistake to think that, in giving a theory of A, we need to eschew all talk in terms of B. In giving a theory couched in terms of better and worse, what we have more or less reason to do, we do not forfeit the ability to talk in other ethical terms about the same situation. Again, consequentialism is not an ethical theory of everything. Our ethical language has the potential to be very rich, and what Bernard Williams calls the 'thick' ethical concepts used in everyday ethical talk can say far more than talk in terms of either betterness or indeed wrongness.

The implication of the challenge against Scalar Consequentialism is that in asserting GC, we have left out too much. But this is obvious, and cannot count as a criticism if GC, or Scalar Consequentialism, is not meant to say everything we want to say about a particular ethical situation. Rather, it is meant solely to make a claim about what was the best thing to be done in the situation, and about the value or choiceworthiness of different courses of action in the situation.

In discussing a case of torture, we of course do not rest easy simply saying that the agent could have acted much better than he did in the particular situation, or even to say that it was a terrible thing to do. In fact, in

everyday ethical talk, it would be strange (though true) to make such a comment. Rather, we would go on to comment on what we thought of the agent, describing him as wicked, sick-minded, or the like. It is quite open to a consequentialist to say these sorts of things as well. Not only do we want to say that it was a terrible thing to do, since the consequences- the effect on the torture victim- were awful, but we also want to comment on the viciousness of the perpetrator.

However, in responding to the apparent challenge, I think the consequentialist can clearly go one step further. Even though my consequentialist rejects MC, the claim that we are morally obliged to bring about the best consequences, there is no reason why he cannot avail himself of moral language, of the concepts of moral wrongness and moral obligation. The scalar consequentialist can happily say that lying is usually wrong, that we are obliged to keep our promises and so on. And in the case of torture, we can say that torture is wrong in general, and that it is wrong in this case. There are clear-cut cases of moral wrongness and moral obligation (for example, almost all cases of torture), and the scalar consequentialist loses nothing by using moral terms to describe such cases.

What then is the relation between choiceworthiness talk and moral wrongness talk? Are they completely unrelated? Of course not. For the consequentialist, if we treat an action-type as morally wrong, part of the explanation for this is that it generally has bad consequences. So if the agent accurately appraises the situation, and is fairly sensitive to the importance of the relevant consequences, for example for human welfare, he will be blameworthy if he ignores these, and, for example, rides roughshod over the welfare of others. Of course, when we say that an action was morally wrong, we usually intend to say more than that it wasn't the best thing to do in the circumstances. We usually intend also that the agent was blameworthy in performing it. But there is nothing to prevent the consequentialist complementing his judgement that an act was unchoiceworthy with a further judgement that the agent was blameworthy (his intentions, efforts and motivation were bad). So the consequentialist can *cover all the ground* that is covered by our moral judgements.

The relation between judgements of choiceworthiness and specifically moral judgements can be seen by considering the following case. Imagine a

situation where I come home to find my friend, who has been depressed recently, lying on the sofa with an empty bottle of pills by him, and I'm unable to wake him. I have reason to believe he has attempted an overdose.[4] Now, in fact, my friend is just in a deep sleep after a long day at work. Now many people, I think, would say that I have a moral obligation to phone for an ambulance, that it would be morally wrong of me not to. For the consequentialist, clearly the *best* thing I can do (in the 'objective' sense I outlined at the beginning of Chapter 5) in the situation is in fact not to phone the ambulance, since my friend is in fact in no danger, and calling out an ambulance will in fact at best embarrass all concerned, and at worst prevent an ambulance from going to help someone else who is in genuine danger. So there seems to be a pressure to say that, in some circumstances, for example those where ignorance or false belief is involved, it can be morally wrong to do what is in fact the best thing to do, what we have most reason to do. The basis of our judgement that failure to phone the ambulance is wrong is that we think I would be blameworthy or 'unvirtuous' if I failed to phone. I would be a very poor kind of friend, and indeed, a poor human being, if I failed to call an ambulance when I judged that my friend was in mortal danger. The question 'Ought I to phone the ambulance in this situation?' might appear to be an awkward one for the consequentialist. This is because, on the one hand, there is in fact nothing to be said for phoning the ambulance, and much to be said against it, while on the other hand, we think ill of, or feel blame towards, the agent who, ignorant of the fact that his friend is in no danger, fails to phone the ambulance. But separating out the two separate 'ought' questions- the ought of moral obligation, and the ought of choiceworthiness- removes any confusion.

We must remember that moral obligation is (on most construals, at least) connected not with the actual reasons there are to perform a particular action, but with the warrantable reasons. In other words, whether someone has acted morally wrongly in doing X is determined not by whether he did what in fact there was reason to do, but rather by whether he did what he had reason to believe there was reason to do. So we can simply say that I *am* morally obliged (I morally ought) to phone the ambulance, since I was warranted in believing that that was what I had most reason to do. After all, I would seem to be blameworthy for failing to

---

[4]This example is courtesy of Simon Robertson and John Skorupski.

phone the ambulance, and there is, as I have said before, a very close link between moral obligation and blameworthiness. But this is compatible with the consequentialist claim that what I ought to do (what it is best for me to do) is in fact not to phone for the ambulance.

The tight connection between the two types of ought judgements is especially evident in contexts of advice. Our application of the concept of moral rightness and wrongness is not limited simply to retrospective judgements evaluating an agent's worth in performing a particular act, but rather it also has a prospective function, recommending some acts, and warding us off others. Thus, we often say, 'You can't do that; it would be morally wrong'. I take it that it is part of the concept of moral wrongness when used prospectively that to regard certain actions as morally wrong is in part to regard them as less choiceworthy than actions one takes to be morally permissible. The point is that in contexts of advice, to say that act A is morally wrong while act B is morally right is, in part, to recommend B over A, to say that there is more reason to do B than A. But this tight connection should not mislead us into thinking that the consequentialist is doing something unusual in offering an account of choiceworthiness and a separate account of blameworthiness. In the case I have just considered, we can say everything we want to say about the situation by asserting that the (objectively) choiceworthy thing to do is not to phone the ambulance, since there is in fact nothing to be said for doing that, and in fact a lot to be said for not doing it; but that, given the agent's ignorance of the facts, he would be blameworthy for failing to phone the ambulance.

