'OUGHT' IMPLIES 'CAN'

by

JAMES GRIFFIN

The Lindley Lecture
The University of Kansas
October 5, 2010

The E. H. Lindley Memorial Lectureship Fund was established in 1941 in memory of Ernest H. Lindley, Chancellor of the University of Kansas from 1920 to 1939. In February 1941 Mr. Roy Roberts, the chairman of the committee in charge, suggested in the *Graduate Magazine* that

> the Chancellor should invite to the University for a lecture or a series of lectures, some outstanding national or world figure to speak on "Values of Living" – just as the late Chancellor proposed to do in his courses "The Human Situation" and "Plan for Living."

In the following June Mr. Roberts circulated a letter on behalf of the committee, proposing in somewhat broader terms that

> The income from this fund should be spent in a quest of social betterment by bringing to the University each year outstanding world leaders for a lecture or series of lectures, yet with a design so broad in its outline that in the years to come, if it is deemed wise, this living memorial could take some more desirable form.

The fund was allowed to accumulate until 1954, when Professor Richard McKeon lectured on "Human Rights and International Relations." The next lecture was given in 1959 by Professor Everett C. Hughes, and has been published by the University of Kansas School of Law as part of his book *Student's Culture and Perspectives: Lectures on Medical and General Education.* The selection of lectures for the Lindley series has since been delegated to the Department of Philosophy.

'OUGHT' IMPLIES 'CAN'

by

JAMES GRIFFIN

White's Professor of
Moral Philosophy, Emeritus
Oxford University

'Ought' Implies 'Can'

James Griffin
White's Professor of Moral Philosophy, Emeritus
Oxford University

1. The larger question

There is a neglected part of ethics that needs exploring. The part I have in mind is how factual limits to human capacities affect our moral obligations. There are limits to human motivation and action, and many of us think that 'ought' implies 'can'. There are also limits to human understanding; it is no good a philosopher's settling on a decision procedure for, or criterion of, moral right and wrong if it requires intellectual feats that are beyond us, at least beyond yielding answers that we would rely upon. What is more, a society would want its moral demands to be understood, widely accepted, and effective in practice. These constraints of understanding, acceptance, and effectiveness are likely also to constrain the contents of a society's moral demands.

This larger question would require a whole book. In this lecture, I want to concentrate on whether 'ought' implies 'can', a question that, while not as neglected as the others, is not thought about enough.

2. The matter of approach

I have long thought, without giving much attention to the matter, that there must be some interpretation on which the claim that 'ought' implies 'can' is true, difficult though it is to state it. What we have to do, I have thought, is to find that interpretation. If so, we cannot directly proceed, as some writers do, to the refutation of the claim by counter-example. How can one know that an example counters a claim when one does not yet know what the claim is? Nor can we immediately embark on abstract arguments against the claim, though some writers do. How can one aim one's attack when one has not yet located the target? Nor, even, should we start by stipulating a sense for the claim that 'ought' implies 'can', though it may be of interest whether a particular stipulated claim is indeed true. What we want to know, however, is not whether the claim is true on some interpretation, but whether it is true on the best interpretation.

Consider, for instance, this proffered counter-example.[1] I promised Black that I would pay him $50 by noon today, and I also promised White that I would pay her $50 by noon today. But I find now that I have enough money to pay only one of them. If 'ought' implies 'can', then, since I cannot pay both, it is not the case that I ought to. But surely, since I promised, I ought.

But this attempt at refutation is too quick. My abilities in this case have not yet been fully enough described. True, I cannot now pay both Black and White by noon today, but I could have avoided getting into this fix. Is that enough to justify saying that, since I promised, I ought?

Or take this example. Brown, a recovering drug addict, tries, but occasionally still fails, to stay off drugs. Brown's resistance has not yet become strong enough to master every temptation that comes his way. Still, in a larger sense, Brown can master his addiction; it is clearly within human ability. Perhaps that is enough to justify saying that Brown ought not to have given in. It depends upon how we understand the 'can' in the claim that 'ought' implies 'can'. Is an act's being within normal human capacity enough? We need to feel our way to the best interpretation of the claim.

