

THE DEMANDINGNESS OF MORALITY

Keith John Horton

M.Phil Philosophy

University College London

ProQuest Number: 10045591

All rights reserved

INFORMATION TO ALL USERS

The quality of this reproduction is dependent upon the quality of the copy submitted.

In the unlikely event that the author did not send a complete manuscript and there are missing pages, these will be noted. Also, if material had to be removed, a note will indicate the deletion.



ProQuest 10045591

Published by ProQuest LLC(2016). Copyright of the Dissertation is held by the Author.

All rights reserved.

This work is protected against unauthorized copying under Title 17, United States Code.
Microform Edition © ProQuest LLC.

ProQuest LLC
789 East Eisenhower Parkway
P.O. Box 1346
Ann Arbor, MI 48106-1346

Abstract

In chapter one I introduce the problem of the demandingness of morality; the problem, that is, that we seem to have good grounds for adopting a form of morality that is considerably more demanding than our ordinary, common sense morality. I argue that such grounds need not depend on any particular, contested moral theory, but can be drawn from our ordinary, pre-philosophical moral consciousness, and hence that the problem is one that ought to concern us all. Given that, I claim that the best way to make progress with the problem is by developing a framework for discussion that as many of us as possible can share, a framework that puts into clearer focus the kinds of issues on which an eventual resolution of the problem will depend.

The rest of the thesis is dedicated to the attempt to develop such a framework. In chapter two I examine some moral theories, and their responses to the problem of demandingness, in the light of a fundamental principle which I claim is common to each of them: that, at some fundamental level, morality should take everyone into account, and take each of them into account in the same way. I argue that the most plausible interpretation of this principle, and the one that offers us the best hope of making progress with the problem, is one that requires us to take everyone into account distinctively in their role AS AGENTS. Doing so may, I claim, give us legitimate moral reasons to limit the demands of morality, but only if those demands are considered in the light of what I call an AGENT-CENTRED MODEL OF EXCELLENCE. I explore this idea, and the kind of role that it might play in a moral theory, in chapter three, and I argue that it raises

neglected issues that are absolutely central to making progress with the problem, issues which we can, indeed, all discuss together, whatever our views about moral theory.

CONTENTS

CHAPTER I THE PROBLEM OF DEMANDINGNESS

i/ Introduction	5
ii/ Common Sense Morality	10
iii/ The Principle of Moral Equality	21

CHAPTER II MORAL THEORY

i/ Consequentialism	32
ii/ Contractarianism	37
iii/ Nagel and Scheffler	43
iv/ Other Responses	54

CHAPTER III THE MORAL AGENT

i/ The Idea of an Agent-centred Model of Excellence (1)	60
ii/ The Idea of an Agent-centred Model of Excellence (2)	68
iii/ A Stringent Model?	77
iv/ Tying the Threads Together	90

BIBLIOGRAPHY	99
--------------	----

CHAPTER 1: THE PROBLEM OF DEMANDINGNESS

i/ Introduction

The problem of the demandingness of morality, as I shall construe it, is this: that we seem to have good grounds for adopting a form of morality that is considerably more demanding than what I shall call 'common sense morality' ('CSM' for short) - the customary moral norms and standards that most of us feel answerable to. I'll call any such demanding form of morality (demanding, that is, by the standards of CSM) 'stringent'. The existence of such grounds is considered to constitute a problem not only because it would constitute a challenge to our present moral practices, but also because the demands of some stringent forms of morality are felt by many to be unreasonable or unrealistic.

Because of these difficulties, a lot of attention has been dedicated in recent moral philosophy to examining whether the apparently plausible grounds for stringency can, in fact, be undermined. There are basically two kinds of responses of this type: 1/ that of denying that there are, in fact, any such good grounds even to suppose that morality should be stringent; and 2/ that of accepting that there are such grounds, but arguing that, in the broader picture, there are other considerations that counterbalance them, considerations which vindicate something like CSM. I shall refer to 1/ and 2/ as 'the first response' and 'the second response' respectively. If neither of these responses are found to be adequate, then the conclusion will be that we should, indeed, adopt a stringent form of morality. Someone who accepts this conclusion will then incur the responsibility of

explaining both why CSM has gone wrong, and how the claim that the demands of such a stringent form of morality are unrealistic or unreasonable can be answered.

What, then, are the grounds for adopting a stringent form of morality? In the philosophical literature on this subject, a variety of grounds have been put forward, many of which are dependent upon one particular moral theory or other. This creates the following difficulty: if we formulate the grounds for stringency in terms that depend on the acceptance of one particular moral theory, then those who reject that particular moral theory will not need to accept either those grounds or the argument for stringency that is based on them. We will only be able to conduct a kind of in-house debate on the issue. Thus if, for example, you agree with John Stuart Mill that morality requires an agent to be, "as between his own happiness and that of others...as strictly impartial as a disinterested and benevolent spectator" (J.S.Mill (1), p.17), then a strong argument for stringency seems to follow pretty directly, but if you do not believe that morality requires any such thing, then you have, as yet, been given no reason to worry about the issue. If, furthermore, all grounds for stringency were internal in this kind of way to some highly contested moral theory or other, then until you had been given a good reason to accept any such theory, you would not have been given a good reason to worry about stringency. So it may seem that, before being able even to properly examine whether there are good grounds for stringency, we will have to decide which moral theory to adopt.

That task would be forbidding enough, but matters are made even worse by the fact that in recent moral philosophy there have been

various strong criticisms, not only of this or that particular moral theory, but also of certain common elements in many of the moral theories in terms of which arguments for stringency have been framed. Thus some philosophers, whom one might refer to collectively as the 'anti-impartialists', claim that many of these theories impose a condition of impartiality on our moral thinking in an unacceptably strong form, one which cannot be supported by our ordinary moral practice. And another group of philosophers, the 'anti-theorists', claim, among other things, that the very conception of a moral theory employed by many modern moral theories - that is, as consisting of a principle or set of principles that would enable us to determine how to think about any particular moral situation - is itself at fault. It is claimed that, in fact, NO principle or set of principles could capture the complexity and context-sensitivity of our actual moral thought, because no such set of principles could replace the role played by sensitivity, or judgement, or practical wisdom, in determining what is the right thing to do. So in both cases, it is claimed that certain alleged defects in different moral theories can be traced back to assumptions that are common to them all. More positively, there are also some philosophers, the 'virtue ethicists', who attempt to develop an alternative picture of ethics which is not subject to these criticisms. In relation to the anti-impartialists, they claim that we should start, in our moral thinking, from the perspective of the particular moral agent, asking the question of how she should live, rather than from some abstract conception of impartiality; and in relation to the anti-theorists, they suggest that what a moral agent needs, in order to act rightly, is not a theory or

a set of principles, but certain virtues. Although it seems to me important to distinguish these three groups of philosophers, it is true that many philosophers do fall into two or all three categories, and it will sometimes be useful to group them together; when I do so, I will call them the 'critics of moral theory'.

Now, because many of the grounds standardly given for stringency depend on the acceptance of one or another type of moral theory that fits the model that is subject to these criticisms, those who reject this model will also have a reason to reject those grounds. So, if the criticisms of this model of moral theory are good ones, it is not so much that we will have to decide which moral theory to accept before deciding in what terms we should frame the problem of demandingness; it is rather that we will have to find a way of framing the problem that doesn't depend on the acceptance of ANY moral theory in the proscribed sense. I believe that this kind of consideration may have led some to suppose that the first response (that is, that of denying that there are even good prima facie grounds for supposing that morality should be stringent) is more plausible than it in fact is. Because the problem of demandingness is so often framed in terms that are subject to these criticisms, it might then appear that the problem is itself a creation of the questionable assumptions made by the moral theorists, so that, if we resist these assumptions, the problem will simply disappear. I don't think that I have ever seen this claim being explicitly made, but evidence that something like it is lurking in the background might be drawn from the fact that it does tend to be the proponents of certain particular moral theories who actually discuss, and seem concerned about, the problem of demandingness, while the

critics of moral theory tend either to give it very cursory treatment, or to neglect it altogether.

With this difficulty in mind I begin, in section of ii/ of this chapter, by arguing that there are considerations drawn from our ordinary, pre-theoretical moral consciousness that give us good reasons to worry about stringency, considerations that do not depend on the kind of assumptions questioned by those who are critical of moral theory. If I am right about this, then it will not, after all, be the case that the problem is merely the creation of a particular moral theory, or of a general, arguably suspect, theorizing impulse in moral philosophy. It will, instead, be a problem that ought to concern us all. But if it is a problem that ought to concern us all, then we should try to find a framework for the discussion of it that all, or at least as many as possible, can be party to. The main aim of this thesis is to try to establish such a framework: one that at least enables us to see more clearly the kinds of issues on which an eventual resolution of the problem will depend. Since any attempted resolution of the problem is unlikely to be successful if it doesn't have these issues properly in focus, I concentrate on the prior task of trying to establish such a framework, rather on trying to put forward a particular solution of the problem.

I take the first steps in trying to provide such a framework in section iii/ of this chapter, with the introduction of a certain principle which I call the Principle of Moral Equality (or 'P1' for short). With regard to the first difficulty that I mentioned above, that of apparently having to decide which moral theory to adopt before being able to tackle the problem, I claim that this principle would be

accepted, though developed differently, by many different influential moral theories, so that it represents a constraint on the selection of acceptable moral principles that would be recognised by them all. So we can at least begin with a principle that is neutral between these different theories, and then seek to evaluate them in terms of the adequacy of their interpretations of this principle. With regard to the second difficulty, that posed by the critics of moral theory, I claim that this principle may not be as vulnerable to what is most persuasive in their criticisms as might initially appear.

I then go on, in chapter II, to examine the ways in which some influential modern moral theories develop the principle, and the resources their interpretations permit them in pursuing the second response. I consider a number of theories, and argue that the kind of approach pursued by Thomas Nagel and Samuel Scheffler, though seriously flawed, may point the way to the most satisfactory interpretation of P1, one that may enable us to make progress with the problem. In chapter III I go on to suggest how I think those flaws can be rectified, in a way that puts into focus a neglected issue that is, I claim, absolutely central to the problem. I also claim that the kind of framework that puts this issue into focus is compatible not only with P1, which seemed to be the source of much of the appeal of many modern moral theories, but also with much of what is most compelling in the writings of the critics of moral theory.

ii/ Common Sense Morality

In order to show that there are prima facie grounds for stringency derived from our ordinary moral consciousness, grounds that we can

recognise without drawing on the kind of assumptions questioned by the critics of moral theory, it will be necessary to find considerations that, firstly, are widely recognised, from the perspective of CSM, to constitute defeasible moral claims; and, secondly, which there is good reason to think should be taken much more seriously than they currently are, in such a way as to support the adoption of a stringent form of morality.

The strongest candidate for such a type of consideration seems to me to be the following. There are vast numbers of people in the world today whose vital needs are unmet (call them the 'needy'), and also vast numbers of people, far more than at any previous point of world history, who not only have enough resources to meet their own needs, but a considerable surplus (call them the 'comfortable'). These two groups, the comfortable and the needy, are not quite insulated from one another: modern communications make the comfortable vividly aware of the plight of the needy, and there are various effective aid agencies that are available to transfer some of their surplus resources to the needy. Given this situation, the first thing to establish is whether the needy make defeasible moral claims on the comfortable that are widely recognised from the point of view of CSM.

It seems to me that CSM does recognise these claims. Various aspects of the complex moral tradition in which the comfortable have been brought up reinforce whatever natural feelings of sympathy or compassion or justice that they possess, and call out for action. Although the response is sporadic and inadequate, it does not seem to me that the comfortable do not hear or acknowledge the call, that is, that they do not recognise that the plight of the needy makes

defeasible moral claims on them; it is rather that these moral claims are recognised, but, if inaction is to be defended at all, are indeed defeated by other considerations. It is not altogether easy to decide how one might substantiate this claim further, in part because of the lack of research that has, as far as I know, been carried out into CSM, but if there are any doubts, perhaps the following considerations may help to alleviate them.

In the first place, the findings of one such piece of research that has been carried out, by James Fishkin, support this claim, and, indeed, the claim that what Fishkin calls "ordinary moral reasoners" are very sensitive to grounds that appear to support the adoption of a stringent form of morality (see J.Fishkin (1)). Secondly, one might point out how susceptible "ordinary moral reasoners" seem to be to arguments for stringency, on the basis of the unmet vital needs of others, put forward by philosophers such as Peter Singer (see P.Singer (1) and (2)), whatever reservations some philosophers may have about them. Such arguments, although they introduce some moral theory, seem to depend for their troubling power on moral intuitions that people have prior to their exposure to moral philosophy, intuitions derived from the ordinary moral consciousness of the culture in which they have grown up. Thirdly, one might point to the readiness of so many people, under appropriate prompting (such as that provided by organisations like Oxfam, Comic Relief, and so on) to respond practically to the needs of others. Such phenomena also show that the call that comes from the needs of others isn't restricted to those with which the potential donors already have a special relationship of some kind, even that of simply being citizens of the same country. And

finally, it does not usually seem to be the case that, when the issue of the needy is raised, the comfortable simply deny the existence of the relevant moral claims. More commonly, the comfortable either point to countervailing considerations; or claim that there is, in fact, nothing that they can do; or, indeed, concede that they are at fault. These kinds of reactions do not suggest that moral claims based on the plight of the needy are not recognised from the point of view of CSM at all, but rather that they are recognised, but that they are either defeated by other considerations, or that the structures that would enable them to be acted on are lacking, or that the indifference of the comfortable is indeed culpable.

Lest there be any confusion, let me emphasise that I have NOT claimed that, from the perspective of CSM, the needs of unknown others should count equally with the needs of those to whom the agent stands in a special relation, nor that physical proximity should be regarded as irrelevant, or anything of that sort. All I have claimed is that CSM generally recognises defeasible moral claims on the comfortable stemming from the needs of others, where it is not necessary for the comfortable to stand in any special relationship with those others for them to feel the force of those claims. Note also that, in order to support this claim, I have not relied on any abstract thesis about the nature of morality, such as its involving impartiality; nor I have I relied on any principle - such as 'if you can provide a great benefit to someone else at little cost to yourself, you should do so' - that might elicit the criticisms of the anti-theorists. All I have drawn upon is the common moral feelings and reactions of CSM.

Assuming, then, that CSM does recognise that the needy make

defeasible moral claims on the comfortable, the next thing to establish is that we might have reason to think, still without going beyond CSM, that those claims should be given considerably more weight than they are currently given, in such a way as to constitute grounds for stringency. Once again, it seems to me that there is good reason to think this; many people who are quite innocent of moral theory are struck by the difference between the life expectations of the comfortable and those of the needy, and moved to think that the comfortable should, indeed, be doing more. What prompts this thought, characteristically, is not any kind of moral theory, but rather the standard moral feelings of the culture in which those concerned have grown up; a sense of justice, perhaps; or a feeling of sympathy or compassion; or, as David Wiggins says, "a particularized form of what Hume called 'the resentment of the misery of mankind'" (D.Wiggins (1), p.30). So while many of the grounds for stringency put forward by moral theorists may indeed be deeply questionable, it seems to me that there are prima facie grounds for stringency that are given by the ordinary moral consciousness of our culture, and which, therefore, we can all recognise.