Thus consequentialists can clearly cover the same ground that is covered by the two main strands in our moral talk- choiceworthiness and blameworthiness.

## SECTION 4: Shaping Moral Norms

The challenger to consequentialism we are considering says, 'Yes, I can see that this is better than that, but now tell me what are our moral obligations?' The consequentialist should react as any other cautious philosophical person should to such a general question, before giving an unambitious answer like, 'We are typically morally obliged to refrain from

harming others, and to helping out, to a reasonable level, those less well-off than ourselves'.

But I want to say more to counter this challenge to the consequentialist to offer a theory of our moral obligations. In order to do so, let me return to my discussion in Chapter 5 of Bernard Williams's claim that modern moral philosophy has typically been over-ambitious. He takes Sidgwick as his prime example of someone who makes this kind of mistake. One part of this mistake was, I suggested, to think that we can *derive the* correct set of moral rules from a single theoretical starting-point, such as the principle of utility. This is often accompanied by a tendency to project the norms and needs of one's own society onto mankind wherever it exists in time and space. If the challenger to scalar consequentialism is expecting us to give a systematic account of what moral rules follow from GC, then he should be disappointed. GC will of course say that the optimal set of moral norms for any given society will be the one that results in the best consequences. But this does not get us very far in working out a moral code.

This is not to say however that GC has nothing much to say about our moral practices. Rather than attempting to derive our moral obligations from our theory of value, I think the role of GC in shaping our moral theory and practice should instead be to mould what we take to be our moral obligations by criticising particular parts of our theory and practice in the light of GC. So, we can criticise particular ethical judgements which prevail in a given society, such as the view that slavery is worthwhile or that racism is to be permitted. According to GC, if treating slavery or racism as unacceptable has better consequences, then a set of norms which treats them as such is a better set of norms (all else being equal). So not only can the consequentialist condemn slavery on the basis that it does more harm than good, but it can also endorse a set of norms which treats slavery as unacceptable.

So consequentialism can recommend changes to the norms we recognise-changes which would produce better consequences than the status quo. Two obvious examples concern the welfare of the starving poor and of our treatment of the natural environment. It is arguable that these two issues are the most pressing ethical issues facing humanity. And yet common sense moral thought is at best unclear on what to say about them, on the

extent of our duties in these areas. So while we all have strong intuitions about everyday cases of obligation, about how to treat those around us- our families, friends, neighbours and communities- our intuitions about the distant poor and the environment are less refined and less dominant in our thought. This is understandable, of course, since problems about how to interact with those around us are vivid to us in a way that the plight of the distant poor, and distant future people, who might suffer because of neglect of the environment today, are not.

Given the amount of human welfare at stake regarding each of these issues, consequentialism would recommend not just that it would be better if we spent more time and thought on addressing these issues, but further it might be better if we regarded apathy about the environment and the welfare of the distant poor as something shameful or blameworthy in a similar way to other things we take to be morally taboo.[5]

That the norms we endorse *can* change in the way I am suggesting is apparent from our changing moral attitudes to homosexuality, drink-driving, and perhaps smoking in public, or smoking in front of children. If these changes in norms result in better consequences, then they are endorsed by GC. These examples are intended to be possible instances of the sort of piecemeal criticism of our current practices that should constitute the interface between our consequentialist theory of value and our moral practice.

It may seem tempting to try to modify MC in various ways in order to bring it closer to our moral intuitions. In other words, rather than settling for consequentialism as a theory of value and using it piecemeal to correct our current practices and norms, we might attempt a project of turning consequentialism into something more like a full-blown moral theory. It seems to me that that project may be misconceived, and that consequentialism has a more modest, but still significant role to play in shaping our moral obligations.

---

[5]Furthermore, it is very difficult for someone, individually, to take these issues really seriously, on a consistent basis, *given* that those around her are not. If we all recognise a norm to do more for the distant poor or to preserve the environment, it becomes easier for us to make sacrifices, when those around us are also making sacrifices, and expect it of us. As things stand, it is often the case that others even expect us (in a non-normative sense) *not* to contribute seriously to aiding the distant poor and preserving the environment.

SECTION 5: The Hybrid Theory

Now I wish to turn to two theories to which do offer theories of moral obligation, and to which my account may be thought to bear significant resemblance, namely, Samuel Scheffler's hybrid theory, and Rule Consequentialism. I hope, in this section and the next, to further clarify my position by distinguishing it from these two views, while outlining points each shares with my view, and to suggest that my position is well-placed to avoid the difficulties faced by the hybrid theory and Rule Consequentialism respectively.

The title of the main work in which Scheffler sets out his hybrid view, *The Rejection of Consequentialism*, might imply that his view has little in common with my view, which is a thoroughgoing consequentialism.[6] However, Scheffler's 'rejection' is clearly rejection of what I have called MC, rather than GC. The hybrid view is one which claims that we are always permitted to maximise the good, but that we are not generally required to do so. Thus, Scheffler is rejecting MC, which claims that we are required to maximise the good. Although Scheffler takes himself to be rejecting consequentialism, one clear rationale for his claim that we are always permitted to maximise the good would be that maximising the good is always the most choiceworthy thing to do, what there is most reason for us to do. Scheffler does not explicitly endorse GC, but it is at least compatible with his hybrid view, and indeed seems to make sense of its claim that we are always morally permitted to maximise the good.