The claim has been made for millennia.[2] In our day it is primarily associated with Kant.[3] But this is because he cites it often, without either explaining or justifying it. For example, in the First Critique he remarks parenthetically: 'The action to which the "ought" applies must be possible under normal conditions.' I understand him to say here: 'ought' implies 'can within the limits of the laws of nature'. If I spot a child teetering perilously on the parapet of a skyscraper, I might wish I could fly, Superman-like, to snatch the child from death. But it seems to

me a misuse of the word to say in this case that I 'ought' to do so. But that does not yet establish the content of the claim that 'ought' implies 'can'. It may well be that there are further understandings of 'cannot' besides Kant's particularly strong one here on which it is also true that if 'cannot' then 'not the case that ought'. We still have to feel our way to the most satisfactory interpretation of the claim. As a first move, let us look at the three words that make it up.

3. 'Ought'

We can start by reminding ourselves of the wide variety of uses of the word 'ought'[4]:

1. You ought not to steal.
2. You ought to quit smoking.
3. When you add 17 to 34, you ought to get 51.
4. The train ought to be in now.
5. It ought to be sunny tomorrow.
6. You ought to have known; ignorance of the law is no excuse.
7. If you wanted to break in, you ought to have brought a jemmy.
8. The cricket pitch ought not to have behaved like that.
9. I ought to be at the meeting tonight, but I can't manage it.
10. He ought, with his training, to be able to dismantle this simple engine.

We do not use the word 'ought' merely to indicate the presence of some kind of reason – a moral reason, a prudential reason, an epistemic reason, and so on. It may be the case that I have a reason to visit old Jones (he is bed-ridden and lonely), but that my connection to him is so slight and remote that it would be incorrect, because too strong, to say that I *ought* to visit him. The judgement that it is not the case that I *ought* to visit him may on some occasions be intended as a moral judgement. But it may also, instead, be meant as a lexicographical judgement: namely, given the circumstances of the utterance (i.e. the nature of my relation to old Jones) it would be a misuse of the strong term 'ought'. Unsurprisingly, the grounds for the moral judgement and those for the lexicographical judgement will overlap. The word 'ought' is here being used as an action-directing term of the moral sort, so morally relevant conditions will inevitably be part of the criteria for correct and incorrect use of the term.

A more successful semantic analysis is this: the word 'ought' is used when there is a standard or regularity in the background and what is said ought to be is what would accord with the standard or regularity.

3

Behind the second example on my list, the standard is prudential: smoking damages one's health. Behind the third there is a mathematical standard of correctness: 17 plus 34 equals 51. Behind the fourth there could be various sorts of standards. You could appropriately say that the train ought to be in now, because that is what the timetable says. I could then inform you that British trains are nearly always ten minutes late, and on the basis of that regularity you could say, ten minutes later, 'well, it ought to be in now'. If the station loud-speaker announces that the train has been delayed a further thirty minutes because of work on the tracks, then after half an hour you could correctly say again 'it ought to be in now'. It is not that the announcement over the loud-speaker indicates a reason for you to believe or to expect that the train will arrive in thirty minutes; you might have decided that station announcements in Britain are pure fiction. None the less, the announcement sets a standard against which you can correctly speak of what 'ought' to happen. Behind the fifth example there is a causal regularity: an area of high pressure generally brings sun. In most legal systems there is a rule requiring those subject to laws to acquaint themselves with them; that is why, as in the sixth example, you ought to have known that it was illegal. There is a rule of prudence: if you want a certain end you must want the necessary means to it; so, as the seventh example says, you ought to have brought a jemmy. I shall stop there: the same general form of explanation fits the remaining examples too.[5]

4. 'Implies'

The word 'implies' covers several different relations. In his discussion of '"ought" implies "can"' Walter Sinnott-Armstrong canvasses four possible relations that implication might be thought to cover here[6]:

1. Semantic entailment: semantic because it holds by virtue of the meanings of the terms involved or the truth conditions of judgements containing the terms; and entailment because, if an agent cannot do such-and-such, then, by *modus tollens*, it is *false* that the agent ought to do it.
2. Semantic presupposition: semantic as above; a presupposition because, if an agent cannot do such-and-such, it is *neither true nor false* that the agent ought to do it (the inability asserted in the consequent rules out the antecedent's having a truth-value).
3. Moral implication: if an agent cannot do such-and-such, it would be morally wrong (e.g. unfair), and *false*, to assert that the agent ought to do it.
4. Conversational implication: not semantic implication but

4

pragmatic, having to do with the intended or actual effect of an utterance; if an agent cannot do such-and-such, it is *pointless* to assert that the agent ought to do it.