Given that stringency is defined in opposition to (as being 'considerably more demanding than') CSM, it follows, of course, by definition that CSM does not currently accept stringency, and so there may appear to be an air of paradox in arguing for stringency on the basis of CSM. But for such an argument to be possible, all that is necessary is that, among the complex web of elements that make up CSM, we have reason to think that some should be given more weight than they currently are, in such a way as to motivate a move away from CSM

in the direction of stringency. This would involve a certain sort of 'error theory'; that is, the claim CSM has currently got the balance wrong, and an explanation of why it might have done so; but the considerations that support stringency are, nevertheless, derived from CSM. And there seem to me to be very strong reasons to explain why CSM might have currently got the balance wrong, including the following.

In the first place, it may be noted that CSM, based as it is in deeply entrenched habits of thought and feeling, may be slow to change in response to changes in circumstances, particularly when the pace of change is very fast, as in recent times. So it may be that CSM is in need of overhaul, in particular in terms of the demands it makes on the comfortable with respect to the needy, because it has failed to keep pace with these changes. The two most important types of relevant change are the following: firstly, changes in our VALUES, particularly with regard to the increased emphasis placed in recent times on the importance of equality; and secondly, changes in our PRACTICAL CIRCUMSTANCES, such as improvements in communications that have made the world seem a much smaller place, and the fact that so many of us in the West now find ourselves in the ranks of the comfortable. It may indeed once have been the case, as J.S.Mill claimed, that "the occasions on which any person (except one in a thousand) has it in his power...to be a public benefactor, are but exceptional" (J.S.Mill (1), p.20), but it is assuredly not the case now. So perhaps CSM needs reform because it has failed to keep track with these changes.

Relatedly, it may be that some elements of the complex moral tradition that we have inherited, and which make up CSM as it stands today, would be less likely now than others, on reflection, to

continue to hold our allegiance, and that this too would provide support for a reform of CSM. It may be the case, for example, that some elements of CSM are more subject ^{to} ideological critique than others; that is, we may feel that we are less inclined to support them once we have a clearer idea of their origins, or of the kind of social and psychological mechanisms that explain their hold on us. It is at least arguable that that part of our common sense moral thought which is responsive to claims of need is more likely to stand firm under such ideological critique than some other elements; and that too would support giving such claims a larger role in our general moral economy than CSM currently gives them.

Another related point is that, as Samuel Scheffler says, "the reliability of the 'common sense' moral instincts shared by the members of a given society may depend to a significant extent on the justice or injustice of the society in question. In a seriously unjust society, well-intentioned people may have internalized moral attitudes that make the injustice of their institutions genuinely difficult for them to perceive" (S.Scheffler (2), p.143). If, as most people would agree, the societies of the past have, in general, been considerably more unjust than today, and if much of the material of CSM was formed in those societies, that might give us another reason to suppose that we have good reason to reform CSM, possibly in the form of making demands on the comfortable that are, by the standards of CSM, stringent.

Finally, we should be aware of all the elements of special pleading and rationalization that are likely to be present when it is we, that is, largely, the comfortable, who are asking whether we should do more

for the needy. Because it may serve our interests to maintain the status quo, which asks so little of us with regard to the needy, we should ask whether the status quo is being maintained, at least in part, precisely BECAUSE it serves our interests. This all too human tendency may go some way to explaining our refusal ^{to} recognise the grounds for stringency, rather than any deficiency in those grounds themselves.

So it seems to me that the first response - that is, denying that there are good grounds even to suppose that morality should be stringent - fails from the perspective of CSM. Given the amount of avoidable suffering in the world, and the fact that the comfortable can make a significant contribution at relatively little cost to themselves, it does seem to me that we have good grounds for thinking, without any prompting by moral theory, that the comfortable should do more, in such a way as to constitute stringency. Just how much would have to be required of the comfortable in order to constitute stringency is, given the way that I have defined stringency (as being 'considerably' more demanding than CSM), somewhat indeterminate. Just to have a relatively concrete example before us, if, for instance, we felt that those whose income is above a certain level should be morally required to give, say, 10% of their disposable income in aid to the needy, then, given that CSM presently doesn't actually seem to require ANY such contribution (though it regards as admirable those who DO contribute), would certainly seem to qualify as 'considerably' more demanding than CSM.

But there remains the second response. That is, one may accept that there are prima facie grounds for stringency, but argue that, in the

broader picture, there are, in fact, considerations that counterbalance those grounds, and vindicate something like CSM (considerations which are not, presumably, subject to the kind of objections listed above). Most of the discussion of the second response has been in the context of one or another of the currently influential moral theories, in part, no doubt, because some of these theories are explicitly committed to premises that cast the problem in a particularly acute form; it then becomes an urgent matter for advocates of these theories, if they are not to accept a highly stringent form of morality, to find a successful version of the second response. But I think that there is a second reason why most of the discussion of the second response has been in the context of one or another moral theory, and that is because the critics of moral theory have failed to develop the kind of resources necessary to tackling this question.

Sometimes, as I have said, it seems to be presumed that, once the errors of the moral theorists have been exposed, the problem of demandingness simply disappears; but we have already seen that we do not need ~~to~~ the prompting of moral theory in order to be concerned about the problem. Sometimes it is claimed that philosophy, in particular, has nothing to contribute to the problem. This seems to be the view of John McDowell, who says that "reflective thinking about specific norms (is) an activity that is not particularly philosophical" (J.McDowell (1), p.95). But it is not clear to me how McDowell would support this exclusionary claim about what is and is not philosophical. He does not say to what other field of enquiry the patient, honest, and responsible examination of our moral norms should

be consigned. No alternative candidate immediately presents itself. Sometimes, again, the response is to mount a staunch defence of those elements of CSM that militate against stringency, but this does not in itself tell us how those elements are to be balanced against those elements that DO support stringency (see, for example, J.Cottingham (1), and the (unresolved) problems he acknowledges he has with "justice as fairness and equity" (p.68)). And sometimes the problem IS acknowledged by the critics of moral theory, but even in these cases, not in such a way as to enable us to make progress with it.

Let me illustrate this tendency with a passage from perhaps the leading critic of moral theory, Bernard Williams. After criticizing the way in which the challenge to our present practices posed by the distant needy is construed by certain moral theorists, Williams writes, "Of course this point does not dispose of the challenge itself. We should be more concerned about the sufferings of people elsewhere" (B.Williams (4), p.186). So it is apparently just obvious to Williams, at least, both that the "challenge" remains, and, indeed, that our present patterns of concern are not justifiable. So what kind of pattern of concern WOULD be justifiable? The frustrating thing in Williams, as in so many of the critics of moral theory, is that no progress is made in answering this question. One might have supposed that, once the issue was framed in a way that avoided the errors of moral theory, one would then be shown a better way to think about the problem, in such a way that, however approximately or provisionally, one would be able to determine what kind of role SHOULD be given to concern for the needy in our general moral economy. But no such account is forthcoming. Williams does make some comments elsewhere

that might be thought to have a bearing on the issue; for example, he says that "a prejudice of the racist or sexist kind is usually a belief guarded against reflection because it suits the interests of the believers that it be held" (B. Williams (4), p.117). This is particularly interesting, because advocates of a stringent form of morality would presumably apply the same sort of analysis to our beliefs about the responsibilities of the comfortable to the needy; but Williams doesn't explore this possibility.

Of course, Williams has his own agenda, and is at liberty to discuss or not to discuss whatever issues he chooses. But I think his neglect of the issue that concerns us is more significant: it seems to me to be emblematic of a general tendency among the critics of moral theory not to examine the issue constructively. And the effect of this tendency is, indeed, conservative: one that resists the effective challenging of the status quo.

In this spirit, then, and given the urgency of the problem, I suggest that we put our doubts about the theorizing impulse on hold, and take a look at what the advocates of some moral theories have to say about the problem. But before doing so I will introduce the principle that I referred to in section i/ above, a principle which is, I claim, at least neutral between many of these different moral theories. If, as I shall also claim (in section iii/ of this chapter), this principle may not in fact ^{be} as vulnerable to what is most compelling in the objections of the critics of moral theory as may at first appear, and that (as I shall argue in chapter III) it can be developed in a way that both meets most of those objections, and offers a way to make progress with the problem, then we will have all

the more reason for taking this step.

iii/ The Principle of Moral Equality

The principle that I want to claim is common ground to many modern moral theories, and the source of much of the appeal of each of them, is the following:

(P1): At some fundamental level, morality should take everyone into account, and take each of them into account in the same way.

A principle of something like this kind is virtually a commonplace of modern moral philosophy (see, in particular, C.Taylor (1), p.130 for a similar formulation). Its introduction here invites the following questions: firstly, why we should accept it; secondly; how it is compatible with many different moral theories; thirdly, why it might appear to give grounds for stringency; fourthly; why that appearance might, after all, be misleading; fifthly, what relation it might have to the idea that morality should be impartial; and lastly, how it might stand in relation to the critics of moral theory. Discussion of these questions will take us right through the rest of this thesis, but I shall try to say something about each of them, in an introductory way, now.

Why, then, should we believe that morality should, at some fundamental level, take everyone into account in the same way? Presumably because everyone is equal, in the relevant, moral sense; that is, because everyone is equally worthy of moral consideration. And the support for this claim, in turn, is that if we have any moral

standing at all, it is in virtue of some property or properties that we possess; and whatever such properties are taken to be fundamentally relevant to our moral standing are likely to be properties that are shared by us all. Different properties are put forward by different moral theories as being those which are fundamental to our moral status - for example, some emphasise the capacity for pleasure and pain, while others emphasise rationality - but the important point here is that whatever PLAUSIBLE candidate is suggested, it is likely to be a property that all human beings possess. If we agree with this, then, even if we disagree about which properties do in fact give us our moral status, we can agree that we are morally equal, and use this claim to support P1: if we are all equally the possessors of those properties, whatever they are, that are fundamental to our moral standing, then it seems that morality should reflect that moral equality by taking each of us into account, at some fundamental level, in the same way.

And indeed it does seem to me difficult to deny the claim of our moral equality. In order to do so we would have to find some property that was a plausible candidate for playing the role of having fundamental moral relevance, but that was not universally shared. It is not easy to see what such a property might be. THIS battle, at least, seems to have already been won. Even those who try to justify such morally abhorrent practices as, for example, racial discrimination, do not try to do so by reference to skin colour alone - surely an implausible candidate for a property that might be fundamentally morally relevant - but, as Bernard Williams says, by "reasons that seek to correlate the fact of blackness with certain

other considerations which are at least candidates for relevance to the question of how a man should be treated: such as insensitivity, brute stupidity, ineducable irresponsibility, etc" (B.Williams, (1) p.233). Such correlations are not, of course, borne out, but what the attempt to employ them shows is that even racists must agree that there are limits to what properties can be considered to be fundamentally morally relevant; as Williams says, such racists "are paying, in very poor coin, the homage of irrationality to reason" (ibid).

At a rather different level, we might find further support for P1 in the role that its acknowledgement has played in movements, such as those against sexism and racism, that few of us would not view positively now. By acknowledging everyone's equality of moral status, and hence the need to take everyone into account, at some level, in the same way, it becomes much harder to justify the kind of differential treatment of different groups of human beings that, until quite recently, was commonplace. P1 is indeed a strong weapon against many kinds of injustice and prejudice.

As for the second question, P1 is neutral across many different moral theories, in part because it does not specify what property it is that is fundamentally morally relevant; it just says that morality should take everyone into account in the same way, because WHATEVER property is selected is likely to be common to all human beings. Nor does it specify HOW everyone is to be taken into account. Different moral theories, as noted above, offer different candidate properties, and they also, relatedly, offer different procedures for taking everyone into account. Consequentialism takes everyone into account in

the same way, at the fundamental level, by giving their interests equal weight; Kantianism does so by endorsing only maxims that everyone can act on; certain forms of contractarianism do so by endorsing only principles that would be accepted by all from an original position of equality; and so on. The point is that these and other different moral theories, while disagreeing about HOW to take everyone into account, can agree about P1.

On to the third question: although the resulting moral theories are very different from one another, it seems to me that each of them faces a version of our problem, for each of them is likely to find, in one way or another, that people's needs are of particular importance from the moral point of view, and that morality will therefore have to assign a particular importance to the satisfaction of those needs. Once again, the way in which this is done will vary across different moral theories, but whatever the way in which people's needs come into the equation, it is not hard to see how, in a general way, P1 might give rise to pressure for stringency: if morality should take everyone into account, and there are so many people in the world whose needs are unmet, and meeting those needs is of particular moral importance, then there are likely to be prima facie grounds for supposing that those with the resources to meet those needs should do so.

On the issue of why, however, we should not immediately assume that P1 will commit us to stringency, note that the claim is that MORALITY, as a system or structure, should take everyone equally into account, not that THE INDIVIDUAL MORAL AGENT should do so. This distinction is sometimes obscured by an equivocation of over the notion of the 'moral point of view': when we are speaking of the 'moral point of view' we

might mean EITHER the point of view from which we determine what moral principles are acceptable, OR the point of view that the agent is to take up in her ordinary moral reasoning. Failure to make this distinction seems to me to cause no end of confusion, one of the consequences of which is the common presumption that by accepting a principle like P1, we will automatically be placing unreasonable demands on the individual moral agent - such as that she has to consider everyone's needs equally in deciding what to do. But note that P1 is supposed to function at the level of 'the moral point of view' in the former sense; that is, the level at which we determine which moral principles are acceptable. It does not automatically follow that the individual moral agent is required to take everyone into account in her ordinary reasoning; what the individual moral agent is required to do is to take up the 'moral point of view' in the second sense, UTILIZING the principles that have emerged from the 'moral point of view' in the first sense, but not necessarily OCCUPYING that point of view. So it may not inevitably follow that the acceptance of P1 leads to such a stringent form of morality as might at first appear; we shall examine in more detail how some moral theorists have attempted to substantiate this claim in chapter II.

The next question is about the relation between P1 and the claim that morality should be impartial. This is a complex matter, but for now I'll just sketch a way in which we might link P1 and the claim that morality should be impartial. What it is to be impartial, in any particular context, seems to depend on the type of considerations that are considered to be relevant in that context. A referee of a football match, for example, in order to be impartial, will have to take into

account the rules of football, but he should not take into account other considerations, such as which team he prefers, or whether one team is trying harder than another, or whether it would bring about more happiness if one team rather than another won. A judge, on the other hand, will have to take into account a different set of considerations, and rule others out as irrelevant to a court of law; but like the referee, impartiality will be a matter of taking into account certain considerations, and excluding others as irrelevant. Similarly for a bureaucrat, and so on. This suggests a general understanding of what it is to be impartial: to be impartial is to take into account only those considerations that are considered to be relevant to the kind of case in hand, and to disregard other considerations. Something like this conception of impartiality has been put forward by a number of writers: for example, J.S.Mill writes that "Impartiality, in short, as an obligation of justice, may be said to mean, being exclusively influenced by the considerations which it is supposed ought to influence the particular case in hand" (J.S.Mill (1), p.47). (See also H.Sidgwick (1), p.268 for a not dissimilar formulation.)

Now although, unlike many characterizations of impartiality, this characterization seem to me at least to have the merit of being accurate, it may not seem to be very helpful. Until we know, to take Mill's formulation, which considerations ought to influence the particular case in hand, we will not know what it is to be impartial in any particular context. But this is where the Principle of Moral Equality may be able to help us. It might be claimed that, in a moral context, the question of which considerations are to be regarded as

relevant is to be determined by reference to a process that, at some fundamental level, takes everyone into account in the same way. And the justification for this claim, in turn, is that, as I have argued above, as a matter of fact all human beings are equal in their possession of whatever fundamental properties it is that might be plausibly be considered to be relevant to their moral standing. So the claim is this: firstly, that impartiality, in general, is a matter of taking only relevant considerations into account; and secondly, that those considerations that are, in particular, MORALLY relevant are those that would be taken account of in an adequate moral system that is compatible with P1.