Like my position, the hybrid view is motivated by the thought that MC is implausibly demanding, and that in determining what our moral obligations are, we do not start with what would be optimal, but rather look at lives as they are lived, from the agent's point of view. Although Scheffler does not go into detail about the precise grounds for endorsing an agent-centred prerogative, he refers to it as, 'a structural feature whose incorporation into a moral conception embodies a rational strategy for taking account of personal independence, given one construal of the importance of that aspect of persons'.[7]

---

[6]Scheffler (1982); Scheffler (1992)
[7]Scheffler (1982), 67

While Scheffler's starting-point is similar to mine, his hybrid view *does* attempt to outline an explicit moral theory, one that purports to be less demanding than MC. In 'Prerogatives Without Restrictions', Scheffler outlines how a criterion for moral requirement might be constructed on the hybrid model, suggesting that we are permitted to give our own interests M times more weight than the interests of others.[8] Thus, he presents a moral theory which requires less self-sacrifice than MC, and is thus less demanding.

As I argued in Chapter 3, providing moral justification of one's actions is a matter of showing that one was *sufficiently* impartial, that one made *sufficient* sacrifices. Thus, it seems plausible to say that acting in a morally permissible way involves not giving *too much* weight to one's own interests. We might then, like Scheffler, think that there is some permissible degree to which we can favour ourselves over others. However, to say this is to stop short of saying that there is some single weighting applicable to all situations, which determines the extent to which we may favour ourselves.

One obvious worry about such a formula as proposed by the hybrid theory is that the level of weighting it recommends will appear arbitrary. It seems obvious to question whether there really is some determinate extra weight that we are morally permitted to give to our own welfare over that of others. And if there is such a determinate weight, it seems mysterious what exactly it is that determines the value of the weighting.

Also, the model seems specifically tailored to cases which involve the sacrifice of one's time and money. It may turn out to be better suited to cases of imperfect duty to the very badly off than to other cases of moral obligation.

A more pressing problem is that, in order to produce a moral theory which is not extremely demanding, the requisite weighting, M, would have to be pretty high. If we are permitted only to give our own welfare twice, or even ten times as much priority, then we may be obliged still to give up almost all of our time and money to helping the distant poor if we are in fact able to make a very significant difference to the welfare of some of the

---

[8]Scheffler (1992b)

poor. However, if the weighting is indeed much higher, if the value of M was, say, 100, then the formula will yield implausible results with regard to how we are permitted to treat those in our own communities. We would, for example, be morally permitted to break a promise so long as the benefit to me of doing so was at least 1/100 of the benefit to others of my keeping the promise. The most viable form of the hybrid theory then might be one in which the weighting is variable across different situations.[9] In local, everyday situations, we might be permitted to give far less extra weight to our own welfare. After all, we might say, the fact is that morality is a phenomenon which has evolved to regulate our conduct at a local, personal level, so we should expect that in trying to formalise an extension of morality from the local to the global, radically different levels of expectation must apply.

Making the value of M flexible in different situations may go some way to making our moral judgements on this model more plausible. One serious concern about such a strategy is that it seems unclear then what the order of determination in our ethical framework is. Our intuitions about whether there is a moral obligation appear to determine our judgement as to what the relevant weighting is in the situation. The story about the weighting may then seem to be eliminible. In discovering what our moral obligations are, our intuitions (guided by our natural sentiments of blame and guilt) do seem to play a central role. We do not naturally feel that blame is appropriate in direct proportion to the discrepancy between the welfare of the agent (given some special weighting) and the welfare of those affected. It is not the theory that is the driving force of our moral judgements, but our moral intuitions, guided by our natural sentiments. If our theory becomes so fine-grained as to fit neatly upon our intuitive judgements about when a moral obligation exists, then it seems to be in danger of losing its status as a general theory of moral obligation altogether.

Thus, while I think that Scheffler's motivations are well-placed, and that he correctly emphasises that judgements about moral requirements are sensitive to the costs to the agent, I am sceptical that any straightforward general moral theory can be given along the lines Scheffler proposes.

---

[9]For further problems facing a hybrid theory with a single variable, see Mulgan (2001), Chapter 6.

Like Scheffler, I favour a framework which endorses agent-centred options, but which rejects agent-centred constraints. The focus of my thesis has not been on the rejection of moral constraints on maximising the good. However, it should be noted that there are different ways a consequentialist might react when faced with cases where consequentialism permits (and recommends) actions forbidden by common sense morality. What I have in mind are cases where common moral intuition forbids us from maximising the good, for example, by punishing an innocent person. In insisting that, in the rare circumstances where we know that punishing the innocent man will in fact lead to the best consequences, the best thing to do is to punish the innocent man, the consequentialist is confronted with two options. Firstly, he might say that punishing the innocent man is morally wrong, but that we should do it anyway since that will result in the best consequences. Or alternatively, he might say that punishing the innocent man cannot in fact, contrary to our common sense moral intuitions, be wrong since this course of action is the one that will lead to the best consequences. My preference is for the second reaction, to see disagreements between consequentialist judgements and common sense moral judgements as genuine disagreements, both about which course of action is the most choiceworthy, and about whether a particular course of action is wrong. This is precisely because to call a prospective action wrong is, in part, to say that it is not choiceworthy, often that it is an especially bad (unchoiceworthy) thing to do. In other words there is conceptual connection between the (prospective) judgement that an action is wrong and the judgement that the action is not choiceworthy.

Thus, someone who knows that an action X, which is commonly judged morally wrong, is in fact the most choiceworthy thing to do does not generally say 'X is wrong, but you should do it'. Rather, he says, 'X can't be wrong, since this is what you should do'. Therefore, it seems that where consequentialism recommends something common sense morality calls 'wrong', it is better to see it as a dispute within morality, to ask 'Is it really wrong after all?'. In purporting to correct our common sense moral judgements, consequentialism need not be rejecting morality itself. Instead, it can argue that it cannot generally be wrong to do something you know will have the best consequences. Thus it is better to say we should change particular common sense moral judgements, rather than say that it is better to do something which is in fact morally wrong. A clash between

common sense intuitions and consequentialism is not best portrayed as somehow being a clash between morality itself and consequentialism.