The first three relations are, Sinnott-Armstrong says, universal: they claim that *all* 'oughts' applied to persons imply 'cans'. The first two are universal because semantic; they depend upon the meanings of 'ought' and 'can', whenever they are asserted of persons. The third depends upon an assumed universal moral principle: if 'cannot', it would be wrong (e.g. unfair) to assert 'ought'. The assertion of a universal relation can be refuted, of course, by a particular negative, and, as I shall come to, there do seem to be counter-examples to '"ought" implies "can"'. The only one of these four relations that Sinnott-Armstrong thinks holds is the last. But it is, as he points out, weak. That it may be *pointless* to say 'ought' to an agent who 'cannot' does not mean that it would be either *meaningless* or *false* to say it. Therefore, he rejects the principle.

But Sinnott-Armstrong's four interpretations are not exhaustive. We are concerned here with the possible presence of an implication. So we should start with 'implication' in its generic sense expressed by the ordinary-language connective 'if p, then q': the truth of p is a sufficient condition of the truth of q.' The implications that Sinnott-Armstrong canvasses are four sub-classes of implication in this generic sense, and there are others. The most plausible claims that 'ought' implies 'can', I should say, are not any of his. The idea that the word 'implies', as it appears in the claims, implies some form of universality should be resisted. It is much more likely that only some sorts of 'cannot' defeat 'ought'. If '"ought" implies "can"' is to be accepted as a principle in ethics, the 'ought' must obviously be taken to be restricted to agents, and the 'can' as referring to the ability of agents. Does the principle cover only moral 'oughts'? That would be too narrow. The principle also covers some prudential 'oughts' ('you ought to quit smoking') and some epistemic 'oughts' ('you ought to have got 51 as the sum') and perhaps other sorts as well.

5. 'Can'

The putative principle must also be taken to cover a variety of 'cans'. It seems clearly to cover inabilities present in all human beings in all natural circumstances – e.g. flying Superman-like to snatch the child from death. Such abilities are the stuff of fantasy. It may also cover behaviour that is beyond the capacity of all but the rarest human beings – all but, say, saints and heroes and exceptional achievers. At least, that is a venerable doctrine of some religions. It may also cover ordinary people facing extraordinary sacrifice. Again, it is an enduring and widely held belief that one does

5

not have to save a drowning child at considerable risk of one's own life. The belief in this case may be that saving the child in those circumstances is more than morality demands even if one could manage it, or that it is more than morality demands simply because ordinary human beings cannot manage it.

The principle may also cover sacrifices less great than life. It may even cover much of the everyday life of ordinary people. That is why the principle may be far more important than it is usually thought to be. Human beings quite naturally form relations of love and affection. They also quite naturally form conceptions of a worthwhile life. These relationships to particular persons and commitments to certain causes or projects are central to their living a good life. But these relationships and commitments may render them incapable of meeting certain demands – for example, duties demanding a high degree of impartiality. Again, it is a common and strong intuition that we ordinary human beings have a domain of permitted partiality, and that this domain includes much of our ordinary life.

Then, there are cases that the principle seems not to cover (e.g. the possible counter-example I alluded to earlier: my combined promises to Black and White).

This gives us the rough sketch of a spectrum that extends from inabilities that the principle seems comfortably to cover, to inabilities that it may well cover, to inabilities that it seems clearly not to cover. We must go deeper.

6. An ability behind the inability

Look again at cases in which we cannot do such-and-such though the 'cannot' seems not to defeat the correlative 'ought'. I cannot pay both Black and White, but I could have avoided getting myself in that impossible situation. Avoiding such situations is well within human capacity. I am sympathetic with the recovering drug addict because of the strength of his addiction, but I still know, and he may know too, that many ordinary persons, neither saints nor heroes, have managed to stop. Again, there is a human ability of ordinary level behind this inability. And a judge in a finely balanced case might reach a faultless judgement of guilt (no one could have done better), but then later, when by chance new evidence comes to light, discover that he had been mistaken. He might rightly say: 'I ought not to have convicted', despite the fact that he justifiably believed at the time that he could not have done otherwise. But it is, in general, within human capacity to collect evidence, to know the law, and to reach correct decisions. There are general human abilities behind the judge's inability to

6

succeed in this particular case. The background ability in all three cases, I believe, explains why the 'cannot' they use does not defeat the 'ought'.

But in the intermediate cases – cases in which our intuitions suggest that the 'cannot' does defeat the 'ought' – there is no background ability. Does this explain why here the 'cannot' does indeed defeat the 'ought'? Let us look more closely at these cases.