Now, frequently a rider to the general characterization is added, to the effect that the agent should not be swayed by personal considerations (desires, preferences, interests, etc.). Sometimes it is assumed that these are the only types of considerations that might sway us from impartiality, but I am not sure if this is true; however, the rider seems to be apt at least in that personal considerations are a particularly frequent and understandable reason for being swayed from an impartial judgement. But, once again, there is scope for equivocation in the claim that, to be impartial, the agent should not let herself be swayed by personal considerations, an equivocation which, like that over 'the moral point of view', causes no end of confusion. To say that an agent should not let herself be swayed by personal considerations might mean EITHER 1/, that she should not let personal considerations sway her from the appropriate verdict, the one that takes account of all and only relevant considerations; OR 2/, that personal considerations are altogether irrelevant to the type of

case in hand. Note that 1/, unlike 2/, does not exclude the possibility that personal considerations ARE relevant in certain contexts; it may be that, in certain contexts, personal considerations SHOULD be taken into account by the agent, but only to a certain point; and the claim that we should not allow ourselves to be swayed by personal considerations means, not that we should not disregard them altogether, but that we should not give them MORE importance than the type of case demands. (As Sidgwick says, we should not allow ourselves to be "UNDULY influenced by personal preferences" (H.Sidgwick (1), p.268, 'italics' added). Which considerations are MORALLY relevant is, according to the conception that I have proposed, to be decided by a procedure that is compatible with P1; but it does not automatically follow that personal considerations will be irrelevant to moral judgement; that is, that we must interpret the claim that one should not be swayed by personal considerations in the second rather than the first sense, so that whenever a moral agent takes personal considerations into account, that constitutes a failure of impartiality. The mistaken assumption that it DOES follow is, I believe the source of much of the disquiet of the anti-impartialists.

This takes us on to the sixth question, of how the critics of moral theory are likely to be disposed with respect to P1. The first point to note is that, whatever it is the anti-impartialists are attacking, it is surely not the assertion that in our moral thinking, we should take account only of relevant considerations. Nor should they be attacking the view that we should not be swayed by personal considerations in the first sense; that is, beyond what is appropriate for the type of case in hand. Rather, the whole issue comes down to

what considerations are to be taken as relevant, and to what extent, in moral thinking. So perhaps what the anti-impartialists might be interpreted as opposing is either the idea that the question of which considerations are morally relevant is to be determined by reference to a process that takes everyone into account in the same way, or P1 itself. If that is so, then they will have to try to undermine the apparently strong support for P1 that we found at the beginning of this section. In fact I suspect, as I mentioned above, that the source of much of their concern is the assumption that for morality to be impartial, it would have to embrace the second version of the rider about personal considerations. If that is so, and if we can find a plausible interpretation of P1 that gives what we feel to be a reasonable place to personal considerations among the considerations that are considered to be relevant to moral judgement, then we might have found a way to diffuse their criticisms.

As for the anti-theorists, while it is certainly true that many moral theories do develop P1 in a way that runs into their criticisms, it does not seem to me necessary to do so. Clearly P1, even though I have called it a 'principle', is not a 'principle' in the sense that has elicited to the criticisms of the anti-theorists; that is, something that can tell us how to think about any particular moral situation; nor is there any necessity that it issue in any such principle or set of principles. P1 is rather a condition on the acceptability of our moral thought, a condition that does, indeed, appear to be supported by the common moral consciousness of our time. Nor does it seem to me that P1 is vulnerable to certain related criticisms of the anti-theorists - though, again, some of the ways in

which it is sometimes developed certainly are. In particular, some moral theories are criticised for their futile attempts to "justify ethical considerations from the ground up" (B.Williams (4), p.28), or to "reconstruct the moral point of view on the basis of one idea" (D.Wiggins (2), p.78); but it does not seem to me that the mere acceptance to P1 need commit one to either of these things. While P1 does impose a condition on the acceptability of our moral practices - that of showing how they are compatible with the equal moral standing of all - it does not inevitably imply that we must rethink or justify them from the beginning, even less on the basis of just one idea. We can rather remain within our customary moral consciousness, with just the thought that IF it should prove that certain of our practices are incompatible with P1, that would provide strong grounds to change them.

Similarly, I do not think that P1 is necessarily incompatible with what is most persuasive in the perspective offered by the virtue ethicists. This is basically because I don't think that we need to be forced into an either/or choice between thinking of morality in terms of impartiality, and thinking of morality from the perspective of the particular agent. In other words, I do not think that thinking of morality as impartial in the sense sketched above need be incompatible with taking the perspective of the historically situated agent, with her own life to lead, seriously. And this is because, in taking everyone equally into account, I believe that morality will have to take the agent's point of view into account. But I will defer a more detailed discussion of this issue until chapter III.

So, if what I have said above is correct, we have a principle that

is neutral between many different moral theories, plausibly the source of much of the appeal of each of them, and which the critics of moral theory, too, may not have to be too suspicious about. I shall say more about the latter question in chapter III, but for now I'd like to use this principle to structure an examination of the responses of some moral theorists to the problem. Two related questions will be in our minds as we make this investigation: firstly, whether the ways in which the different theories interpret P1 enables them to mount a strong version of the second response; and, secondly, whether the ways in which they interpret P1 are, in themselves, satisfactory.

CHAPTER II: MORAL THEORY

i/ Consequentialism

Much of the discussion of the problem of demandingness has been in the context of consequentialist theories of one sort or another, which is not surprising, not only because these theories have, in recent times, been so influential, but also because they face the problem in a particularly acute form. They do so for two reasons. The first is that these theories appear to give grounds for a particularly stringent form of morality. This is because all versions of consequentialism (as I shall understand the word) determine the moral value of whatever the appropriate primary object of moral evaluation is taken to be (whether individual actions, characters, rules, or whatever) exclusively in terms of their contribution to the promotion of good consequences, impersonally construed. In this way P1 is accommodated: everyone is taken into account in the same way, at the fundamental level, because everyone's interests are given equal weight in the evaluation of consequences that is ultimately criterial for moral evaluation. Different versions of consequentialism, as well as selecting different primary objects of moral evaluation, offer different criteria for the evaluation of consequences; but whatever plausible criterion is selected, strong grounds for a very strong form of stringency are likely to follow. To focus on the case we introduced in the last chapter, it is likely to appear that the comfortable can promote better consequences, impersonally construed, by giving to the needy than by dispensing of their surplus resources as they currently do. Not only that, but it appears that the comfortable should keep giving

until they reduce themselves to the level at which to give more would lead to worse overall consequences. So consequentialism seems to lead to an argument for a particularly extreme form of stringency.

The second reason why these theories face the problem in a particularly acute form is that the only resource available to them to limit this stringency is the very consideration - contribution to good consequences, impersonally construed - that led to such a strong version of stringency in the first place. Because contribution to good consequences, for the consequentialist, is uniquely criterial for moral evaluation, there is simply nothing else to which she can ultimately appeal as a reason for reducing the stringency of morality. For a consequentialist, then, it must somehow be the case that, contrary to what we might at first have thought, a less stringent form of morality will actually result in better consequences than the highly stringent form that we seemed to have strong grounds for. In other words, the only strategy available to consequentialism for limiting the demands of morality is to argue that it is COUNTER-PRODUCTIVE IN TERMS OF THAT VERY THEORY to adopt or espouse a form of morality that is highly stringent.

Different consequentialists have different opinions about how powerful this strategy is. Some claim that, even taking it into account, morality will still be highly stringent (Peter Singer and Shelley Kagan are notable examples). Others seem to think that if the consequentialist strategy can be shown to give us reason to reject certain extreme forms of stringency, the presumption is that something like CSM is vindicated. This seems to me to be an error, because there is a lot of space between an EXTREMELY stringent form of morality and

CSM. Others, again, try to occupy the middle ground. I will limit myself by trying to show that it is unlikely that the consequentialist strategy will give us reason to reject stringency as I have defined it.

One version of the strategy focuses on the goods that will, allegedly, be lost, if agents adopt a highly stringent form of morality: common candidates are the values of love and friendship, or of integrity, or those involved in certain cultural activities. Although this version may have some force against a highly stringent form of morality, I do not think that it will be very effective against stringency as I have defined it; although the importance of these goods is such that it may give us a consequentialist reason to reject a form of morality that precluded their realization, it seems wildly implausible to claim that the comfortable would be prevented from realizing these goods if they were required to give, say, 10% of their income to the needy. Given the value of the needy having their needs met, then, it seems that, from the consequentialist viewpoint, the comfortable should be required to do so.

A more promising version of the strategy is to claim that there are features of human nature, rather than certain specific goods, that give us reason to limit the stringency of morality; it is claimed that if we set the moral standard too high, the consequences would in fact, because of these features of human nature, be likely to be worse than if we set a less stringent standard, and so we have strategic reasons not to set the standard too high. An influential example of this version of the consequentialist strategy can be found in Peter Railton's article, 'Alienation, Consequentialism, and the Demands of

Morality'. Railton is an act-consequentialist; that is, he believes that the consequentialist criterion should be applied to individual actions, where the right action is the one that has the best consequences, impersonally construed. This does indeed appear to set an extremely demanding standard. Railton's response is to argue that an agent who "has certain traits of character, or commitments to persons or principles, that are sturdy enough that (she) would at least sometimes refuse to forsake them even when this refusal is known to conflict with making some gain...in total utility" may end up, overall, doing more good than an agent who tries to maximise utility with every act (P.Railton (1), p.18). The strength of Railton's argument, like that of similar examples of this version of the consequentialist strategy, depends on the force of two empirical claims: firstly, that the more an agent acts in a directly consequentialist way, the less she will be capable of forming satisfying personal relationships, or of otherwise meeting her personal and psychological needs; and secondly, that, at a certain point, such an agent would be likely to suffer some kind of psychological demoralization or breakdown that would lead her to do less good overall (see P.Railton (1), pp.120-1). And the general objection is this: even if we concede that, at some point, the attempt to maximise the good may lead to such a breakdown, it is highly unlikely that that point would be reached at such a level as to yield anything like CSM. In other words, though we might concede, in a general form, Railton his empirical assumptions, they would have to be made in an implausibly strong form to save us from stringency. In particular, it seems that, given the empirical circumstances in the

world, the comfortable could in general do MUCH more for the needy, without having to forego the opportunity of pursuing satisfying personal projects and relationships, and without therefore being likely to suffer the kind of psychological breakdown that would lead them to do less good overall.

It seems to me that that kind of point can be generalized to cover other versions of the general consequentialist strategy, and therefore that it will be hard to find a version of consequentialism that is able to avoid stringency. And indeed we may feel, on reflection, that it would be verging on the fantastical if things were otherwise; if CSM, in its present form, derived as it is from so many different moral traditions and influences (some of which may be subject to the kind of criticisms noted in section ii/ of chapter I), should just happen to coincide, even approximately, with the system that would bring about the best consequences, impersonally construed.

Now this, of course, is no objection to consequentialism, unless it can be established that the kinds of demands placed on the moral agent by whatever form of consequentialism we prefer are indeed unreasonable or unrealistic. Some reasons why we might think so will be explored in later sections. But for now notice how contingent the consequentialist strategy is, as a way of limiting those demands: for the consequentialist, there is no reason OF PRINCIPLE why the moral agent should not be required to do anything at all, provided only that others' lives could thereby be sufficiently improved. Railton hopes that the kind of considerations he discusses will show "how being moral (as the consequentialist understands it) might be compatible...with living a desirable life" (P.Railton (1), p.133). But

nothing in the consequentialist machinery guarantees such a result; it will only follow if, conveniently, the sums come out right. If, on the other hand, we felt that there were reasons of principle for limiting what can be required of moral agents, reasons that do not depend on the kind of strategic calculation that consequentialists rely on, and which were compatible with P1 and so could be impartially acknowledged, then we would not have to rely on the consequentialist strategy - and we would also have a reason for rejecting consequentialism as such. Such a possibility, too, will be examined in later sections.

ii/ Contractarianism

In speaking of 'contractarianism' in this section I will be referring exclusively to the type of view which seeks to reflect satisfy P1 by having moral principles chosen from some original position of equality, where this original position is understood in a Scanlonian rather than in a Rawlsian sense; that is, everyone retains knowledge of their identities, but is assumed to have the desire to reach reasonable agreement with others. Thus everyone is taken into account in the same way, at the fundamental level, by having an equal say in the selection of moral principles. Once again, it is not hard to see such a process might lead to the selection of a highly stringent form of morality: those who do badly under current arrangements are unlikely to accept CSM, and are likely to feel that they are justified in pressing for a more stringent form of morality that ensures that their needs are met.

Nevertheless, once again, there are those who argue that a

successful version of the second response may be available to contractarians; that is, that although contractarianism can accept that there are prima facie grounds for stringency, it can also draw on other considerations that counterbalance those grounds, and yield something like CSM. I will take Brian Barry as my representative advocate of this strategy. He argues that moral principles may be 'impartial' in roughly the sense indicated in section iii/ of chapter 1 of this thesis - that is, they may be compatible with P1 - without thereby being 'impartial' in the sense that has drawn the criticisms of the anti-impartialists; that is, without requiring the agent to treat everyone in the same way, regardless of her own interests, relationships, or projects (note that this would be equivalent to interpreting the claim that, to be impartial, the agent must not be swayed by personal considerations, in the second sense that I distinguished in section iii/). Barry calls the former type of impartiality 'second-order impartiality', and the second type 'first-order impartiality', and his claim is that, as he puts it, "second-order impartiality does not entail universal first-order impartiality" (B. Barry (1), p.12). He argues that it is only the presumption that it does do so that has led many to criticise the very idea that morality should be impartial.

Barry himself shows how his distinction between first-order and second-order impartiality might be employed, not only by the contractarian theory he favours, but also by other 'impartialist' moral theories like utilitarianism and what he calls 'Kantianism'. Suppose we take a principle like this:

(A) Each agent is entitled to give a certain degree of precedence to their own children's interests over the interests of other people's children.

Such a principle is clearly not impartial in the first-order sense; it permits you precisely to be 'partial' to your own children, and not because they are necessarily any more worthy or needy than other people's children, but just because they are YOUR children. But it might nevertheless be claimed that (A) would satisfy the criterion of second-order impartiality; that is, that it would be endorsed from a standpoint that took everyone into account equally. We have already seen how a consequentialist might attempt to justify this sort of claim: he could say that the consequentialist end of securing the greatest wellbeing for all would, all things considered, be best served by permitting agents to act on principles like (A), rather than by insisting on their always following a norm of first-order impartiality. And the contractarian too could seek to justify (A), in the following way: he could say that people in an original position of equality would select (A), because a contrary principle could reasonably be rejected by all as being incompatible with the need for a life involving meaningful personal relationships, relationships which inevitably involve a certain amount of 'partiality'. Even the poor will want such personal relationships, and will resist principles that prevent them.