SECTION 6: A Comparison with Rule Consequentialism

My preferred form of consequentialism might also be thought to bear significant resemblance to Rule Consequentialism, insofar as both purport to stand outside of our particular moral practices and evaluate particular moral codes in terms of the consequences they produce. In this section, I will make explicit the ways in which my theory comes apart from Rule Consequentialism.

Rule Consequentialism (hereafter RC) evaluates moral codes, or sets of moral rules, according to the consequences of these rules being accepted or obeyed. In Brad Hooker's terms, the 'ideal code' is the one whose acceptance has the best consequences. Thus far, GC is broadly in agreement. However, Rule Consequentialists go on to make the further claim that the moral wrongness of an act is to be determined by reference to whether it obeys this ideal code- an act is wrong if it violates the code whose acceptance would have the best available consequences.[10] My favoured form of consequentialism does not endorse this second claim.

A clear commonality between RC and my theory is the denial of MC. For Hooker, one of the central strengths of RC is just that it gives a much more plausible account of the demandingness of morality than Act Consequentialism. Hooker describes Act Consequentialism (AC) as the view that 'an act is morally right if and only if the act maximizes (actual or expected) value, impartially construed'.[11] However, in criticising AC, Hooker clearly has as his target a view upon which it is morally *required* to perform the act which maximises value; thus his target is what I have called MC, rather than GC. Hooker argues that RC yields far more plausible

[10]Hooker's full preferred formulation of Rule Consequentialism is: 'An act is wrong if it is forbidden by the code of rules whose internalization by the overwhelming majority of everyone everywhere in each new generation has maximum expected value in terms of well-being (with some priority for the worst off). The calculation of a code's expected value includes all costs of getting the code internalized. If in terms of expected value two or more codes are better than the rest but equal to one another, the one closest to conventional morality determines what acts are wrong.' Hooker (2000), 32

[11]Ibid., 142

moral demands (than a position like MC) concerning aid to the world's poor, since he thinks it will produce a rule requiring contributions of between 1 and 10 per cent of annual income from the relatively well off.

One central difference, however, between RC and the form of consequentialism I have outlined is in the concentration of the former on *rules*. In *Ideal Code, Real World*, Brad Hooker says 'A theory is consequentialist if and only if it assesses acts and/or rules (or motives, social codes, virtues, or ways of life) in terms solely of the production of agent-neutral value.'[12] My theory assesses the value of all the things Hooker mentions in just such terms, and is thus what Kagan calls 'Direct Consequentialism'.[13] The best character traits or motives for any individual to have are those which result in the best consequences, and the best way of life to adopt is the one which will result in the best consequences.

By contrast, RC limits itself to the assessment of rules or codes directly in terms of consequences. It then goes on to evaluate acts, not in terms of the consequences they produce, but in terms of their obedience to the preferred code. For a consequentialist with what Hooker calls 'an overall commitment to maximising the good', it may seem unclear what motivates the Rule Consequentialist to adopt this schema for rules, but reject it for acts and so on. RC seems to abandon a central claim in which much of the appeal of consequentialism is invested, namely the claim that it is always better to do A than B if doing A will have better consequences.

At this stage, it is useful to note two fundamentally different forms of RC- one which does endorse an overall commitment to maximising the good, and one which, like Hooker's, does not endorse such a commitment, but rather attempts to defend RC on the basis of its general ability to provide intuitive results about what our moral obligations are. Rule Consequentialists of the first kind, who attempt to give a consequentialist justification for their Rule Consequentialism are often faced with the objection that their theory simply collapses into Act Consequentialism. After all, they appear to claim that we should obey the rules whose acceptance will have the best consequences, on the basis that doing this

---

[12]Ibid., 110
[13]Kagan (2000)

(obeying these rules) will itself have the best consequences. The problem for such a form of RC emerges in the apparently common situation that obeying the ideal code will not, in *this* particular instance, result in the best available consequences.

Rule Consequentialists who offer a consequentialist justification for their theory seem to be committed to endorsing one of the following three claims:
(1) We should break the optimal rules in the ideal code whenever doing so will lead to the best consequences.
(2) The rules in the ideal code will be sufficiently fine-grained to always approve the action that has the best consequences.
(3) Obeying general rules will always lead to the best consequences.

The first two seem to collapse into direct consequentialism. On the first view, the central claim that we should do what is recommended by the ideal code becomes redundant when supplemented by the proviso that whenever obeying the ideal code is suboptimal, an exception is granted and we should do the directly optimal thing. On the second view, the notion of a rule seems empty, and again RC collapses into equivalence with the Direct Consequentialist claim that we should do whatever will bring about the best consequences. The third view is empirically implausible.

A view like Hooker's then seems the most plausible independent form of RC, one that is genuinely distinct from Direct Consequentialism. Hooker says, 'the best argument for rule-consequentialism is *not* that it derives from an over-arching commitment to maximise the good. The best argument for rule-consequentialism is that it does a better job than its rivals of matching and tying together our moral convictions, as well as offering us help with our moral disagreements and uncertainties.' (101) In explicitly abandoning the overall commitment to maximising the good, Hooker can portray the ideal code as one which 'will contain rules of limited number and limited complexity. Such a code would not be extensionally equivalent with act-consequentialism'.[14]

---

[14]Hooker (2000), 97

The cost of Hooker's view, of course, is that it loses what seemed like one of the main attractions of consequentialism, its simple claim that if A will produce more good than B, then there is more reason to do A. Against Hooker's view, the Direct Consequentialist will persist, 'Why should we perform action B if we *know* that doing A will produce more good? Indeed, doing B may produce no good at all, or *far* less than A, even though B is required by the ideal code, and thus morally required by the lights of Hooker's form of RC.