7. The limits of human motivation

Prudential deliberation, I have suggested elsewhere[8], ends up with a list of values, such as enjoyment, understanding, accomplishment, deep personal relations, autonomy, and liberty. A striking feature of many items on that list is their long-term, life-structuring character. To have deep attachments to particular persons is to acquire motives that shape much of one's life and carry on through most of it. To accomplish something with one's life requires dedication to particular activities that typically narrow and absorb one's attention. Many prudential values involve commitments – to particular persons, institutions, causes, and careers. One cannot live a prudentially good life, one cannot fully flourish, without becoming in large measure partial. That partiality then becomes part of one; it may be something that one cannot psychologically enter into and exit from at will. It involves becoming a certain kind of person. Even short-term pleasures have finally to be judged in a fairly long-term, character-fixing way, because a person has to decide how much place to give to living for day-to-day pleasures seen up against competing ways of life.

On this view of prudence, one should become deeply partial. That partiality is, I think, bound to be in some tension with the moral point of view. I doubt that we shall ever find a way to dispel totally the tension between prudence and morality, even if there is, as I think there is, considerable interpenetration between the two.

One might hope that the tension could be reduced if, say, we made impartial benevolence our central project in life; then one could accomplish something with one's life (a prudential value) by behaving impartially (a moral ideal). But this would be merely to realize one prudential value at the expense of many others – at the expense, say, of deep personal relations, of many forms of enjoyment, and, if this project takes much of one's time, of a lot of understanding. The tension arises even within the aim of impartially maximizing the good. A world in which everyone's life was as good as possible would be a world in which people were full of commitments. The impartial ideal, then, would be a world populated by agents who may be incapable of promoting the impartial ideal. And what one comes to see as one's own form of flourishing becomes a large part of

7

who one is. That is why, in his autobiography, John Updike says: 'We are social creatures but, unlike ants and bees, not just that; there is something intrinsically and individually vital which must be defended against the claims even of virtue'.[9] Somehow, these two parts of ethics, the demands of others and the goal of individual flourishing, must be rendered, if not entirely harmonious, at least combinable in one normative point of view, and in one human personality.

This brings us up against an apparently empirical question on which most of our ethical beliefs rest, but which we largely ignore: What are the limits of human motivation?

Evolution has planted in us both a crude form of self-interest and a form of limited altruism.[10] One obvious way to combat our biological inheritance is by increasing people's knowledge; we can make them appreciate far more fully and vividly the plight of others.[11] We know how one photograph of a starving child can make tens of thousands reach for their cheque-books. Still, I doubt that the problem could be just a deficit in knowledge. Well-intentioned famine-relief workers, whose field of vision is filled with starving victims, no doubt make great sacrifices to help them, but do not generally sacrifice themselves to the point where their marginal loss equals the others' marginal gain. It is true that there are usually good impartial-maximizing reasons for those aiding to have more than those aided, but relief workers generally do not sacrifice themselves to that point either. And it is hard to believe that it is because their knowledge is still somehow incomplete or faint.

A more hard-headed answer is to impose some stiff behaviour modification. We should not mistake limitations that arise merely from current social conditioning for genuine limitations in human nature. Anyway, human nature is not itself unmalleable. We are naturally partial; but so are we naturally aggressive and carnivorous, and no one suggests that those two features of human nature could not be proper subjects for drastic moral demands. In time of war, hundreds of thousands of perfectly ordinary people go off to defend their country at the risk of their lives. If military training can help motivate them to go into battle, could not a well-conducted moral training do something comparable for us? In most cases, though, soldiers can be brought to accept great potential danger, I suspect, out of fear of the sergeant-major or a court martial or of being shamed in front of their mates. This suggests that we could, similarly, institute a kind of neighbourhood Red Guard to train us as children and to keep us up to moral scratch thereafter. There would, however, be a terrible price to pay. We are willing to pay a comparable price in an emergency such as war, because of the exceptional importance of what is at stake. But perhaps we ought to think that what is at stake in moral life is equally important.