Note, then, that the contractarian rationale for limiting the demands of morality is not, like the consequentialist rationale, that it would be counter-productive in terms of some antecedently given end

to fail to do so; it is rather that it would be unreasonable to expect certain things of people, in a way that everyone can acknowledge. But just as we saw in the previous section that the consequentialist strategy, though perhaps effective against extreme forms of stringency, would be unlikely to yield anything like CSM, so I think we would have to say the same about the contractarian strategy. We can see this by distinguishing two questions about the contractarian strategy: the first is whether it gives us reason to reject a norm of universal first-order impartiality; and the second is whether it gives us reason to reject principles that are considerably more demanding than those of CSM. Barry writes, "What is required is a set of rules of justice (including...both legal and moral norms) that provide everybody with a fair opportunity of living a good life...while leaving room for the kind of discretion in shaping one's own life that is an essential constituent in every conception of the good life" (pp.206-7). But it is pretty clear that current moral and legal norms fail massively to provide everybody with such a "fair opportunity of living a good life", and so those who lose out under present arrangements would apparently be justified in vetoing present norms. The question would then be: how far is it necessary to go, in the direction towards universal first-order impartiality, in order to secure the reasonable assent of all? This is a difficult and interesting question, but one that Barry, fixated as he is (like so many writers in this area) on what he himself calls the "eccentric variants" of impartialism that advocate universal first-order impartiality, fails to give adequate attention to. He does concede that "up to a point, increasing the stringency of first-order

impartiality works to the advantage of the least advantaged groups in society" (p.211); but he continues immediately, "it is an illegitimate extension of this truth to argue that a regime of universal first-order impartiality would therefore be the arrangement most advantageous to the worst-off". But we can agree with this, while wondering whether a good deal more first-order impartiality than CSM requires would, in fact, be to the advantage of the worst-off.

With this in mind, let's formulate another principle that, while not mandating universal first-order impartiality, does take us well beyond CSM. We can use the same example we use above:

(B) The comfortable should be required to give 10% of their disposable income to the needy.

We have already seen how the consequentialist strategy would be unlikely to give us reason to reject a principle like (B): although there may be a level at which to demand more would be counter-productive in consequentialist terms, it is hardly plausible that a principle like (B) would be at that level. But the contractarian strategy doesn't seem to do any better, because one imagines that the needy could reasonably reject a principle that allows the comfortable to go on spending their surplus resources on luxuries while their needs continue to go unmet. Nor does it seem that (B) is a principle that everyone would have a reason to reject in the way that it might be claimed that everyone has a reason to reject (A): although (B) will make considerably more demands on those fortunate enough to be comfortable than CSM does, it does not seem plausible to

claim, as was claimed for (A), that such demands would be incompatible with a life involving meaningful close relationships, or with anything else of comparable importance. Thus, given the desire on the part of the comfortable for principles on which all could reasonably agree, it does not seem that the comfortable could reasonably reject (B).

So we have reached the same conclusion with contractarianism as with consequentialism: although the contractarian strategy may indeed give us reason to reject an EXTREMELY stringent form of morality, it does not seem to give us reason to reject a form of morality that is considerably more demanding than CSM; that is, the second response fails for contractarianism. So it seems that contractarians (as defined here) too, like consequentialists, should be committed to stringency.

Exactly how stringent a form of morality contractarians should be committed to is not easy to establish; in fact, contractarians face an additional problem here that is not faced by consequentialists. For consequentialists, it is, once a particular method of evaluating consequences is accepted, just a matter of getting the sums right; of establishing, that is, what level of stringency will lead, all things considered, to the best consequences. That is a forbidding technical problem, but it is only a technical problem. Contractarians, on the other hand, face a problem of principle: whether there is in fact ANY principle regulating the relations between the comfortable and the needy that everyone can agree to. As Thomas Scanlon himself says, "our different individual points of view, taken just as they are, may in general be simply irreconcilable. 'Judgemental harmony' requires the construction of a genuinely interpersonal form of justification which

is nonetheless something each individual could agree to" (T.Scanlon (1), p.117). The immediate problem for contractarianism, then, is whether there is in fact such a form of justification that would yield principles that everyone could accept; if not, there would be a major breakdown in the contractarian system. But a further point is that the construction of such an 'interpersonal form of justification' thus turns out to be a task that has to be accomplished before we get to the specifically contractarian part of the enterprise, for we will not know what principles can be reasonably rejected until we have such a form of justification that will give an interpretation to the word 'reasonable'. With regard to the moral demands on the comfortable stemming from the needs of others, in particular, we will need criteria for determining what, exactly, it would be 'reasonable' for the comfortable to reject, and an explanation of why those criteria are the right ones. Since I believe that some of the work of Thomas Nagel and Samuel Scheffler can be interpreted as an attempt to answer such questions, now is an appropriate time to turn to them.

iii/ Nagel and Scheffler

Although there are significant differences between the views of Nagel and Scheffler, I think the similarities between them justify my taking them together. (For how the authors themselves see some of the differences, see Nagel on Scheffler in T.Nagel (1), pp.174-5, and Scheffler on Nagel in S.Scheffler (2), p.125. I should also emphasise that the 'Nagel' I will be speaking of is the Nagel of 'The View from Nowhere', not that of 'Equality and Partiality' (nor, for that matter, of 'The Possibility of Altruism'). In 'Equality and Partiality'

Nagel's thought takes a contractarian turn that is lacking, not only in the work of Scheffler, but also in Nagel's own earlier 'The View from Nowhere'. Since, as Nagel candidly concedes, this contractarian element runs into the kind of difficulties mentioned at the end of section ii/ above (see T.Nagel (2), pp.45-52), I will stay here with the Nagel of 'The View from Nowhere'.)

What Nagel and Scheffler have in common are the following elements. Firstly, they both accept that, as Nagel puts it, "there is one important component in ethics that is consequentialist and impersonal" (T.Nagel (1), p.164), and that "such an impersonal element...may become very demanding" (T.Nagel (1), p.189). The reasons why this is likely to be so are already familiar to us from section i/ above. Secondly, they both think that there are non-consequentialist reasons of principle to ameliorate these demands, rather than merely the kind of strategic reasons we looked at in section i/. Like the consequentialist reasons, these reasons of principle are based in certain features of human nature, but unlike in consequentialism, these features are given a directly normative, rather than a merely instrumental or strategic, role in their accounts. While both Nagel and Scheffler agree that human beings are capable of taking up an 'impersonal point of view', from which states of affairs are evaluated regardless of the identities of the persons involved, and that this point of view yields strong consequentialist reasons, they contrast this with another aspect of our nature, a 'personal point of view', from which, to quote Scheffler, "concerns and commitments are NATURALLY generated quite independently of the weight of those concerns in an impersonal ranking of overall states of affairs"

(S.Scheffler (1), p.9).

What are the grounds for giving individuals' possession of such a personal point of view a normative role, rather than merely the kind of strategic role that consequentialism would give it? Scheffler puts it in this way: "if, as Rawls has said, 'the correct regulative principle for anything depends on the nature of that thing', we must surely reject any regulative principle for persons which ignores the independence of the personal point of view" (S.Scheffler (1), pp.57-8). Nagel argues similarly: "even though morality has to emerge from an impersonal standpoint, that standpoint must take into account the kind of complex beings for whom it is being devised. The impersonal is only one aspect of their nature, not the whole of it. What it is reasonable to ask of them, and what is impersonally expected of them, should reflect this" (T.Nagel (1), p.202). So the basic claim is that the nature of the type of creature with which we are concerned is relevant in determining what it is REASONABLE - not merely expedient, as in consequentialism - to ask of it. If we were impersonal evaluators and nothing more, maybe consequentialism (or something like it) would be appropriate; but since our nature is more complex, morality should reflect that fact.

Although this is certainly an appealing idea, I am not sure that either Nagel or Scheffler say enough to support it. Both of them tend to appeal to the idea that, if we stand back from our immediate position in the world, and consider how human beings, in general, should live, we will just see that, given certain aspects of their nature, it is unreasonable to demand too much of them. But the problem is that not everyone seems to see this. The consequentialist, in

particular, will not deny that we possess features that lead us to resist bringing about the best consequences, impersonally construed, but she regards our possession of such features as an unfortunate impediment, to be accommodated strategically, perhaps, but not to be accorded any independent normative role. The debate between Nagel and Scheffler and the consequentialists, therefore, is in danger of collapsing into a stalemate of opposing intuitions about the nature of morality.

I believe, however, that we can draw some support for the Nagel/Scheffler approach by showing how it, too, is compatible with P1, and, indeed, arguably represents the best interpretation of P1. Remember that consequentialism ensures that everyone is taken into account in the same way, at the fundamental level, by counting their interests equally in the evaluation of consequences that is uniquely criterial for moral evaluation, and that contractarianism does so by requiring that moral principles be agreed upon by all from a position of equality. Since the Nagel/Scheffler approach claims there are reasons of principle (rather than just the strategic reasons that consequentialism allows) for departing from the consequentialist standard, and is not based on the contractarian idea of universal agreement, it may not be immediately obvious how the Nagel/Scheffler approach is compatible with P1. But I think it is; let me try to explain how.

Remember that the idea behind P1 was that whatever properties there are that may be considered to be fundamentally relevant, they are likely to be properties that are universally shared. But it is important to recognise that different properties might be relevant in

different ways; in particular, while some properties might be relevant to our role as (for the want of a better word) PATIENTS, that is, to the question of how we can legitimately expect to be treated by others, others might be relevant to our role as AGENTS, that is, to the question of what we can legitimately be expected to do for others. (I don't want to claim that ALL moral relationships should be seen in terms of the agent-patient relation, but clearly some can, and in particular this seems to be the appropriate type of relation in terms of which to consider the issue of the comfortable and the needy). Some moral theories, particularly those that lay the main emphasis on the question of how people should be TREATED, seem to me to neglect the role that we play as agents; which is odd, if we remember that at least one main function of moral principles is, as it were, to connect agents and patients. Moral principles must determine not only how people should be TREATED, but also how people should ACT, and if people, quite generally, possess properties, distinctively in their role as agents, that are fundamentally morally relevant, and affect what it is reasonable to ask of them, then it seems that morality should give an independent normative role (reflecting the independent role we play as agents) to those properties in determining which moral principles are acceptable. Of course, this is not to determine that there ARE any such properties; the point here, rather, is that IF there are, then it is not incompatible with P1 - indeed, it will be essential to an adequate interpretation of P1 - to take such properties into account.

I believe that it is the failure to take human beings distinctively into account in their role as agents that lies behind much of the

vague sense of dissatisfaction with consequentialism. All that is relevant, for the consequentialist, is how good the consequences are, impersonally construed. The agent counts equally with others affected by what he does, but THE FACT THAT HE IS THE AGENT is treated as irrelevant: all that matters, for him as for any others affected by what he does, is whether his level of wellbeing (or whatever) goes up or down as a result of it. To think in this way is to exclude the possibility that there may be properties of human beings AS AGENTS that have moral significance in their own right, independently of how the overall balance of consequences is affected, and independently of the kind of strategic significance that sophisticated forms of consequentialism may give them.

It seems to me to be a virtue of contractarianism that, unlike consequentialism, it does implicitly keep the agent's point of view in focus. If we are asked to choose moral principles from some original position of equality, we naturally take into account both our role as agents (being obligated to act on the relevant principles), and our role as patients (how we will be affected by others acting on those principles). In selecting moral principles we in effect shuttle back and forth between our role as agents and our role as patients in trying to determine whether a given principle is acceptable or not: as patients, we may prefer highly demanding principles, given that we may find ourselves in a position where on such principles will ensure us a decent life; but as agents we may resist such demanding, because of our desire to lead a life that is not overwhelmed by moral demands. We have already seen some of the problems involved in such a procedure, but for now I just want to emphasise that contractarianism takes the

agent's point of view fully and independently into account. It does not collapse the two roles of patient and agent into one currency of 'consequences'.

So while consequentialism takes everyone into account in the same way, at the fundamental level, by giving their interests equal weight, and contractarianism does so by giving them an equal say, from a position of equality, in the selection of moral principles, the Nagel/Scheffler approach, as I interpret it, does so by taking whatever properties may be morally relevant directly into account in deciding what it is reasonable for people to do and to expect. And their claim is that when we do so, we shall be able to see that human beings have properties - in particular, the possession of a 'personal point of view' - that are universally shared, and fundamentally morally relevant, which would lead us to say, quite impartially, that it is unreasonable to expect too much of them. Unlike consequentialism, these properties are given an independent normative status in the theory, and unlike contractarianism, what should be expected of agents is to be established directly, from a kind of 'impartial point of view', rather than via a process that requires the assent of all.

This seems to me, as I shall try to show, the interpretation of P1 that is both most convincing in itself, and most likely to give us a way of framing the problem of demandingness that will enable us to make progress with the problem. But, even though this account is derived from the work of Nagel and Scheffler, it seems to me that there are serious defects in their versions of it, and I should say something about these first.

Nagel and Scheffler claim that there are features of human nature, features that lead us, as agents, to resist a highly demanding form of morality, which should be given an independent normative role in our moral theory. If that is so, then we will need, firstly, a satisfactory account of what the criteria are by which we decide whether any given feature of human nature should be accorded this kind of normative status, and, secondly, an account of what those features are. Unfortunately the treatment given by both Nagel and Scheffler of both these matters is deeply unsatisfactory.

I'll take the second problem first. The features that are supposed to render it unreasonable to expect too much of agents are characterized generally as the possession of 'personal' motives or concerns, or of a 'personal point of view', where the sense of 'personal' here is understood by reference to, that is, in contrast with, 'impersonal' motives or concerns, or the 'impersonal point of view'. Note that, introduced in this way, the question of the specific CONTENT of these concerns is left open: to form part of the 'personal point of view' all that is necessary is that these concerns are weighted differently by the individual concerned from the weight they would have from an impersonal point of view. This is because the impersonal point of view is characterized first, and then concerns - ANY concerns - that don't correspond in importance for the individual to the importance they would have from the impersonal point of view are ascribed to the personal point of view.

'Personal' concerns, then, would include at least the following:

i/ self-interested concerns;

ii/ concerns that one would normally call 'moral', but whose importance for the individual doesn't correspond to their importance from the impersonal point of view (for example, concerns for one's own family or community); and

iii/ concerns that would normally be classed as IMMORAL ones, such as the concern that someone who has crossed you should suffer.

Both Nagel and Scheffler seem insufficiently aware of how this complexity might affect their theories. Both of their formulations are most commonly in terms of the agent's 'interests', which could be read as a kind of shorthand for all of these types of concerns (understood, perhaps, as 'what the agent is interested in'), but is perhaps more naturally read as i/, that is, as self-interest. (Thus, for example, Scheffler defines his 'agent-centred prerogative' as allowing "each agent to assign a certain proportionately greater weight to his own interests than to the interests of other people" (S.Scheffler (1), p.20); and Nagel writes, "It isn't just that my personal interests will cause me to rebel against impersonal demands - though that may happen. It's that this resistance will get some support from the objective standpoint itself" (T.Nagel (1), pp.201-202)).

Given that the idea of the 'personal point of view' encompasses such a wide variety of types of motives and concerns, some of which would clearly not be given any normative status by anyone, it cannot be that we should, as Scheffler claims, grant 'moral independence' to the personal point of view as such (S.Scheffler (1), pp.61-62). The notion of a 'personal point of view' or of 'personal' concerns is so general and indiscriminating as to be completely unhelpful as a way of identifying the features of human nature that might be given an

independent normative status in our moral theory. After all, mean, cruel, and vindictive motives and concerns seem, to be pretty invariant ("naturally generated") features of human nature, and, by not being 'impersonal', they would seem to qualify for Nagel and Scheffler as 'personal', but no one would suggest they should thereby be given any normative status. So we will need to be much more discriminating and specific than either Nagel or Scheffler are about the types of 'personal' concern that we accept as candidates for having normative status. We will need to look carefully at the different types of 'personal' concerns, evaluate them one by one, and decide on their normative status, if any.