Again, the starting-point of RC seems very plausible. Hooker says, 'Rule-consequentialism holds that the code whose collective internalization has the best consequences is the ideal code.'[15] However, Hooker, and other rule consequentialists then go on to claim, more controversially, that moral wrongness is a matter of violating this ideal code. Hooker appeals to our moral intuitions as follows: '"What would be the consequences if everyone felt free to do that?" There is no denying the familiar appeal of this question. And there is no denying that rule-consequentialism is a natural interpretation of the question.'[16] It is surely correct that the question is very relevant to our ethical thought. To say this however is to stop short of expanding this question into a theory of moral wrongness. To see this, consider the following phrasing used by Hooker: 'Rule-consequentialism endorses prohibitions on...[certain]... kinds of acts, since on the whole the consequences, considered impartially, will be far better if such prohibitions are widely accepted.'[17]

There are two sorts of things that we might understand by the phrase 'endorsing prohibitions':
(1) To endorse a prohibition is to say that it would be good if that prohibition were accepted by most people.
(2) To endorse a prohibition is to say that we are in fact morally forbidden to act in a certain way.

And an analogous ambiguity exists with endorsing positive obligations, as opposed to negative prohibitions:
(1) To endorse an obligation is to say that it would be good if that obligation were accepted by most people.

---

[15]Ibid., 2
[16]Ibid., 5
[17]Ibid., 30

(2) To endorse an obligation is to say that we are in fact morally obliged to act in a certain way.

I take it that any consequentialist will endorse a positive obligation or a negative prohibition in the first sense, if the general acceptance of the obligation or prohibition will result in good consequences. However, to acclaim it as a good thing that a norm be recognised is to stop short of saying that an actual moral obligation or prohibition exists. This seems clear in the case of First World aid to the distant poor. It may well be a good thing if we were all to feel morally obliged to give, say, 10% of our income to aiding the Third World, but from this it does not follow that we are in fact obliged to give any particular amount.

Hooker suggests that RC will be plausible if it yields a moral requirement to give between 1% and 10% of our income to help the poor. However, we must ask what it is we are doing when we say that giving less than, say, 10% of our income is morally wrong. Certainly, it would likely be a good thing if we all gave 10% of our income to help the poor.[18] But saying that giving less than this is *wrong* is to say more than this. In giving an account of moral wrongness and of moral requirements, the central aim is to say what sorts of actions people are and are not blameworthy for performing. Are people in the real world, where we are not all philosophers, never mind rule-consequentialists, and where there is no well-entrenched convention or practice that we give a percentage of our income to help the distant poor, really blameworthy for failing to give 10%?

Perhaps if there was some lengthy public discourse in a particular society where it was discussed and established that we would all give 10%, in the style of an informal tax, then it would be blameworthy not to do so. But in the actual world, there is no such discourse (even though it might be good if there were), and so it seems there is no such determinate obligation. A society in which we semi-formally treated giving less than 10% to the poor as blameworthy would probably be a better society. But this is not to say than in the real world giving less than 10% is wrong, while giving more than 10% is acceptable. To this extent then, I think Rule Consequentialism

[18]I am not suggesting that an extra tax upon the income of those in the First World would be the best means to reducing world poverty. Reform of international trade rules and practices, and of power structures in the poor countries themselves, are, quite plausibly, at least as important in reducing poverty as the donation of money.

shares in the mistake of many consequentialists attempting to offer a moral theory proposing a way to derive more or less precise obligations to the poor.[19] There may simply be no determinate answer as to what our obligations 'really are'. I return to this in the next section.

Further problems result when we move from the simple claim that the best code is the one with the best consequences to a full-blown Rule Consequentialist theory of moral wrongness and moral obligation. Again, the Direct Consequentialist can agree that the best code is indeed the one that has the best consequences. However, this does not change what *I* have most reason to do in any particular situation, which, according to the Direct Consequentialist, is a matter of the actual consequences of the various courses of action available to me. From the perspective of Direct Consequentialism, the main weakness of RC is that it is insufficiently responsive to actual consequences. RC grounds moral requirements in the value of hypothetical consequences- it asks what the consequences would be were the majority to accept the relevant rule. Hooker considers, for example, a 'genius' who knows for sure that breaking the optimal simple rules on a particular occasion will result in the best consequences. He concludes that, 'The genius should follow the same rules as the rest. That is, the genius should follow the simple rules even if she could do somewhat more good by following more complicated rules.'[20]

But surely actual consequences are very relevant in determining what we ought to do. It may turn out in fact that in many circumstances, RC is less sensitive to actual consequences than common sense (non-consequentialist) moral thought or than most plausible forms of deontology. Most non-consequentialists assign substantial weight to the actual consequences of our actions in determining what we should do in many circumstances. The central disagreement between common sense morality and Direct Consequentialism is over the latter's claim that *only* consequences matter in determining what we should do, rather than in the claim that the consequences are generally relevant. Potentially, RC could recommend a course of action which has nothing going for it at all, and indeed produces significant harm, simply on the basis of what the

---

[19]Hooker himself says, 'Coming up with an intuitively plausible rule about aiding the needy is something that (as far as I am aware) no other consequentialist theory does.' Hooker (2000), 174

[20]Ibid., 86

consequences would be were the majority to accept a code which required that action. Hooker does include a disaster clause in his preferred form of RC, which allows us to violate the ideal code if 'disaster' would result from obeying it.[21] But this seems insufficient: the question of what we ought to do should be more sensitive to actual consequences in non-disaster circumstances.[22]

One of the strengths of RC is that it rejects the hugely demanding claim, MC. Hooker says 'a comfortably off agent is required by act-utilitarianism (and indeed by most versions of the broader theory, act-consequentialism) to repeatedly make enormous personal sacrifices for the sake of needy people with whom the agent has no special connection'.[23] And this provides him with motivation to reject act consequentialism. However, Hooker's rejection of act-consequentialism is of course not a rejection of GC, but of MC. I have argued in this thesis that such a rejection is well-motivated. However, one can still, as I have shown, deny that we are morally obliged to bring about the best consequences we can, while holding on to the following fundamental consequentialist claim:
If doing A has better consequences than doing B, then it is better (it is more choiceworthy, there is more reason) to do A than to do B.