8

But there are two different sorts of doubts about that whole Red Guard enterprise. First, to produce moral action by fear denies an agent autonomy, and loss of autonomy is the loss of an essential component of morality, at least as most of us now conceive of morality. Can we, in the name of morality, so substantially undermine morality? Anyway, second, the Red Guard approach does not seem to work. Think what forces are gathered on the other side. Our propensity to form bonds of love and affection are vastly stronger than our propensity, say, to eat meat. We think that our personal relations and our commitments to certain causes are central to a good life. These beliefs are not only common, but also sound. And some sound ethical beliefs are, simply because they are sound, very likely to be persistent and recurrent, especially now that societies are much harder to isolate from one another. Are these attachments that we can enter into and exit from at will? We could, of course, try to suppress these commitments or alter these beliefs about the good life, but we are unlikely to succeed for long. The Red Guard enterprise aims at an unsustainable state. Many Chinese youths who were fanatical products of the Cultural Revolution in their childhood turned up later in the tents in Tiananmen Square.

Yet another answer is that, besides increasing knowledge and remoulding agents, we should give them a more inspiring goal. This is, I take it, Iris Murdoch's answer.[12] Modern moral philosophy, she thinks, is unambitious. It sets modest goals; it assumes that our psychological capacities are puny. But goals and capacities are causally connected. Noble aims can turn egoism into something approaching altruism. The good, she says, is 'what makes a man act unselfishly in a concentration camp'.[13]

Are there any such transforming goals? If I thought that I was created by God, that my bodily life was an illusory passage to eternal bliss, that my flourishing consisted in the extinction of my own ego, and if I had the psychological support of a community of believers living the same sort of life, then I could more easily make sacrifices that I cannot now make without great difficulty. I might also hope for some transformation of my will through divine grace. But I do not, nor do many religious believers, hope for that. My, and their, conception of human flourishing is nothing like that. Murdoch's own view of the goal of moral life is something like Plato's Form of the Good, and she sees it as having a magnetic power akin to many religious conceptions. Perhaps it is best to see what she calls the 'sovereignty' of good as something not unlike selflessness or impartiality. But that goal, though inspiring, is not inspiring enough to transform motivation in the necessary way (at least, that is what I concluded a moment ago). The goals that might transform it I see no reason to adopt; and the goals that I see reason to adopt do not transform it.

Despite all of this, it is undeniable that some rare human beings do

9

indeed sacrifice themselves for others. So they can. If they can, human beings can. Does that mean that I can? If so, the question, Ought I?, rises to challenge me. At Auschwitz Father Maximilian Kolbe volunteered to take the place of another prisoner in a punishment detail, and knowingly went to his death. But that Father Kolbe, with his religious beliefs, could sacrifice himself does not show that those with very different metaphysical beliefs can too. I doubt that we can use what a few people are capable of as evidence of what normal human beings are capable of day in, day out, which is what a moral life needs. One special circumstance would be impending disaster. We expect great sacrifices if the alternative is dire enough: I ought, I think, to accept my own death to stop a lunatic getting to the nuclear button. That the threat is so appalling should make motivation follow more naturally.

These intermediate cases, as I have called them, do not have, as the earlier cases had, an ordinary ability behind, and greatly qualifying, the inability. Common-sense ethics accepts that the mass of ordinary human beings 'cannot', that only a rare person or a person in rare circumstances 'can'. For this reason common-sense ethics does not assert that in these cases we 'ought'. It does not assert it not because, though we 'ought', there would be no point in bothering to say so. On the contrary, common-sense ethics has taken the less-demanding form it has, not because human societies reflected on the limits of human motivation and consequently pitched its moral demands at this more modest level, but because it never in the first place entertained the possibility that ordinary human beings 'could'. Still, that something is a feature of common-sense ethics does not make it right.

When I have talked about the limits of human motivation, I have admitted that there are exceptions to the rule. Some few exceptional persons – saints, heroes, those with rare qualities or in rare circumstances – might escape these limits. Then, at the other end of the spectrum, there are no doubt people who are, so far as human values go, dysfunctional or utterly indolent, and since the limits I have been speaking of arise from our having strong commitments to particular persons and causes, these indolent people may not even be subject to such limits. But my claims concern the vast majority of human beings who occupy the space between these two extremes.

8. Competing accounts

My account of permitted partiality is centered so far on facts about human capacities. It regards certain forms of partiality as falling outside morality; such partiality is allowed because it is not morally condemned.

10

There are several accounts that compete with mine, centered on values, not facts about motivation.