But even if we do this, it does not appear, to go back to the first problem, that Nagel and Scheffler give us any adequate criteria for determining WHICH personal concerns should be granted normative status. I have already argued that if there are properties of human beings that, as agents, lead them to resist a highly demanding form of morality, and those properties are universally shared and taken to be fundamentally moral relevant, then it will not contravene P1 to give those properties an independent normative role in our moral theory. But that possibility does not, by itself, show that there ARE in fact any such properties, nor give us any criteria by which to determine whether each of the possible candidate properties should, in fact, be given such an independent normative role. Nor does it seem to me that Nagel and Scheffler offer any such adequate criteria. Once again, the idea of taking a more 'objective' point of view seems insufficient, by itself, to providing any such criteria. Sometimes Nagel and Scheffler attempt to support their view by observations about how people

typically behave or are typically motivated, but it is difficult to see how such facts about what is typical, alone, can support the granting of normative status to those features of human nature. For example, Scheffler writes that "people do not typically view from the world from an impersonal perspective, nor do their actions typically flow from the kind of concerns a being who did inhabit the impersonal standpoint might have" (S.Scheffler (1), p.62). But, as Shelley Kagan says in response, "thus, faced with the fact that people DON'T typically promote the overall good, (Scheffler) responds that morally they're not required to. What is the underlying rationale for this response? Personal independence may constitute an implicit appeal for agent-centred prerogatives - but what is the rationale for granting this appeal?" (S.Kagan (1), p.253). It is this, I think, that has led some to suppose that Nagel and Scheffler are merely granting a moral imprimatur to what is going to happen anyway. While we have seen that there is at least theoretical room, compatibly with P1, for giving an independent normative role to certain features of human nature that lead us to resist highly demanding moral principles, that does not by itself give us criteria for determining what those features are. And the idea that we can just read off those features from how people are typically motivated suggests, indeed, that we are merely capitulating to the status quo.

These are serious problems, and until we have resolved them it will be difficult to say whether the Nagel/Scheffler approach, unlike consequentialism and contractarianism, provides considerations that counterbalance those that support stringency, in such a way as to vindicate something like CSM. Certainly, to the extent that the

suggestion is that we can limit stringency by an appeal to how people are typically motivated, or what sort of patterns of concerns they typically have, this kind of approach would indeed seem likely to deliver something like CSM - but only at the cost of failing to seem like a legitimate moral argument.

I do think, however, that there might be a way to remedy the defects that I have outlined in the work of Nagel and Scheffler. I shall say more about this in chapter III, but first I shall consider briefly three other responses to the problem of demandingness.

iv/ Other Responses

Even though I cannot, of course, comment on all the responses to the problem of demandingness, there are three other responses that I would like to say something about, however briefly.

The first is to agree with Nagel and Scheffler in giving certain considerations that lead us to resist extremely demanding principles a normative role, but to assign them to another area of practical reason, independent of morality. If the former types of considerations are held to yield good reasons that are not always overridden by moral reasons, it may then be claimed that while morality is, indeed, stringent, all-inclusive practical reason is not, because we may have better reason, all things considered, not to act on the moral reasons. As the point is sometimes put, we should react to the problem not by modifying the CONTENT of morality, but by limiting its RATIONAL AUTHORITY.

This kind of line can be found in the work of Henry Sidgwick, who identified morality with utilitarianism, with which he contrasted an

independent area of practical reason, rational egoism (hence Sidgwick's "dualism of practical reason": see H.Sidgwick (1) pp.xii-xiii). A similar position has recently been advocated by Roger Crisp (R.Crisp (1)). This version of the response is subject to the objection that it is doubtful whether all legitimate types of reasons can be squeezed into one or other of these categories; what about limited, but genuine altruism, for example - the type one may show to a friend - which cannot be justified on either utilitarian or egoistic grounds? In response to this specific objection, further types of reasons may be added, as they are, for example, by Susan Wolf, who distinguishes reasons given from "the point of view of individual perfection" (S.Wolf (1), p.437) from both 'moral' and egoistic ones, though once one is in the business of noticing different types of reasons, it is not clear why one should stop at three, or four, or... But whichever version of this response is selected, I think it faces some serious problems.

In the first place, if it is indeed the case, as I have suggested it might be, that we can get the properties of persons that lead them to resist a highly demanding form of morality into view from an impartial perspective, one that is compatible with P1, then it seems that those properties are indeed relevant to morality, rather than to some other area of practical reason. If, furthermore, as Nagel and Scheffler would claim, it is NECESSARY for morality, specifically, to take these properties into account, then a form of morality that failed to do so (such as Sidgwick's or Wolf's 'morality') would to that extent be defective. Secondly, it can be forcefully argued (as it is by Scheffler in Scheffler (2)) that the way in which we actually talk

about morality takes these features into account. If we want to preserve at least some continuity with our ordinary discourse, then, we should consider them to have direct moral relevance, rather than consigning them to some other area of practical reason beyond the boundaries of morality. Thirdly, presumably just about everybody acts consistently against what 'morality', as construed by Sidgwick or Wolf, would suggest; they would say that we have other reasons that authorize our doing so, but in doing so, we do, nevertheless, not act on the moral reasons. So it would seem that we consistently fail to act on the moral reasons. But, if we expect a life that is at least generally responsive to moral reasons to be achievable by more than a handful of saints, then it would seem to make better sense to include the considerations that support those countervailing reasons WITHIN morality. And finally, the 'other reasons' response seems to give us no adequate criteria for deciding between the different types of reasons, moral and non-moral ones. We seem to be left with incommensurable types of reasons corresponding to different aspects of our nature. Hence Sidgwick's famous 'despair', and the unhelpful references to 'intuition' on the part of Wolf (p.439) and Crisp (p.65). If we CAN get all the relevant types of considerations into view from one general perspective, then surely we would do better to do so. I shall have more to say about how this might be possible in chapter III.

The second type of response that I should mention is to say that it should not, in any case, be the responsibility of individuals acting in a private capacity to respond to the needy; it should rather be the

responsibility of national and international institutions to create the conditions in which all people's needs are met in an equitable way, so that private charity is no longer necessary. Though there is no doubt much in this, it doesn't, unfortunately tell us what to do, given that we do NOT live in a just world order. In particular, does the claim that in an ideal world responding to the needy would not generally be a responsibility of private individuals mean that in the far from ideal world in which we live it need not be? I do not see how it does. It may change what we feel to be the most appropriate way to respond to the problem: we may feel that, instead of private charity, we should concentrate on trying to change the relevant political structures, or trying to influence the prevailing culture; or maybe we should do all of these things. But surely the fact that the needy are out there right now, and that we are able to help them, if we choose to do so, means that we still face an ethical challenge.

So I do not see how the claim that the problem is, at root, a political one removes the more narrowly ethical dilemma that we face now. Focusing on what kind of political conditions WOULD remove or at least ameliorate the problem cannot, by itself, tell us what we should do, as individuals, now. Nagel says that "we must change the question from 'How should we live, whatever the circumstances?' to 'Under what circumstances is it possible to live as we should?'" (T.Nagel (2), p.52); but it seems to me that we should not let an emphasis on this, admittedly very important, question distract us from trying to give an answer to a third question: 'How should we live in our present, far from ideal, circumstances?'

Similar points seem to me to apply to another response, one that claims that what each of us, as individuals, should be required to do for the needy should be limited to what each of us would have to do if everyone else did their fair share (for different examples of this response, see B.Hooker (1) and L.Murphy (1)). (Similar points apply because of the structural similarity of this response to the political one: that is, the strategy of appealing, in some form, to how things ideally should be to determine what we should do in our present, far from ideal, circumstances.) Although it may be that, if all of the comfortable contributed, the demands on any particular one of them would not be excessive, but, once again, the immediate question for us now is what we should do in a world in which most others do NOT contribute. Both Hooker and Murphy discuss an adaptation of Singer's famous 'drowning child' case (see P.Singer (1)) in which the failure to do more than what would be required of one if everybody else did their bit would be disastrous: two drowning children, two passers-by, only one of whom responds. If the limits of the responsibilities of the other passer-by are set by what would be required of him if everyone did their bit, then he would not be obligated to go on and save the other child: a claim that ~~on~~ one wants to endorse. But if one is going to require more of the passer-by in this instance, why not require more of the individual comfortable person when other comfortable people do not respond to the distant needy? At this point Murphy appeals to the idea that one's direct confrontation with someone creates a special obligation, a claim that (as Hooker has pointed out) will surprise those who are directly confronted with the destitute on the streets of Calcutta (or London). Hooker's response is

to appeal to a version of the consequentialist 'counter-productiveness' strategy: he claims that, given the "natural biases and limitations" of human beings (B.Hooker (1), p.76), and the consequent costs of trying to inculcate a highly demanding code, it may be cost-effective to select a less demanding moral code, one which requires more of us in rare cases (such as that of the drowning children) than in everyday cases (such as that of the distant needy). I am not sure how stringent a form of morality this version of the consequentialist strategy would lead to, but in any case it is subject to the objection that I have already made of the consequentialist strategy: that there is no reason OF PRINCIPLE why we should not expect anything at all of moral agents. If Nagel and Scheffler are right in saying that there are such reasons of principle, then we would not have to rely exclusively on the consequentialist strategy. I will move on now to try to explain how I think the appeal to human nature can be made in a form that does deliver such reasons.

CHAPTER III: THE MORAL AGENT

i/ The Idea of an Agent-centred Conception of Excellence (1)

Let us return, then, to the debate between Nagel and Scheffler, on the one hand, and the consequentialists, on the other. Recall that both parties accept strong, bona fide consequentialist reasons, reasons to bring about the best consequences, impersonally construed, and that both ~~both~~ parties agree that there are features of human nature that make us resist acting on such reasons. The point of dispute is whether any of those features should be given any independent normative status within our moral theory; while Nagel and Scheffler affirm this, the consequentialist denies it, treating those features as unfortunate impediments, to be accommodated, perhaps, for strategic reasons, but not to be given any normative status in their own right. I argued, firstly, that giving certain of those features such an independent normative status need not, in itself, contravene P1: given the role that we all play as agents, it seems that morality should take everyone into account in the same way specifically as agents, and if human beings, quite generally, possess properties that have normative status distinctively in their role as agents, then morality should take those properties too into account.

But to show that there is this theoretical possibility is not, in itself, to show that human beings DO possess any such properties, properties that have the kind of independent normative status that Nagel and Scheffler claim. In order to show this, it would be necessary to establish acceptable criteria for determining whether any particular property possesses this kind of normative status, and then

to show that certain properties satisfy these criteria. I then argued that Nagel and Scheffler fail to provide any such adequate criteria: they both point to the fact that human beings, quite generally, have 'personal' concerns, but that is not enough to show that any of those concerns should be given a normative status. Indeed, when we recall that the sense of 'personal' is given by contrast to that of 'impersonal', and that therefore among the 'personal' concerns will be types of concern (such as the concern to hurt someone who has crossed you) that would not be given any normative status by ANYONE, the question of how we are to decide which types of 'personal' concern should be given any normative status becomes all the more acute.

I also claimed that, although Nagel and Scheffler fail to provide any adequate criteria, there may be a way in which such criteria can be established. My suggestion is this: that we determine which features of human nature to give an independent normative status to by reference to an AGENT-CENTRED MODEL OF EXCELLENCE ('AME' for short); that is, a model of what would constitute the best expression of an individual's nature, of her human potentialities. The reference to the agent's 'nature' is intended to indicate that the model in question must be responsive to any relevant features of the agent, including, indeed, her possession of a 'personal point of view'; but it does not follow that the model need grant all of those features normative status. It will, rather, grant normative status to those features that are compatible with what is considered to be a kind of life that realizes the best potentialities of that person's nature. Thus we will not be in the position, as Nagel and Scheffler apparently were, of rubber-stamping certain types of concern just because they are

'typical', or 'naturally generated'; we will endorse them, rather, only if and insofar as such types of concern have a place in what we consider to be the best achievable kind of life for a person with such a nature.

Let me try to illustrate how such an agent-centred model of excellence can be relevant to determining which features of human nature to give normative status to with the help of an example. Karl Marx, in 'On the Jewish Question', argues against the 'bourgeois' conception of liberty in the following way: "the liberty we are here dealing with is that of man as an isolated monad who is withdrawn into himself... (The) right of man to freedom is not based on the association of man with man but rather on the separation of man from man. It is the RIGHT of this separation, the RIGHT of the restricted individual, restricted to himself" (K.Marx (1), p.229). Now, confronted with this feature of human nature (that is, the desire for separation), Nagel and Scheffler would presumably have to say that it is one of the range of personal desires that has to be respected, and balanced against the impersonal point of view; but Marx questions whether we should give normative status to this feature of human beings, and NOT, of course, on the consequentialist grounds that it prevents the realization of the best consequences, but rather by reference to a certain AME. He rejects the normative recognition of such a desire for separation on the grounds that it reinforces a debased model of human life, one in which man is separated from man. His idea is that the 'connected' man exemplifies a higher, better form of being, and he determines which features of human nature to give normative status to by reference to that ideal. I am not, of course,

concerned here with whether Marx's model of excellence is the right one; the point here is rather to illustrate the general claim that which features of human nature we grant normative status to will depend on our model of excellence, and that we will need some such model to ground our normative claims. It is not enough, pace Nagel and Scheffler, to point out that a certain pattern of concern is typical, nor to appeal to what strikes us as 'reasonable' on reflection, or from a more objective point of view; in order to grant normative status to a certain trait, rather, we will need to be convinced that it is the RIGHT, or the APPROPRIATE, way for human beings to behave, given the potentialities of their nature, and in order to be convinced of this we shall have to appeal, either implicitly or explicitly, to an AME.

What should a consequentialist say about this? It seems to me that she has two choices. Either she can accept the moral relevance of an AME, and then try to make the case for a highly extreme model: that the best fulfillment of our nature is to be perfect, impersonal good-maximisers; or she can, in the kind of way indicated in section iv/ of the last chapter, deny the distinctively MORAL relevance of an AME. It doesn't seem to me that either option is very attractive.

On the first option, although, as I shall argue in section iii/ of this chapter, I do think that a strong case can be made for an AME that makes moral requirements on an agent that are considerably more demanding than those of CSM, I do not think that this case can be pressed so far as to support the idea that the best fulfillment of our nature is to be such that we maximise good consequences, impersonally construed. Perhaps such an AME would be appropriate to a creature that

was a pure benevolent will, and no more, but it does not seem appropriate to the creatures of flesh and blood that we are. Turning oneself into an instrument of the general good is likely to seem, quite literally, dehumanizing, rather than the best fulfillment of our humanity. (Incidentally, I think much of the pressure for this type of AME comes from the kind of equivocation over the notion of the 'moral point of view' that I described in section iii/ of chapter I. Even if it is the case that the content of morality should be determined from the point of view of an 'ideal spectator' or 'impartial observer', it does not follow that the INDIVIDUAL MORAL AGENT should take up that point of view in her ordinary moral reasoning, in such a way that she should try to act as such an 'ideal spectator' or 'ideal observer' would. Rather, the individual moral agent should act on the principles that emerge from that point of view, principles that might, if the 'spectator' really is an 'ideal' one, acknowledge the distinctive role that human beings play as agents - and, in particular, the fact that human beings are NOT merely 'spectators' or 'observers'.)