In this way, my theory aims to retain the attractions of both RC and MC, while avoiding the main objections to each. In particular, within my framework of endorsing GC, but denying MC, we can do the following:
1. We can, as Rule Consequentialism does, assess our moral rules, codes and practices in terms of the consequences they produce, but at the same time avoid dropping the link between choiceworthiness and actual consequences.
2. We can, as Direct Consequentialism does, retain the attractive claim that if A will produce more good than B, then there is more reason to do A than B, but at the same time avoid the implausible claim (MC) that we are morally obliged to bring about optimal consequences.

In conclusion then, Rule Consequentialism, like the hybrid theory, begins from a plausible starting point, but arguably goes wrong in extending a genuine consequentialist insight into a fully-fledged moral theory.

---

[21]Ibid., 165
[22]For more objections to Rule Consequentialism, see Mulgan (2001), Chapter 3.
[23]Hooker (2000), 28

SECTION 7: Full-Blown Moral Theory

Finally, I shall offer some positive reasons why I think it is a mistake for the consequentialist to get caught up in a project of giving a full-blown moral theory, and why it is not a black mark against my Scalar Consequentialism that it does not attempt to do so.

Firstly, Scalar Consequentialism's focus on choiceworthiness or reasons for action has some distinct advantages over characteristically 'moral' talk. Scalar Consequentialism evaluates *all* our actions, behaviour, and practices, not just some particular area, siphoned off from the rest. If moral theory typically divides actions into two categories- the permissible and the forbidden- then it leaves a lot out in doing so. Amongst the permissible actions in a given situation, some will generally be better than others. Likewise, amongst the forbidden or morally wrong actions, some will be worse than others. To this extent then, consequentialist evaluations of different actions are much richer than narrowly moral ones, since consequentialism recognises important differences amongst different courses of action that are ignored by moral categories. 'Moral philosophy', and even 'ethics' can be misleading names, in that they draw artificial divisions which can lead us to overlook important continuities. Consequentialism affords us a viewpoint from which our ethical thought is continuous with the rest of our evaluation of human acts, behaviour, and practices. Studying the paradigmatically moral independently from the rest of our conduct has perhaps led to some of the dominant, but ill-conceived traditional problems of moral philosophy, such as the question 'Why be moral?', the posing of which implies that most of the time we have reason to pursue our own interests, but that we then need to come up with some explanation as to how we can have reason to perform morally obligatory actions. The study of the value of moral conduct, and even of other-regarding conduct, is not something to be done in isolation from the study of the value of human behaviour more generally.[24]

[24]There may nevertheless be some advantages to compartmentalisation involving the narrowly moral categories in our everyday practices. It is not useful for us to be thinking directly in terms of what is the best thing we can do, impartially considered, every single time we act, as I argued in Chapter 5. Rather, a healthy ethical practice is more likely to involve (1) occasional impartial reflection, and (2) specifically moral motivations kicking in when we are tempted to do something harmful. To some extent, it is inevitable, and desirable, that we siphon of the moral and 'get on with our lives', living by general moral rules and dispositions, rather than consciously aiming to act impartially at every turn.

A second reason why the consequentialist should not try to change his theory of value into a theory of moral obligation is that the moral categories themselves face a number of awkward problems. For example, traditional moral categories seem to get caught up in saying implausible things about the notion of moral luck- as Williams and Nagel argue in different ways. Nagel argues that the notion that we are only morally responsible for what is ultimately within our control leads to the disappearance of the sphere of moral responsibility.[25] Williams suggests, 'There is a pressure within [morality] to require a voluntariness that will be total and will cut through character and psychological or social determination, and allocate blame and responsibility on the ultimately fair basis of the agent's own contribution, no more and no less. It is an illusion to suppose that this demand can be met.'[26] Relatedly, many also argue that the categories of moral obligation and blameworthiness presuppose an implausible model of the self, along Kantian or Christian lines, imagining that deep moral responsibility exists in a way that seems incoherent.[27] While I am not presenting any detailed arguments here against the traditional moral categories, and there may be plenty of scope for defending these categories against objections of this kind, such worries should at least check the urgency of consequentialists to squeeze their theory of value into a shape that fits neatly on top of the categories of moral wrongness, obligation and responsibility.

Finally, another reason why the consequentialist should not seek to outline a peculiarly consequentialist moral theory is that in some of the most important areas of ethics, what our moral obligations 'really are' is, at best, radically unclear, and perhaps even indeterminate. I think this is particularly clear in the case of imperfect duties. Consider again, for example, our duties to help the distant poor and to work towards preserving the environment. From a consequentialist point of view, given some plausible empirical assumptions, both these issues are hugely pressing. Nevertheless, it is not at all clear what the nature and extent of our 'real' moral obligations are here. How much is demanded of us? How

---

[25]Nagel (1979b)

[26]Williams (1985), 194; see also Williams (1995c)

[27]Again, both Williams and Anscombe suggest that common sense recognises uses of 'ought' and 'obligation' and 'wrong' that are quite distinct from the same terms and concepts as they appear in much modern moral theory.

much sacrifice is required of us? Many consequentialist philosophers attempt to outline precisely what our obligations to the distant poor really are. For example, some philosophers try to pick out percentages- say 10% or 25%- of our income that we are obliged to give in order to help the distant poor.[28] To pick out an amount of money or time that is 'really obligatory' seems inappropriately artificial. Imagine, for instance, that we somehow settle on 25% of our income as what is 'really' morally obligatory. The moral categories divide courses of action in such a way that giving 25% is permissible, while giving only 24% is forbidden, morally wrong. Likewise, it acclaims giving 25% and 75% equally as 'permissible'. Surely, the scalar judgements of GC seem much more natural in this context of so-called 'imperfect' duties. To non-philosophers, the claim that there is an exact figure that is morally obligatory is surprising to say the least, in the same way that it is surprising to learn that there is a real determinate point in the spectrum of colours where red becomes not-red. I am not of course suggesting that it wouldn't be good if people gave, say, 10% of their income towards relieving severe poverty. In fact, I take it that GC implies just this sort of claim, again given empirical assumptions.