For instance, this account. It starts with the now well-known distinction between agent-neutral and agent-relative reasons. An agent-neutral reason is one that counts as a reason for any person – for example, not wantonly to harm others. An agent-relative reason is one that counts as a reason for a particular agent whether or not it also counts as a reason for others. Agent-relative reasons, it is thought, can arise from one's deep concerns and commitments. Thomas Nagel once drew a related distinction between personal and impersonal values.[14] Our attachments, not just to persons but also to causes, abstract standards, and institutions, express much of what we value. Some things, Nagel thought, are valuable only from the point of view of an individual's tastes, for example playing the piano well or climbing challenging mountains.

But nothing acquires value of any kind just by someone's caring deeply about it or making it a major project in life. One may want most in life to become a millionaire by the age of thirty but find, when one succeeds, that it does nothing for one. One cannot save this account just by requiring that the relevant desires be informed, that one be fully acquainted with the object of the desire. John Rawls mentions the case of a man with some crazy aim in life, say counting the blades of grass in various lawns.[15] The man accepts that no one is interested in the results, that the information is of no use, and so on. He makes no factual or logical error. It is hard to accept, however, that the fulfillment of his obsessive desire enhances his life – apart, that is, from preventing anxieties or tensions that might arise if his desire is frustrated. But anxiety and tension are not to the point; they are impersonal disvalues.

Nor does anything become valuable simply by its being what Bernard Williams calls a life's 'project'.[16] Persons often have as their 'ground project' morally hideous or merely shallow or shabby ambitions. To ask someone whose project is fueled by resentment or revenge or vanity or one-upmanship to abandon the project may, as Williams claims, be to ask a person to commit a kind of 'suicide'. But most of us would benefit from no small disintegration and reintegration. Undermining a person's 'integrity', in Williams' sense, would often not be a kind of suicide but a kind of ethical growth.

We correctly say that our major commitments express what we value in life. We also correctly say of the things that we care about deeply that we value them. But it is a non-sequitur to move from 'we value them' to 'they are valuable to us'. As my examples show, they might not be. Nor may we conclude that we have a reason, even merely an agent-relative reason, to bring about our major aims. Agent-relative reasons must

11

ultimately be sanctioned by agent-neutral reasons.[17]

Another test of adequacy for such a competing account is whether it can tell us which partialities are permitted and which are not. When the ship goes down, I may save my child instead of two strangers, but I may not, as the judge of a prize, favour my child over a more deserving stranger. I explain the difference between these cases in terms of what most human beings can do. To my mind, the competing account struggles for an explanation. I may care just as much about my child's winning the prize as surviving the sinking boat, irrational though it would often be.

9. Three plausible conclusions

What I have provided so far I offer only as raw material for a conclusion. There are, it seems to me, three different conclusions worth pondering.

(i) Not 'can't' but 'won't'. G.A. Cohen, in his recent book *Rescuing Justice and Equality*, announces that, when it comes to the more strenuous demands of equality, it is not that we *cannot* meet them but simply that we *won't*.[18] He singles me out for reproof, because I have in the past expressed scepticism about demands for complete impartiality. Cohen writes, more in sorrow than in anger, that he regrets that I should have chosen to give solace to the morally slack. He would say, I believe, that not only can we, but we ought, and failure would be blameworthy. One could take a slightly softer line: namely, we can and ought, though since most of us find it so difficult to do, we are not to be blamed if we don't.

(ii) Perhaps 'can'; perhaps 'can't'; but in any case 'not ought'. It is widely thought that morality permits, perhaps often requires, partiality. The 'intermediate' cases, as I am calling them, can largely be parceled between the permitted and the required. So what are crucial may not be empirical issues about 'can' – that is, about the extent of human ability – but independent moral issues about 'ought'. Or perhaps facts do come in – say, facts about the extreme difficulty for most persons of achieving the higher reaches of impartiality – but these facts would not occupy the deepest place in the argument. The deepest judgement may still be moral: it would be unfair, because of the great difficulty, to require such high degrees of impartiality and to blame one for failure.

12

(iii) <u>Difficulty, if great enough, turns into impossibility; and that is so in many intermediate cases.</u> There are indisputable cases of impossibility in the relevant sense: namely, the Superman sort of cases. But Kant's category may be too narrow. It would, at my age, be impossible for me to run a mile in under four minutes (it was probably always impossible for *me* to do it). But it is not beyond *human* capacity; this is not a Superman case. But it is beyond the capacity of almost all human beings; only the rare person – a Roger Bannister – can do it. Many young athletes at the very peak of both their ability and their ambition have tried their utmost but failed. Father Kolbe sacrificed his life in Auschwitz; so *he* can; so a *human being* can. But let me ask again: does this show that *I* can? Surely not. Nor does it show that most human beings can. This might be a case like running a mile in under four minutes: only rare individuals *can*. Perhaps, after all, *my* running a mile in under four minutes does not fall outside Kant's category: instead of looking at only what falls within universal adult human capacities, we should also look at the capacities of individuals and of various sub-classes of them. Perhaps the four-minute mile case is a good analogue of acts that demand the higher levels of impartiality.