There has been a lot of discussion in the literature about the degree to which this consequentialist AME is an achievable one, about whether it is POSSIBLE for human beings to be maximisers of the good, impersonally construed. But another, perhaps more important, question is whether this is an APPROPRIATE model of excellence for human beings, whether, on reflection, we think it is the finest, or most fitting expression of the distinctive nature of human beings to be such good-maximisers. If we do NOT think that this is an appropriate ideal, then, whether or not it is actually possible, we will have to acknowledge that there are reasons for agents not to pursue it. And if

- to turn to the second option mentioned above - the acknowledgement of those reasons is, as I have suggested it might be, compatible with P1 - in particular, with taking everyone into account in the same way in their role as agents - then we should regard those reasons as legitimate moral reasons. It is not that, in taking account of the nature of agents, we are going outside morality to some other area of practical reason; it is rather that we are taking account, in an impartial way, of properties that are fundamentally MORALLY relevant. But if that is so, then we will have to accept an independent constraint on moral principles: that they be evaluated in terms of the degree to which their being acted on by an agent is compatible with that agent leading the kind of life that we consider to be the finest expression of her nature.

Take the case of close personal relationships, for example. The consequentialist will not, of course, deny that such relationships can be of great value, but will, nevertheless, view the pursuit of such relationships negatively, morally speaking, insofar as that pursuit leads to worse consequences, impersonally construed, than an alternative orientation on the part of the agent. She may, as we have seen, accommodate that pursuit on strategic grounds, permitting people to pursue such relationships to the extent that doing so, given certain features of human nature, will lead to better consequences, but she will regard that, from the moral point of view, as a regrettable concession to hard fact. Nagel and Scheffler, on the other hand, want to morally justify the agent's pursuit of such relationships on independent normative grounds, rather than just on the kind of strategic grounds that consequentialism allows. But the

problem is that, although we may find this intuitively plausible, Nagel and Scheffler make it hard to see what such independent normative grounds could be. When Nagel, for example, speaks of a "modification in the demands of impersonal morality...based on tolerance and the recognition of limits" (T.Nagel (1), 201), and of striking "a bargain between our higher and lower selves" (T.Nagel (1), 202), it looks as though all the normative force is on the side of consequentialism, and that we are just reducing the demands of consequentialism in the spirit of a kind of defeatist realism about the potentialities of human nature - a move that may not even have the kind of strategic justification that the consequentialist move has. No wonder so many people have thought that the Nagel/Scheffler approach is more of a sop to our moral weakness than a legitimate moral argument. But if, instead, we justify the agent's pursuit of personal relationships by reference to an agent-centred model of EXCELLENCE, then we are justifying it by reference to something that is ALREADY normative, and so there is no mystery about where the NORMATIVE recognition of such a pursuit comes from. We will not endorse the pursuit of close personal relationships, if we do so, just because that is how people typically behave, and because we feel it is vain to hope for anything else; we will endorse it because that is the FITTING, or the APPROPRIATE, or the RIGHT way for people to behave, given their nature. And since this recognition is compatible with P1, and based on properties that are plausible candidates for having fundamental moral relevance, we can endorse it from a perspective of moral impartiality.

Even if our pursuit of close personal relationships does prevent us

from realizing the best consequences, impersonally construed, it does not follow, then, that it is compulsory to regard this trait, morally speaking, as one of those unfortunate "natural biases and limitations" that Hooker, for example, speaks of (see Chapter II, section iv/). If the normative recognition of such a trait is compatible with P1, and can be justified, in particular by reference to an appropriate AME, then we need not use the pejorative word 'bias' to characterize such a trait, even if it prevents us from realizing the best consequences. Nor need we be forced to see the currency of such a trait, morally speaking, as a 'limitation' - something which we would like human beings to get beyond, if they only could - we might see it instead as a wholly fitting and morally justifiable expression of our nature. We will indeed regard SOME typical human traits as regrettable, but we will not, like the consequentialist, be compelled to regard ALL traits that prevent us from maximising the good in this way.

It is on this territory, then - on the issue of which AME we should accept - that I think Nagel and Scheffler should be fighting consequentialism. And shifting to this territory may indeed break the deadlock in the sense of setting before us an issue - which AME is most appropriate - that is both central to the problem of demandingness, and one that everyone ought to be able to discuss together. The opponents of consequentialism are likely to claim, as I have said, that the consequentialist AME is not appropriate to our nature, and consequentialists will face the difficult job of trying to rebut this claim. And other moral theorists too will be able to join the fray, without contentious premises that cut off their debates from one another. The contractarian will be able to employ of the idea of

an AME in providing criteria for that 'interpersonal form of justification' that will give a shared sense of 'reasonable' which, as we saw, was a necessary precondition on getting the specifically contractarian part of the enterprise going: what it will be 'reasonable' for agents to reject, in particular, will be determined by reference to an appropriate AME. And the critic of moral theory too, as I shall argue in section iv/, should welcome the introduction of the idea of an AME as bringing into focus the kinds of issues that she accuses certain moral theories of neglecting.

Not that progress will then be easy. The question of which AME is appropriate is, by any standards, an exceedingly profound and difficult one, and no doubt it is at least in part for this reason that many moral theorists have tried to shift the ground to some other, more tractable issue. But this seems to me to be a kind of false economy: in the end we will have to face this issue in order to determine both the content of morality, in general, and the limits that should be set on moral demands, in particular. I shall go on to say some more about the idea of an AME, and the kind of role it might play in a general moral theory, in the next section.

ii/ The Idea of an Agent-centred Model of Excellence (2)

As I have said, since an AME is normative, we cannot just read it off from how people actually behave; but, on the other hand, if it is not just to be an empty fantasy, it will have to be constrained by certain factors, including the following:

a/ Certain invariable or relatively invariable features of human nature. To mention just a few examples: our biological nature; our

embodiment; our affective, conative, and cognitive potentialities and limitations.

b/ Features of human nature that vary from culture to culture: the model of excellence that is appropriate for a Masai tribesman is hardly likely to be the same as that which is appropriate for a Western city-dweller.

c/ Features of human nature that vary from individual to individual, even within the same culture; for example, different individuals possess variable temperamental traits, and variable talents, that may be taken to affect the specialization of human excellence that is appropriate for them.

Just from this very sketchy list we can see that both the idea of an appeal to 'features of human nature', and that of a 'model' of excellence, should be handled carefully. On the former point, when determining an appropriate AME for a given individual or for individuals in a certain culture, we will have to take into account not only even variable features of what is normally called 'human nature', but also features of the social world in which those individuals exist. This is another reason why the bland, indiscriminating appeal on the part of Nagel and Scheffler to 'personal' concerns or the 'personal point of view', as a way of substantiating the claim that it is unreasonable to expect certain things of moral agents, is inadequate. When trying to determine what it is reasonable to expect of people, it will be necessary to take into account MUCH MORE than a general tendency on the part of agents to weigh some concerns differently from the weight they would have from an impersonal point of view.

On the latter point, it should be clear that the appeal to 'a model' of excellence is not intended to imply that there is only one model that we need to take into account, appropriate to all times, places, and people. For ease of exposition I will continue to talk of 'an agent-centred model of excellence' that is responsive to relevant 'features of human nature', but the use of these phrases should not be taken to imply that there is only one legitimate model, or that any such model should not be responsive too to features of the social world in which agents live.

When all of these relevant features have been taken into account, the task will be to determine what, given those features, would be an excellent life for one of those individuals, what would be the best fulfillment of their nature: this will then function as an independent constraint on the selection of moral principles, helping to determine, inter alia, what it is reasonable to expect of particular moral agents. In calling the model of excellence 'agent-centred', I should emphasise, I am insisting that whatever moral principles are selected will have to be grounded in, and appropriate to, the nature of the agent; this requirement follows from the fact that the role of (at least some) moral principles is, as I put it above, to connect agents and patients, and the consequent requirement that we take into account any properties that may be relevant to our role as agents, as well as our role as patients, in deciding what principles to endorse. It does NOT necessarily follow that the principles endorsed have to be OPTIMAL from the standpoint of the agent, in the sense of ensuring, if the principles are followed, the best kind of life FOR HER. The point of bringing in the agent's point of view is not to crudely pit the

agent's interests against the interests of others, and the 'agent's point of view' as I use the phrase is not, emphatically not, simply the point of view of the agent's self-interest.

Once again, the possibility of this kind of approach is often obscured by the constant emphasis, in Nagel and Scheffler, and other writers, on the agent's interests as a consideration that will help to reduce the stringency of morality. It then seems that the suggestion is that since human beings, quite generally, care more about their interests than about the interests of others, morality should reflect that fact by allowing them to favour their own interests. And once again the question is: WHY should morality do that? Just because that's how people generally are? And once again the best response to this stalemate, it seems to me, is to appeal to an agent-centred model of excellence: we will permit a certain favouring of her own interests, on the part of the agent, in as much as we do so, and if at all, by reference to a model of what would constitute an excellent life on the part of that person, given her nature, of what would constitute the finest expression of that nature. Considerations of the agent's caring more for her own interests will become relevant, not because it is just obvious (to some, but not to others) that morality must normatively endorse this trait, but because, and to the degree to which, such a trait is appropriate to, or fitting with, our best-considered model of excellence.

So the general picture is this: that we may be able to determine acceptable moral principles by reference to an AME, which is not simply a conception of what would be in the agent's interest, AS WELL AS by considerations that reflect the patient's point of view, AND

that we are able to get BOTH sets of considerations into view from a position of moral impartiality. It seems to me that the possibility of such a position is frequently missed, for a number of reasons. Most commonly, the potential relevance of an AME is not realized, and the problem is then seen as simply one of the agent's interests versus the requirements of an 'impartial' morality. Some philosophers then take the agent's interests as primary, in a way that makes considerations of impartiality seem deeply problematic, while others take a certain conception of impartiality as the indisputable starting point of morality, in a way that makes considerations concerning the agent's interests similarly problematic. Others attempt a kind of ad hoc, unsatisfying accomodation between the two sorts of consideration. Then there are some philosophers who do hit upon the idea of an AME, but either, firstly, fail to see the kind of role it can play; or, secondly, fail to see its distinctively MORAL relevance; or, thirdly, fail to see how considerations reflecting it are compatible with moral impartiality.

On the first point it is notable that Railton himself, in an early passage in Railton (1), draws upon the idea of what I call an AME when he writes that "some of the very 'weaknesses' that prevent us from achieving this moral ideal (of 'sub specie aeternitatis abstraction') - strong attachments to persons or projects - seem to be part of a considerably more compelling human ideal" (P.Railton (1), pp.99-100). But when Railton goes on to attempt to justify departures from a life of impersonal welfare-maximixation he does so, as we have seen, on the grounds that to try to live such a life would be counter-productive in consequentialist terms, rather than arguing that such a life is

incompatible with a "more compelling human ideal". In this way he seems to me to fail to see the kind of role that the idea of an AME might play in a moral theory.

On the second point, some philosophers do, as I have said, recognise the relevance of something like an AME to our practical reasoning in general, but seem to me to fail to recognise its distinctively MORAL relevance. Susan Wolf, who^m I discussed briefly in section iv/ of the last chapter, is a case in point: her idea of a "point of view of individual perfection" (S.Wolf (1), p.437) has some affinities with the idea of an AME, but, as we have seen, she assigns the reasons it gives to an area of practical reason independent of morality. The idea is that what morality requires is already fixed, and then we have to take into account competing kinds of reasons. But in my picture, morality does not have a content that can be fixed without taking into account an AME. And this is because, as I have said, the appropriate content of morality is determined, following P1, by a process that takes everyone into account, and, given that one function of morality is to connect agents and patients, this will include taking everyone into account as agents, and so any moral properties that people possess distinctively in the role of agents will also have to be taken into account. Since which properties of agents ARE morally relevant is, I claim, to be decided by reference to an AME, it cannot be that we can fix the content of morality independently of reference to such an AME. So the appeal to an AME will not give us non-moral reasons for limiting or counterbalancing the pre-set demands of morality; rather, it is central to determining the content of morality itself.

On the third point, David Wiggins has argued that a Humean "public

standard determined by the common point of view might be one that stressed not utility (in the sense of the greatest public happiness or whatever, thus requiring that the identities of putative recipients of it be treated as irrelevant) but, say...a policy of 'self-referential altruism'; and that "the public standard need not require these kinds of conduct because each one's looking after his own is taken to be the most efficient way of maximizing 'general happiness'...but because of some relatively unsystematic connexion linking them with a historically conditioned, agent-centred ideal of human personhood and human excellence" (D.Wiggins (2), pp.61-62). Here we have a clear statement of the idea that the content of MORALITY, specifically, might be determined, at least in part, by reference to the notion of an AME; but the problem this time is that there is no insistence that the 'common point of view' be one of what I have called 'moral impartiality'. Without this claim it will always be possible that the public standard systematically favours the powerful; it is true that, on Wiggins' Humean approach, the weak too must be expected to concur on the standard in question, but, without the condition of moral impartiality (that is, that everyone is taken into account IN THE SAME WAY), it will always be possible that they do so under ideological pressures. (For example, the weak may agree that it is fitting for those of 'noble' birth to have access to certain priveleges, but do so only because they have been bedazzled by an ideology that systematically favours the powerful.)

Now, even if the type of picture I have suggested was accepted, and even if some consensus was reached on what AME is appropriate, there would still remain the difficult issue of balancing the agent's and

the patient's points of view - for the idea, as I have put it forward, is that an AME will set an INDEPENDENT CONSTRAINT on the selection of adequate moral principles, not that it will REPLACE considerations that reflect the patient's point of view. I have been emphasising the agent's point of view, because it seems to me to have been inadequately taken into account by consequentialism, in particular, but the suggestion is not that we REPLACE an 'impersonal' or 'agent-neutral' morality with an 'agent-centred' one. Some philosophers who have been impressed by the importance of agent-centred considerations have indeed advocated something like this, but this is to over-react: what is required is a form of morality that adequately reflects, and somehow balances, both the point of view of agents AND that of patients - for, after all, in order to reflect P1 morality will have to take into account people AS PATIENTS too. This is indeed a difficult task, but the point I want to emphasise is that we cannot even BEGIN to tackle it adequately until we have given some serious attention to the question of which AME is appropriate, and that the importance of this question has not been grasped by most contemporary moral philosophers.

For one thing, the degree of tension between the agent's and the patient's points of view would vary considerably depending on the question of which AME we felt to be appropriate. If, for example, we took the best development of our human potentialities to be one that resulted in a relatively 'selfless' orientation, one in which the agent was highly responsive to the needs of others, then the tension between the two points of view would be reduced: taking up the patient's point of view, in a world such as ours, in which so many

people's needs go unmet, would provide moral pressure for highly demanding moral principles, and the appeal to such an AME would actually REINFORCE (at least to some degree) that pressure - for the claim would be that, in responding to the needs of others, the agent would be realizing the best potentialities of her own nature. I will examine how such a claim might be supported in the next section. But if, on the other hand, we felt that the best development of our human potentialities was one, say, in which the agent concentrated on developing her particular talents, then the tension between the two points of view would be much more acute. What is clear, however, is that we will not know how strong that tension will be until we have addressed the question of which AME is appropriate.

iii/ A Stringent Model?

My main aim in this chapter is to try to bring into focus the kind of issue - that is, the question of which AME is most appropriate - on which I believe an eventual answer to the problem of demandingness will ultimately depend; an issue which, though it seems to me to be absolutely central to the debate, is sidelined in many discussions of the problem. I do not, of course, claim to be able to settle this issue here, but I would like to say something about the possibility that I mentioned in the last section: that is, the idea of a what I shall call a 'stringent model', a model which requires considerably more of us than CSM does. So far, the presumption has generally been that reference to an AME would bring in considerations that countervail against those considerations that support a stringent morality, and insofar as we conceive that stringent morality as a consequentialist one, that presumption seems to me, as I have said in section i/ of this chapter, to be correct. But if we take a less extreme form of stringency, then it does seem to me important to consider the possibility of models that actually REINFORCE the idea of a stringent morality; that is, ones that makes the case, FROM THE AGENT'S POINT OF VIEW, for a form of morality that is much more demanding than CSM.