Certainly, GC implies that it would be good if we treated concern for the environment or the distant poor as obligatory, as expected, and apathy or inaction about these issues as worthy of shame, guilt or blame, if this pattern of sentiments would result in good consequences. A society in which we, in wealthy western countries, take seriously the plight of those in the Third World, and regard ourselves as having substantial obligations towards them, is, from the consequentialist point of view, a better society than our own. However, it is not clear that the fact that it would be better if we treated X as morally obligatory implies that our moral obligations are therefore such-and-such. Even if a particular set of norms N is the optimal one for a given society, this does not amount to saying that these *are* the moral obligations that people have in this society. We can pick out the factors that are relevant to setting a level of what can be expected of us- the value at stake, and the cost or sacrifice to the agent- but there is no formula for determining the level of our 'actual' obligations.

The ambition of many moral philosophers to find an answer to the question of what our moral obligations 'really are' may be vain. Brad

---

[28] See e.g. Singer (1993), Chapter 8

Hooker says that 'we frequently wonder whether we are doing what morality requires of us in terms of helping others [especially the needy]. And it is far from obvious that any familiar moral theory has an intuitively acceptable line on this matter.'[29] In the light of this repeated failure of moral theory, a refocusing of our ethical thought away from questions of moral obligation looks prudent. In the continuing absence of any plausible systematic theory of moral demands, a mild scepticism about the merits of systematic moral theory itself will be attractive. Perhaps a more useful ethical picture is one in which we recognise that we do in fact treat some things as required or expected (on pain of sanctions), and other things as forbidden, and then evaluate *these* practices as better and worse.

In it is no coincidence, I think, that these problems of radical unclarity, or even indeterminacy, are most vivid in the case of our moral duties to the distant poor and to the environment.[30] As I suggested above, it may help here to think in terms of the evolution of our moral practices and moral sentiments. These practices and sentiments evolved to deal with problems and conflicts that are moral local than the ones I have been discussing. It is unclear how the moral categories that shape our ethical thought in these more local matters are supposed to carry over to global issues.[31] The consequentialist sensibly says that the welfare of someone distant to us is just as important as that of someone who lives nearby. However, it is unsurprising that the plight of the distant poor fails to motivate us to action in the way that similar plights, were they to occur on our streets, presented to us vividly, would motivate us to action. Were we to pass emaciated and disease-ridden children on our way to work every morning, we would be hugely more motivated to act. As things stand, we are in fact affected far more deeply by having said an ill word to someone, a minor insult, than by failing to help the distant poor. And we shun someone who breaks a promise to us, rather than someone who does nothing to

---

[29]Hooker (2000), 159

[30]I think Scanlon takes there to be some sort of indeterminacy in such cases; or at least, he thinks that whether a particular level of sacrifice is morally required or not is not determined by (in his case, contractualist) moral *theory*. 'Contractualism, as a theory, does not tell us exactly what level of sacrifice is required by this principle [a Rescue Principle, concerning situations in which 'you can prevent something very bad from happening, or alleviate someone's dire plight, by making only a slight (or even a moderate) sacrifice'].' Scanlon (1998), 224-5

[31]See Fishkin (1982), for a discussion of how substantial theoretical problems arise when we try to extend our moral intuitions about local, small-scale cases to global, large-scale cases.

help the needy in Africa. Further, given that the distant poor yield very little power, and future people who may be harmed by our neglect of the environment yield none, recognition of their problems and welfare is not entrenched in our practices. So there is a huge gulf between our considered judgement as to what is really valuable and what motivates us to action. An impartialist consequentialism insists that this struggle to be motivated when the plight is far away does not make helping the distant needy or future people any less choiceworthy. Even when we fail to be motivated to any great extent, there is still as much at stake, and there is still the same amount to be said for helping the needy.

It is essential, of course, to be realistic about the extent to which we can change our motivations. Perhaps the best we can realistically hope for is consistent incremental changes in our awareness of, and motivation to act upon, considerations grounded in the welfare of distant others. It is obvious but true to say that humans are naturally local, embedded creatures, entrenched in their everyday lives, and motivated by the everyday concerns and relationships that are vivid to them. This is why, even when we judge that a quite different pattern of motivation and set of norms of social expectation would be better than our current motivations and norms, it sounds very strange to say that in fact our 'real' obligations are to do what is optimal.

## Conclusion

To conclude then, I want to summarise what I hope to have said in response to the three-part 'challenge' that I outlined at the beginning. Firstly, the consequentialist can happily say much of what common sense says in making ethical judgements, since most true common sense ethical judgements are quite compatible with GC, with consequentialism about the value or choiceworthiness of our acts and behaviour. While my form of consequentialism does not give a systematic account of what our moral obligations are (for the reasons I have outlined), this does not amount to denying that we have moral obligations, under an everyday, non-technical construal. Secondly, I have tried to show why consequentialists should not engage in a project of deriving a moral theory from their view of value, and in general should not be preoccupied with the task of fitting consequentialism into the traditional deontological categories of right,

wrong and obligation, and thereby offering a distinctively consequentialist moral theory.