How would one show that the intermediate cases are cases of 'cannot' in the relevant sense? Would the argument turn out in the end to be empirical?

It is not enough to point out that there is a biological basis for our partiality. We are by nature profoundly self-interested; we are of only very limited altruism. To become entirely impartial we should have to mobilize our beliefs and desires and motivation to a very high pitch. And that is the main obstacle to the goal. One cannot have a good life without becoming in large measure partial. And we should try to have good lives. Even if a psychotropic drug were developed that greatly reduced the difficulty of becoming completely impartial, the obstacle would remain.[19] Should one take it? Should one give up most of what makes a life worth living? Should one give up great love and affection and friendship? Should one give up deep commitment to certain institutions and causes and projects? It would not matter if the psychotropic drug were so cleverly targeted that one also did not mind whether one's life was worth living. What matters is not whether one minds it but whether it really is not worth living. Part of a life worth living is a capacity for understanding, including, importantly, understanding what makes a life worthwhile. Admittedly, one could lose

13

some of the things that make a life worth living and the life still be worth living. But if one lost one's autonomous choice of a worthwhile life and freedom to pursue it, one would have lost exactly what is usually meant by the 'dignity' of the human person.

Am I being melodramatic? Would becoming entirely impartial require stripping oneself of so much that life would no longer be worth living?

Let me quickly make two acknowledgements. First, we could certainly become much more impartial than most of us are now without losing anything that makes life worth living – anything at all. Second, we could probably become even more impartial than that, when the loss in the quality of our own lives is outweighed by the importance of what our self-sacrifice would achieve. These two acknowledgements would, I think, command wide assent. But we are interested in stronger moral demands than the two that I here accept. One stronger demand would take the form: sacrifice yourself up to the point where your loss in quality of life equals the consequent gain of the others. But this much stronger demand does not address the problem of what level of well-being one would thereby be reduced to and whether we should allow the demands on moral agents to reduce them to a life not worth living.

I think that, of the three plausible conclusions, the last is the best: 'cannot' so 'not ought'. And this seems to me, at its base, an empirical judgement, not, as in the second of the three plausible conclusions, a moral one. It is a psychological judgement about 'cannot' and a lexicographical judgement that 'ought' implies 'can'. Certainly several value judgements appear in the premises of my conclusion. First, a judgement about what makes a life good is a value judgement. Second, that there are no moral reasons not to try to live a prudentially good life is a value judgement. Third, that morality may not demand of moral agents what would make their lives not worth living is a puzzling kind of judgement, probably a value judgement. So these value judgements figure in the all-things-considered judgement: 'cannot' so 'not ought'. But they figure as establishing the absence of beliefs that would be necessary to produce the exceptional kind of motivation that might – just might – bring about complete impartiality. We see no reason to give up certain of our deep partialities. Indeed, we see strong reasons not to give them up. Reasons are the kind of thing that we merely acknowledge; they are not subject to our will. We find ourselves unable to acknowledge reasons that are necessary for us to hold if we are to have any hope of achieving complete impartiality. The final judgement here seems to be an empirical one about human capacities.

I say that this judgement seems to be empirical, but the borderline between human actions that are 'impossible' for us and those that are

'difficult in the extreme' is so unclear that any conclusion on the matter is bound to be insecure – indeed, permanently insecure. How should we incorporate this insecure conclusion into ethics? Well, the ultimate aim of ethics is not empirical truth; it is not explanation of human behaviour but, rather, its regulation. Of course, regulation of human behaviour is, in general, more likely to succeed if it is not based on falsehood. But given the situation we now find ourselves in, it seems reasonable to accept that in these intermediate cases we *can't* so it is not the case that we *ought*. But if we use the word 'reasonable', as I have just done, we must be willing to explain what the reasons are. The most important one is that it seems in fact that in most intermediate cases we *can't* reach certain levels of impartiality. And it is at this point that moral considerations may also enter. Our judgements in this domain are sunk in such deep obscurity that it would seem arbitrary, harsh, and unfair to use them as the basis for dauntingly difficult demands on humankind.