In fact, this idea has already been foreshadowed in the example I gave in section i/ from Marx: there the idea was that we should not give normative status to the agent's desires for 'separation', for her own good defined in opposition to the good of others, because the model of the 'connected man' constitutes a better fulfillment of the agent's (perhaps 'real', or 'deep') nature. What if the best

fulfillment of our nature was indeed a relatively 'selfless' one, so that our most deeply considered AME, instead of providing considerations that militated against a highly demanding form of morality, actually reinforced it?

The examination of this possibility is of the first importance not only for the reason that I gave in the last section - that is, that if a stringent model can be motivated, it will reduce the tension between the agent's and the patient's points of view - but also for a number of other reasons. Firstly, because, if the way that I have characterized the debate is correct, the question of whether a highly demanding form of morality is defensible will largely depend on it. Secondly, because, despite this, this kind of issue, in the writings of most philosophers who have written about the problem of demandingness of morality, tends to be off the page somewhere. It is as if most philosophers have already taken a stand on the issue of which AME is appropriate before setting pen to paper, usually apparently derived from the conventional wisdom of the culture in which they live, and then simply use these assumptions without argument and without explicitly bringing them into the open in their writings. Thirdly, even those writers who do advocate a stringent form of morality often seem unaware both of the resources available for the defence of a stringent AME, and of how the support of such a model may help their cause, so that they need not be reduced to saying, in effect, 'I know it's tough, but you just have to accept it'. And finally, given that the general tenor of our age is, I think it is fair to say, one that is fairly sceptical about moral ideals, I believe that we should take particular care, in order to guard against a too ready

assimilation of the prejudices of our own time, to examine the case for a stringent model.

So how might one attempt to justify an AME that is responsive to any relevant features of human nature, and yet involves moral requirements that are considerably more demanding than those of CSM? The best general strategy, I suggest, would be the following: to attempt to show that there are features of human nature such that the best fulfillment of that nature would involve a life respecting those requirements. It need not, indeed, be denied that there are other features of human nature pulling against such a form of life, but it will nevertheless be maintained that the agent has good reasons to repress such features, insofar as she can, and to nourish those tendencies endorsed by a stringent model. What sort of reasons might those be? It seems to me that we can make at least a rough distinction between two sorts of reasons, and distinguish correspondingly two particular strategies.

In the first place, we might claim that such a life, a life corresponding to a highly stringent model, is, contrary to what we might first suppose, better FOR THE AGENT. That is, we might claim that it is in the agent's enlightened self-interest to live such a life (call this 'the self-interest strategy'). And secondly, we might claim that such a life, to the degree to which it is not in harmony with even the agent's enlightened self-interest, is nevertheless choiceworthy from the agent's point of view. That is, she may have reasons to choose a certain form of life because it is the best kind of life, in the sense of being the best (perhaps 'highest' or 'worthiest') fulfillment of her nature, but not necessarily in the

sense of being the best life FOR HER, from the point of view of her self-interest (call this 'the commitment strategy'). The availability of this second strategy follows from the fact that, as I said above, the point about taking account of the agent's point of view is that we should make an effort of identification with the agent, as well as with the patient, before deciding which moral principles are acceptable. Making such an effort of identification is supposed to ensure that whatever moral principles are chosen MAKE SENSE from the point of view of the agent, that they are compatible with the agent's nature - not necessarily that acting on them proves optimal from the standpoint of the agent's self-interest.

The first strategy might begin from the types of considerations that are sometimes put forward to show that it is in the individual's enlightened self-interest to live a moral life at all; the task would then be to extend those considerations in order to try to show that it is in fact in the agent's enlightened self-interest, not so much to live a standardly moral life, as to live a life that involves acting on requirements that are considerably more demanding than those of CSM. There are a number of very familiar arguments of the former type; perhaps the most useful for our purposes would be those based around the claim that the practice of certain moral virtues either constitutes, or is a means to, certain ('higher') goods; goods that the agent might not happen to want, but should want, because it is in her enlightened self-interest to have these goods. (The former version is, of course, associated with Aristotle ("happiness IS an activity of the soul in accordance with perfect virtue" (Aristotle (1), I, viii

← whose (or rather what)
is this?

(my 'italics')); the latter is advocated by, for example, by John Stuart Mill (J.S. Mill (1), part I.) For this claim to support the self-interest strategy effectively, it would have to be the case that the value for the agent of realizing certain goods that are internal to a stringent model outweighs the value to her of the other goods that leading this sort of life will lead her to forego.

Just how far such an argument might take us in a stringent direction is a matter for debate; but one that has not, unfortunately, received the attention it deserves in recent moral philosophy. In its stronger forms, it might be taken to support a kind of ideal of 'selflessness' - where that word is understood by reference to a 'lower' self that we would, it is claimed, be better to be without. Understood in this sense, this ideal need not have any puritanical or ascetic implications - though it often has been so taken, even by some of its advocates. Sensual or other ordinary pleasures are not, properly speaking, the target of its attack, but rather the aspect of the self that seeks to grasp such pleasures, or is entirely enslaved by them, to the detriment of the individual's overall wellbeing. Sometimes a kind of asceticism is advocated, but (again, properly speaking) this should be seen as a strategic device rather than as a deprecation of ordinary pleasures: if the individual is so focused on these pleasures as to be enslaved by them, it may be best for him to deny himself them, in order to change his character in a way that will, overall, be to his benefit - but there need be no implication that they are bad in themselves.

Versions of this thought are prominent in many religious and philosophical traditions, but it has not received much attention among

recent philosophers. (It is worth pointing out that one version of this thought is found in Aristotle - see, in particular, Aristotle (1) IX, viii - if only because Aristotle is so often taken as a kind of patron saint by opponents of stringency.) One exception is Iris Murdoch, who takes a view of both morality, and, relatedly, of the higher potentialities of human nature, that might be taken to provide support for the self-interest strategy. She writes, "One might start from the idea that morality, goodness, is a form of realism... The chief enemy of excellence in morality...is personal fantasy: the tissue of self-aggrandizing and consoling wishes and dreams which prevents one from seeing what is there outside one" (I.Murdoch (1), p.59). The reason why this sort of 'realism' is relevant to morality is that "true vision occasions right conduct... The more the separateness and differentness of other people is realized, and the fact that another man has needs and wishes as demanding as one's own, the harder it becomes to treat a person as a thing" (p.66). The link with the agent's good is established by the claim that "it (is) better to know what is real than to be in a state of fantasy or illusion" (p.64); the 'narrow' self cuts one off from the reality around one, "the great surprising variety of the world" (p.66). So it is better FOR THE AGENT not to be enclosed within a narrow world of personal anxiety and fantasy, but with the opening of her eyes comes a greater readiness to act on the needs of others. Thus it is not hard to imagine such a picture supporting a stringent form of morality; but, in this picture, such attention to the needs of others is not justified by reference to any abstract considerations about the nature of morality, but rather directly from the agent's point of view, by

reference to the agent's own good.

My main aim here is not to defend such a picture, but just to indicate that the kinds of issues it raises, though they are seldom discussed by contemporary moral philosophers, are vitally important to the issue of the demandingness of morality. Partly, no doubt, this is just a reflection of the sceptical, perhaps cynical, outlook of the age; an outlook that seems to me to be reflected uncritically by many contemporary philosophers. Susan Wolf, for example, writes of the "Loving Saint", whose "happiness...would truly lie in the happiness of others" (S.Wolf (1), p.420), that "when one reflects on (his) easily and gladly giving up his fishing trip or his stereo or his hot fudge sundae at the drop of the moral hat...one thinks that, if he can give these up so easily, he does not know what it is to truly love them. There seems, in other words, to be a kind of joy which the Loving Saint, either by nature or by practice, is incapable of experiencing" (S.Wolf (1), p.424). I find this an odd piece of reasoning. Why not suppose that the Loving Saint is quite capable of enjoying a fudge sundae, but that, on occasion, he regards other things as more important - and furthermore, that he is not thereby the loser, all things considered by having such an orientation? In most contemporary discussions it is just assumed that the giving up of many personal interests (in the narrow sense) will constitute a 'cost' for the agent; the possibility of there being BENEFITS for the agent in such a reorientation is hardly ever raised.

Getting repetitive

See back

In addition, there is the genuine difficulty of the degree to which a stringent model might actually make sense from the standpoint of the agent, beginning from where she happens to be. As Murdoch herself

says, "a moral position higher than our own may only be imagined as a deprivation" (I.Murdoch (2), p.307): it might appear to us now that we would lose out if were to become more open to others and their needs, but that appearance may be no more than a symptom of our present, depraved condition.

However far the self-interest strategy might take us in the direction of stringency, the commitment strategy might be expected to take us even further; and, on reflection, it might seem the better strategy to emphasise. This is because prudence (understood as the point of view of enlightened self-interest) has only one ultimate criterion: the good of the agent; it gives no INDEPENDENT weight, at the fundamental level, to the interests of others. Although we have seen how a perspective of sophisticated prudence might lead the agent to give considerable weight to the interests of others, the ultimate justification of her doing so has always come back to HER good. But the perspective of morality, by contrast, does give independent weight at the fundamental level to the interests of others. Although I have been looking at the question of the self-interested advantages of leading a moral life, there has been no suggestion of a REQUIREMENT that such a life be shown to be optimal from the standpoint of the agent's self-interest. And indeed nor could there be, if the standpoints of morality and of prudence are to retain their proper independence, the former giving independent weight, at the fundamental level, to the interests of others, and the latter failing to do so. We have been interested in the question of the degree to which, contrary to what might at first appear, the two perspectives might not be

totally antithetical, and, even if the first strategy is finally acknowledged to have its limits, our efforts may at least have shown that there are SOME important goods available to the agent by following a highly stringent model. If we ARE convinced of this, we might then apply our findings to the second strategy and say that, even if a certain stringent model isn't optimal from the point of view of the agent's self-interest, it might nevertheless still be choiceworthy for her, and not unreasonable to expect, because the goods that it offers for her at least satisfy what we might consider a reasonable minimum of self-interest, a minimum that we might want to make a condition on such a model being 'choiceworthy'. We might then claim that such a life, though it is not the BEST kind of life, from the point of view of the agent's self-interest, is nevertheless GOOD ENOUGH from that point of view to make it more choiceworthy than the best life for her, given the advantages FOR OTHERS that it involves.

Nor need there be any reason to think that there is anything paradoxical about the claim that there might be reasons for an agent to choose a form of life that is not the best form of life FOR HER. All that is necessary is that the agent be committed to values (hence the term 'commitment strategy') that may require her to sacrifice a certain amount of her self-interest, and this possibility should not be controversial. It IS controversial what kinds of considerations may constitute reasons for individuals to act; for example, whether mere desires do; but it is relatively uncontroversial that sincerely held axiological commitments do so.

The best way to develop the commitment strategy seems to me to be by utilizing the same kinds of considerations that were put forward above

?
Good reasons come
from correct
axiological
commitment
whether
sincere
or not.

under the self-interest strategy, but arguing that the agent's reasons to try to realize such higher goods, or to lead such a life, are not restricted to reasons of enlightened self-interest. The claim will then be, not that it will be better for the agent to realize these goods, though it will certainly be good for her to do so, but rather something like this: that it would be nobler for her to do so, or more worthy of admiration and respect, or more appropriate to the higher potentialities of human nature. In fact, I think it may not always be easy to separate the self-interest strategy from the commitment strategy (hence my saying above that the distinction is a 'rough' one), because, at a certain point, it seems to be difficult to say whether the agent is gaining or losing.

In any case, I believe that the power of such ideals is often underrated by moral philosophers; no doubt it is tempting for the friends of morality to emphasise its self-interested advantages, but perhaps they have underestimated the side of our nature that can be strongly drawn to such ideals without needing the claim that it would be in our enlightened self-interest to try to realize them. Mill, for example, might have supported his claim that "a being of higher faculties...can never really wish to sink into what he feels to be a lower grade of existence" (J.S.Mill (1), p.9) not by the claim that such a being will thereby have access to higher pleasures, but rather by the tendency in human nature to be drawn to such ideals. Similarly, Sidgwick might have emphasised, not the "delight" that accompanies the "sacrifice of sensual inclination to duty" in the "truly virtuous man" (H.Sidgwick (1), p.150), but the compatibility of such a sacrifice with an ideal that our nature and upbringing make us potentially

responsive to.

Once again, many philosophers seem to me to reflect the prevailing culture uncritically in being very suspicious of such ideals; but once again, there is usually little discussion or argument on this point. See S.Wolf (1), p.424 again, for example, this time on the "Rational Saint; or Michael Slote, who writes, "Selflessness...even if it was highly admired in Victorian times, nowadays - perhaps because we are less priggish and high-minded, perhaps because we think we know more about human psychology and are less willing to take things at face value - selflessness is a pretty automatic object of suspicion" (M.Slote (1), p.106). I am not sure if we should share uncritically Slote's confidence about the superior wisdom of our own age. Perhaps we shouldn't let our suspicion be quite so "automatic": perhaps its being so is more a reflection of OUR unexamined prejudices, absorbed uncritically from the social milieu in which we have grown up, than of any good argument or clear moral vision. In the same passage, Slote goes on, "consider, for example, how we tend to regard a character in literature like Pere Goriot, who pauperizes and then destroys himself in order to allow his daughters to go on leading their frivolous, luxurious existences" (ibid). Now, a moment's thought is enough to show this example to be doubly irrelevant: in the first place, no one is suggesting that we should 'destroy' ourselves, and, in the second place, no one is suggesting that we should make sacrifices so that others may lead "luxurious, frivolous existences". The question immediately arises, then, of why Slote can even have thought that this example is relevant; and the most likely answer will be that he is so spellbound by the prejudices of HIS time that he does not even find it

necessary to think hard about what the ideal of selflessness involves before dismissing it.

None of this implies, of course, that it is NOT possible for individuals to sacrifice themselves in unfulfilling or pointless ways, nor that we need deny that Marx and Freud, among others, have done much to reveal how this kind of compulsive or alienated form of morality can come into being. But until we establish that the attempt to follow a stringent model NEED take such a form, perhaps we should regard our 'automatic suspicion' itself with some suspicion.

Although I have only had space to say very little about the possible justification for a stringent AME, I hope I have said enough to let the following points emerge. Firstly, as I said above, the advocate of a stringent form of morality may have many more and stronger resources available for the defence of her view, even when the importance of the agent's point of view is taken into account, than is often supposed - more resources, indeed, than even the defenders of stringency usually seem to be aware of. Frequently the only resource that such defenders of stringency seem to be aware of is some tendentious application of the notion of 'rationality', as if one must be compelled into the acceptance of stringency on pain of irrationality - not a strategy, apart from anything else, that even seems to have much motivating force.