## BIBLIOGRAPHY


Anscombe, G. E. M. - (1958) 'Modern Moral Philosophy', in Crisp and Slote (eds.) (1997)

Ashford, Elizabeth- (2000) 'Utilitarianism, Integrity and Partiality', *Journal of Philosophy*, XCVII

Bennett, Christopher- (2003) 'Personal and Redemptive Forgiveness', *European Journal of Philosophy* 11:2

Bratman, Michael- (1994) 'Kagan on "The Appeal to Cost"', *Ethics* 104

Brink, David O. - (1988) 'Sidgwick's Dualism of Practical Reason', *Australasian Journal of Philosophy*, 66

Cottingham, John- (1983) 'Ethics and Impartiality', *Philosophical Studies* 43

Crisp, Roger- (1996) 'The Dualism of Practical Reason', *Proceedings of the Aristotelian Society*, 96

- (1997), with Slote, Michael (eds.) *Virtue Ethics*, Oxford

Driver, Julia- (2001) *Easy Virtue*, Cambridge

Eriksson, Bjorn- (1997) 'Consequentialism For Sinners', *American Philosophical Quarterly*, 1997

Fishkin, James- (1982) *The Limits of Obligation*, Yale

Foot, Philppa- (1985) 'Utilitarianism and the Virtues', *Mind* 94

Frankena, William K.- (1992) in Schultz (ed.) (1992)

Gibbard, Allan- (1990) *Wise Choices, Apt Feelings*, Oxford Clarendon

Goodin, Robert E.- (1991) 'Utility and the Good' in Singer (ed.) (1991)

Hooker, Brad- (2000) *Ideal Code, Real World*, Oxford

Hooker, Brad; Mason, Elinor; Miller, Dale E. (eds.)- (2000) *Morality, Rules, and Consequences*, Edinburgh

Kagan, Shelly- (1989) *The Limits of Morality*, OUP

- (1994) 'Defending Options', *Ethics* 104

- (2000) 'Evaluative Focal Points', in Hooker, Mason, Miller (eds.) (2000)

Kupperman, Joel- (1981) 'A Case For Consequentialism' in Pettit (ed.) (1993)

Mackie, J. L.- (1977) *Ethics: Inventing Right and Wrong,* Harmondsworth: Penguin

- (1992) 'Sidgwick's Pessimism' in Schultz (ed.) (1992)

Melden, A. I., (ed.)- (1958) Essays in Moral Philosophy, Seattle: Washington Press

Mulgan, Tim- (2001) *The Demands of Consequentialism,* Oxford

Nagel, Thomas- (1979a) *Mortal Questions,* Cambridge

- (1979b) 'Moral Luck' in Nagel (1979a)

Nietzsche, Friedrich- (1990) *Beyond Good and Evil,* Penguin

- (1994)*The Genealogy of Morality-,* Cambridge

Pettit, Philip- (1991) 'Consequentialism' in Singer (ed.) (1991)

- (ed.) (1993) *Consequentialism,* Dartmouth

Raphael, D. D., 'Hume and Smith on Justice and Utility', *Proceedings of the Aristotelian Society,* 1972-3

Raz, Joseph- (1999) *Engaging Reason,* Oxford

Scanlon, T. M.- (1998) *What We Owe to Each Other,* Harvard

Scarre, Geoffrey- (1996) *Utilitarianism,* Routledge

Scheffler, Samuel- (1982) *The Rejection of Consequentialism,* Oxford

- (ed.) (1988) *Consequentialism and its Critics,* Oxford

- (1992a) *Human Morality,* Oxford

- (1992b) 'Prerogatives Without Restrictions', *Philosophical Perspectives,* 6

Schneewind, Jerome- (1977) *Sidgwick's Ethics and Victorian Moral Philosophy,* Oxford

Schultz, Bart (ed.)- (1992) *Essays on Henry Sidgwick,* Cambridge

Shaver, Robert- (2004) 'The Appeal of Utilitarianism', *Utilitas* 2004

Sidgwick, Henry- (1907) *The Methods of Ethics*, 7th Edition, Indianapolis; Hackett

Singer, Peter (ed.)- (1991) *A Companion to Ethics*, Blackwell

- (1993) *Practical Ethics*, Cambridge

Skorupski, John- (1998), 'Morality and Ethics', *Routledge Encyclopaedia of Philosophy*, London

(1999a)- *Ethical Explorations*, Oxford

(1999b)- 'Freedom, Morality, and Recognition: Some These of Kant and Hegel', in Skorupski (1999a)

- (2001) 'Three Methods and a Dualism', *Proceedings of the Pritish Academy* 109

- (2004) 'Externalism and Self-Governance', *Utilitas* 16

Slote, Michael- (1984) 'Satisficing Consequentialism' in Pettit (ed) (1993)

- (1985) *Commonsense Morality and Consequentialism*, London; Routledge and Kegan Paul

Smart, J. J. C. and Williams, Bernard (1973) *Utilitarianism: For and Against*, Cambridge

Stocker, Michael- (1976) 'The Schizophrenia of Modern Ethical Theories', in Crisp and Slote (eds.) (1997)

Urmson, J.O.- (1958) 'Saints and Heroes' in Melden (ed.) (1958)

- (1988) *Aristotle's Ethics*, Blackwell

Waldron, Jeremy- (1994) 'Kagan on Requirements; Mill on Sanctions', *Ethics* 104

Williams, Bernard- (1973) with Smart, J. J. C., *Utilitarianism: For and Against*, Cambridge

- (1976) 'Moral Luck', *Proceedings of the Aristotelian Society*, Supp. Vol. 50

- (1985) *Ethics and the Limits of Philosophy*, Fontana

- (1993) 'Moral Luck: A Postscript' in Williams (1995a)

- (1995a)*Making Sense of Humanity*, Cambridge

- (1995b) 'The Point of View of the Universe: Sidgwick and the Ambitions of Ethics' in Williams (1995a)

Wolf, Susan- (1982) 'Moral Saints' in Crisp and Slote (eds.) (1997)