Secondly, it should have become clear that the notion of an 'appeal to cost', so frequently put forward as a reason to limit the stringency of morality, needs to be handled much more carefully than it usually is. Both the defenders and the opponents of stringency tend

to take largely for granted what would constitute a 'cost' for the agent, and then the only question is how much she can be expected to pay. But the discussion above should make it clear that we cannot decide what in fact constitutes a 'cost' for the agent, as opposed to what may appear to her (or indeed to us, on first examination) to constitute one, without reference to a conception of her enlightened self-interest, and that this question, in turn, cannot be separated from the issue of what AME is appropriate. To advert to the quotation from Marx again, it is too easily assumed that the frustration of desires that reinforce the 'separation of man from man' would constitute a 'cost' for the agent, when it might be argued that the frustration of such desires, as part of a coherent project of re-orientation, might serve the agent's enlightened self-interest. Of course, I do not claim to have established that this is indeed so, but only to have indicated some of the types of considerations that might support this line of argument. The important point is that the opponent of stringency will have to engage with these arguments, rather than merely assuming a certain picture of the agent's self-interest. I have focused on how an advocate of a stringent form of morality might support his position by developing a stringent AME, but it is the task of her opponent not only to criticise the stringent model, but also to argue for a less stringent model. Perhaps she can convince us that a kind of robust individualism is a more appropriate expression of our nature. Perhaps she can provide considerations that would make the kind of selflessness advocated by writers like Murdoch, and by many religious traditions, seem inappropriate to our nature. The main point I want to make here is that the resolution of the

question of demandingness seems to me largely to depend on such issues, and yet such issues are usually off the page somewhere.

iv/ Tying the Threads Together

In chapter II two main strategies for dealing with the problem of demandingness emerged. Both appealed to certain features of human nature, and of the human agent in particular: the consequentialist strategy appealed to these features to motivate the claim that it is COUNTER-PRODUCTIVE to ask too much of people; and the Nagel/Scheffler strategy appealed to these features to motivate the claim that it is UNREASONABLE to ask too much of them. (The contractarian strategy also appealed to the idea that it is unreasonable to ask too much of people, but as I argued that it would need supplementation by a theory that provided criteria for what is 'reasonable' or 'unreasonable', and as Nagel and Scheffler have attempted to supply such criteria, I have concentrated on them.) I argued, firstly, that the general kind of approach that Nagel and Scheffler employ is compatible with P1, and so may be seen as a legitimate moral argument; in particular, that the appeal to what it is reasonable to ask of people is one that can be impartially justified, insofar as compatibility with P1 guarantees impartiality. It is compatible with P1, because it rests on taking everyone into account in the same way, but specifically in their role AS AGENTS: if there are features of human nature that are plausible candidates for having fundamental moral relevance, and render us, as agents, resistant to extreme moral demands, then it will not infringe P1 to take those features into account.

In fact, it seems to me that this approach is not only compatible

with P1, but represents the best interpretation of P1: given that we do all play these distinctive roles of patient and agent, and that one central function of morality is to connect agents and patients, it would seem necessary for morality to take all of us into account in the same way in both of these distinctive roles. Consequentialism, by contrast, seems to require a kind of willed forgetting of the kind of role that moral principles are to play in the lives of agents, with their particular nature: when it is found that human beings seem peculiarly ill-adapted to do what consequentialism requires of them, features of that nature are then belatedly accommodated in the kind of strategic way that we have seen. But if those features are quite general, and plausibly candidates for having fundamental moral relevance, then it seems that we should get them into focus at the ground level of our moral theory, rather than introducing them through the back door at the end.

I then argued that, though it would not infringe P1 to give an independent normative role to features of human nature that render us, as agents, resistant to extreme moral demands, Nagel and Scheffler fail to provide adequate criteria for determining what those features are. The indiscriminating appeal to 'personal' motives or concerns, understood by reference to 'impersonal' ones, seems to let in too many features of human nature that we would not want to endorse. And so I suggested that we determine which features of human nature to give such an independent normative status to by reference to an agent-centred model of excellence, a model that would give us the criteria for what should be taken as 'reasonable' that Nagel and Scheffler (and the contractarian) so desperately need. What it is

'unreasonable' to expect of people is, in a sentence, what is quite incompatible with what we feel to be the best achievable fulfillment of their nature as human beings - and as particular, individual human beings, it should be borne in mind, in a particular social setting and culture.

So I have three broad suggestions to contribute to the debate about the demandingness of morality. The first is the idea of using P1 as a framework for discussion of the problem; the second is the idea that, in interpreting P1, we should take everyone into account specifically in their role as agents (as well as as patients); and the third is that we will need to address the question of which AME is appropriate in order to make progress with the problem. Returning now to some of the concerns of chapter I, I would like to say something more about the claim that P1 might not be so vulnerable to the objections of the critics of moral theory as may at first appear; in particular, because the way that I have interpreted it, as requiring us to consider the question of which AME is appropriate, introduces many of the kinds of issues that those critics have wanted to emphasise, and have accused many moral theories of neglecting.

It will be remembered that I argued in section ii/ of chapter I that, given the difficulty of pursuing the question of whether the place currently given to concern for the needy by CSM is one that we can justify, it may be worthwhile, on heuristic grounds at least, to look at what some moral theorists have to say about the problem. With this in mind, I introduced, in section iii/, the Principle of Moral Equality, a principle that I claimed was at least neutral between many influential moral theories. I also claimed that this principle was

not, in itself, at least obviously vulnerable to some of the central objections of the critics of moral theory. Nevertheless, those sympathetic to those critics will no doubt have had their patience tested by some of what has followed. Much of the debate has been framed in terms of an acceptance of consequentialist reasons, which set an extremely demanding standard, and the strategy has been to see whether we can find any reasons to justify modifying that standard. The critics of moral theory will have at least two criticisms of this procedure: firstly, they may either reject consequentialist reasons altogether, or, if they accept them at all, give them a far less central place in morality than even Nagel and Scheffler do; and secondly, even if they accept consequentialist reasons, they will feel that it is loading the dice to put those types of reasons into place first, and then go looking for something to set against them.

But, even though the procedure I have followed may be suspect to the critics of moral theory, it has had results to which they may not be so unsympathetic. It has turned out that, even if we give consequentialist reasons the kind of central place that Nagel and Scheffler do, there will still be strong reasons, based on the need to reflect P1 adequately by taking into account the distinctive role people play as agents, to bring in some of the types of considerations that those critics themselves want to emphasise. For example, focusing on the question of which AME is appropriate will inevitably bring into the forefront issues concerning integrity, and concerning the fact that each moral agent has a particular character; issues that Bernard Williams, in particular, has accused some moral theories of neglecting (see B.Williams (2) and (3)). I myself want to remain neutral on the

question of the validity of consequentialist reasons: my strategy has been to argue that, EVEN IF we begin with such reasons, we shall need, in order to reflect P1 adequately, to take the agent's point of view seriously, and to give properties that are relevant to that point of view an independent normative status in our moral theory.

The importance of the issue of which AME is appropriate can then be separated from the route we took to establish it: it will then be the case that, even if we do NOT accept consequentialist reasons, we will still need to consider any prospective moral principles in the light of the nature of the moral agent, and what we consider to be the best expression or development of that nature. This finding can then be applied to the quandary we found ourselves in in section ii/ of chapter I, when we acknowledged that there were strong prima facie grounds for stringency derived from CSM, but found that it was difficult to know how to establish, without going beyond the resources of CSM, whether there were legitimate countervailing considerations strong enough to give us reason to reject stringency. I believe that the notion of an AME may at least help us to set about tackling this problem: the question of whether the comfortable, specifically, are justified in their present patterns of concern will be determined, in part, by reference to the question of whether we feel that those patterns of concern are appropriate to the best achievable fulfillment of the nature of those individuals. Keeping this question at the forefront of our minds should also help to prevent the kind of mental logjam that occurs if we attempt to extract a principle from one particular case (for example, the principle that 'if you can provide a great benefit to someone else at little cost to yourself, you should

do so'), and then find that repeated applications of such an apparently unobjectionable principle seem to lead to unacceptable consequences (see, for example, J.Fishkin (2)). Keeping in mind the issue of what kind of role moral principles are to play in the life of a moral agent, and the degree to which any suggested principles are compatible with what we feel to be an appropriate AME, should reduce the temptation to set the issue up in this kind of context-free way.

An eventual solution will depend, then, on the model of excellence that we feel to be most appropriate, and that, of course, is a profoundly difficult question. But the points that I want to emphasise here are, firstly, that it is at least the RIGHT question, the one we need to address in order to make progress with the problem; secondly, that it is a question that ALL of us need to address, whether we are advocates of specific moral theories, or opponents of them all; and, thirdly, that CSM itself has considerable resources for discussing this question. Much pre-philosophical debate about morality seems to me indeed to be about the sort of person one should be, the kind of life one should lead, what kinds of ideals are appropriate or viable, and so on. Even if moral philosophers have generally, at least until very recently, tended to neglect such questions, the wider culture has not, and can draw on many resources - for example, from literature, from various religious traditions, from psychology, and so on - in trying to answer them.

Similarly, the importance of an AME for morality can be separated from P1 itself, though it is by reference to P1 that I have argued that giving an independent normative role to agent-related properties is impartially justified. The critics of moral theory have tended to

be less concerned than many moral theorists about the question of the impartial justification of our moral principles, but, at least insofar as their doubts about P1 are prompted by the assumption that its adoption will lead inevitably to a moral theory that is subject to their criticisms, we can find reasons to allay their doubts. With regard to the criticisms of the anti-impartialists specifically, I have already pointed out, in section iii/ of chapter I, some considerations that may lead us to doubt this assumption: very briefly, that the MORAL AGENT won't necessarily have to take everyone into account, and that she won't necessarily have to disregard all personal considerations if she is to act in a way that is impartially justified. And thinking of what agents should be required to do in the light of a plausible AME would, I suggest, make these troubling and extreme ideas seem all the more inappropriate.

I also pointed out that the acceptance of P1 does not in itself commit one to a 'moral theory' in the sense criticised by the anti-theorists; that is, to a principle or set of principles that would enable us to determine how to think about any particular moral situation. And, once again, the suggestion that we will have to give an independent role to agent-related properties, and that our doing so will have to be justified by reference to an AME, makes this prospect all the more unlikely. But the point that I'd like to emphasise now, with regard particularly to the 'virtue ethicists', is that some of the ideas put forward in this thesis may give us reasons to question the kind of basic division in approach to morality that exists among contemporary moral philosophers. Sometimes it seems to be presupposed that there are, very broadly, two kinds of ways of thinking about

morality, ways that may seem to be irreconcilable (see R.Wollheim (1), pp.197-198 and S.Scheffler (2), pp.50-51). On the one hand, one can begin with the historically situated agent, with her particular psychology, embedded within a network of personal projects and relationships, in a way that is often seen to make considerations of justice or impartiality inherently problematical (see, for example, P.Foot (1), or J.Cottingham (1)). On the other hand, one can begin in a more abstract or theoretical way, perhaps with the idea that morality involves some kind of condition of impartiality, and then the presumption usually is that we are going to find it difficult to accomodate facts about individual psychology, or the fact that each person has their own life to lead (this, of course, is the way that I have approached matters for most of this thesis: Nagel, Scheffler, Railton, and Kagan would all provide examples). But if we can show that we have to get agent-centred considerations into focus right from the start, at the ground level, in an approach that begins from the Principle of Moral Equality; and if, on the other hand, we can show that considerations emphasising the reality of other persons and the importance of their needs may need to occupy a central place in a form of thinking about morality that starts resolutely from the point of view of the particular individual, then we might be able to show that the two broad approaches to morality mentioned above are not quite so antithetical as they might at first appear to be. In particular, it has tended to be the critics of moral theory, particularly as 'virtue ethicists', who have emphasised agent-centred considerations - but insofar as any doubts they may have about P1 may be prompted by the assumption that any form of morality which accepts P1 would be unable

to find an appropriate place for such considerations, we have now seen reasons to believe that their doubts can be allayed. Perhaps it is too early to allow ourselves the hope of a grand synthesis, but we might at least have indicated that there may be a way of framing the issues that is more enlightening than the spectacle of advocates of one or the other broad approach pressing their fundamental intuitions against advocates of the other approach.

BIBLIOGRAPHY

- ARISTOTLE (1): Nicomachean Ethics (Indianapolis, 1985)
- BARRY, B. (1): Justice as Impartiality (Oxford, 1995)
- COTTINGHAM, J. (1): 'Partiality and the Virtues' (in Crisp (2))
- CRISP, R. (1): 'The Dualism of Practical Reason' (Proceedings of the Aristotelian Society, 1996)
- CRISP, R. (2): (ed) How Should We Live? (Oxford, 1996)
- FISHKIN, J.S. (1): The Limits of Obligation (Yale University Press, 1982)
- FISHKIN, J.S. (2): Beyond Subjective Morality (Yale University Press, 1984)
- FOOT, P. (1): 'Moral Beliefs' (in Foot (2))
- FOOT, P. (2): Virtues and Vices (Berkeley, Calif, 1978)
- HOOKER, B. (1): 'Rule Consequentialism and Demandingness' (Mind, 1991)
- KAGAN, S. (1): 'Does Consequentialism Demand too much?' (Philosophy and Public Affairs, 1984)
- KAGAN, S. (2): The Limits of Morality (Oxford, 1989)
- MARX, K. (1): 'On the Jewish Question' (in Marx (2))
- MARX, K. (2): Early Writings (ed L.Colletti, Cambridge, 1970)
- MCDOWELL, J. (1): Mind and World (Cambridge, Mass, 1994)
- MILL, J.S. (1): Utilitarianism (Everyman, 1993)
- MURDOCH I. (1): 'On "God" and "Good"' (in Murdoch (2))
- MURDOCH, I. (2): The Sovereignty of Good (Routledge and Kegan Paul, 1970)
- MURDOCH, I. (3): Metaphysics as a Guide to Morals (Penguin, 1992)
- MURPHY, L. (1): 'The Limits of Beneficence' (Philosophy and Public

Affairs, 1992)

- NAGEL, T. (1): The View from Nowhere (New York, 1986)
- NAGEL, T. (2): Equality and Partiality (New York, 1991)
- RAILTON, P. (1): 'Alienation, Consequentialism, and the Demands of
Morality' (in Scheffler (3))
- SCANLON, T.M. (1): 'Contractualism and Utilitarianism' (in Sen and
Williams (1))
- SCHEFFLER, S. (1): The Rejection of Consequentialism (Oxford, 1982)
- SCHEFFLER, S. (2): Human Morality (New York, 1992)
- SCHEFFLER, S. (3): (ed.) Consequentialism and its Critics (New York,
1988)
- SEN, A. and WILLIAMS, B. (1): Utilitarianism and Beyond (Cambridge,
1982)
- SIDGWICK, H. (1): The Methods of Ethics (London, 1907)
- SINGER, P. (1): 'Famine, Affluence, and Morality' (Philosophy and
Public Affairs, 1972)
- SINGER, P. (2): Practical Ethics (Cambridge, 1986)
- SLOTE, M. (1): From Morality to Virtue (New York, 1992)
- SMART, J.J.C. and WILLIAMS, B. (1): Utilitarianism: For and Against
(Cambridge 1973)
- TAYLOR, C. (1): 'The Diversity of Goods' (in Sen and Williams (1))
- WIGGINS, D. (1): 'Claims of Need' (in Wiggins (3))
- WIGGINS, D. (2): 'Universalizability, Impartiality, Truth' (in Wiggins
(3))
- WIGGINS, D. (3): Needs, Values, Truth (Oxford, 1987)
- WILLIAMS, B. (1): 'The Idea of Equality' (in Williams (2))
- WILLIAMS, B. (2): Problems of the Self (Cambridge, 1973)

WILLIAMS, B. (3): 'A Critique of Utilitarianism' (in Smart and
Williams (1))

WILLIAMS, B. (4): 'Persons, Character, and Morality' (in Williams (5))

WILLIAMS, B. (5): Moral Luck (Cambridge 1981)

WILLIAMS, B. (6): Ethics and the Limits of Philosophy (Cambridge,
Mass. 1985)

WOLF, S. (1): 'Moral Saints' (Journal of Philosophy, 1982)

WOLLHEIM, R. (1): The Thread of Life (Cambridge, Mass. 1984)