If the third conclusion is indeed right, then the principle 'ought' implies 'can' is much more central to ethics than philosophers have thought: it is at the heart of the question of the sort of impartiality on which ethics stands. That deserves more than just Kant's parenthetical mention.

However, we still have not managed to formulate the elusive principle that 'ought' implies 'can'. I have proposed cases of one kind of 'cannot':'cannot' because of limits on the will arising from commitments to particular persons, institutions, and causes. There are other potential sources of relevant kinds of 'cannot'. So my lecture is only a start.

10. Is there a role for impossible ideals?

To those who doubt that impossible ideals are of any practical relevance, Leo Tolstoy had a ready reply.[20] Ethics, he said, has two parts. There is an ethics of rules with which we are expected to, and can, comply: for example, the Ten Commandments. But there is also an ethics of ideals, such as the one Jesus set us: 'Be ye therefore perfect, as your Father which is in heaven is perfect'.[21] Of course, we cannot be perfect. But none the less this ideal, Tolstoy maintained, has an obvious role in life: we are to strive to come as close to it as we can.

Tolstoy's example of an impossible ideal comes from religion. I am interested in secular ethics. Secular ethics might regard total impartiality, for example, as an ideal. Human beings cannot, in general, achieve total impartiality. But should we, at least, strive to come as close to it as we can?

We must, in any case, take care how we formulate the secular ideal. One cannot try to do what one knows one cannot do. I know I cannot jump a mile straight up into the air without mechanical aid. If a lunatic put a

pistol to my head and demanded that I do just that, I could not even try to do it, although I would try hard to look as if I were trying.

But we could express our ideal, instead, as coming as close to total impartiality as we can. There is nothing impossible about aiming at that. The problem is, rather, that there are constraints on what can be considered ideals. Perhaps the Christian ideal of the imitation of Christ has been developed over time into an intelligible ideal. But the secular ideal of complete impartiality has not been. As we have just seen, to become as impartial as we can is to lose a lot of what makes a human life good. We should, as much as we can, have to turn ourselves into egalitarians detached from commitments to any particular persons and projects.

But most of us have to be prepared to raise children, or at least to have successful relations with other people, and more generally to be loyal and cooperative members of a community, and to care enough about our work to be productive. A few people may turn out quite different from this; a very few of them, the ones who manage to salvage some sort of sanity, might even be capable of effective impartial concern for all. But what is in the accessible psychological repertoire of a minute minority might well not be in the repertoire of the vast majority. In any case, very few of us indeed would be willing to raise our children to be utterly impartial. We should want to raise them to be capable of love and affection for those around them – that is certainly hard enough. We should not know how to produce someone emotionally detached to an extreme degree, yet sane. We are incapable of such fine-tuning. We should be too likely simply to produce an emotional wreck.

Total impartiality is not an ideal, so not an ideal that we should come as close to as we can.

Notes

1. Walter Sinnott-Armstrong, *Moral Dilemmas*, Oxford: Blackwell, 1988, pp. 6–7.

2. See also Kant, 'On the Common Saying: "This May Be True in Theory, but It Does Not Apply in Practice."': '... it would not be a duty to strive after a certain effect of our will if this effect would not be possible in experience (whether we envisage the experience as complete or as progressively approximating to completion),' introduction. For an early appearance of 'ought' implies 'can' see Aristotle, *Nicomachean Ethics* book 3, ch. 2, translated Roger Crisp, Cambridge: Cambridge University Press, 2000: '... there is no rational choice of what is impossible, and someone claiming that he was rationally choosing this would be thought a fool. ... No one ... rationally chooses things like this [immortality or that some athlete win – things that one can wish for], but only things that he thinks might come about through his own efforts.' I am grateful to Terence Irwin for this reference.

3. *Critique of Pure Reason*, A548/B576.

4. I borrow example 3 from Sinnott-Armstrong, *op. cit.*, and examples 8–10 from Alan Montefiore, '"Ought" and "Can"', *Philosophical Quarterly* 8 (1958).

5. Alan R. White offers a semantic analysis of 'ought' somewhat like mine: 'Any problem about what ought to be ... arises in a situation identified in terms of a set of circumstances, a requirement, a set of alternatives and an aspect. What ought to be is, as its etymology in several languages suggests, what amongst the alternatives is *owing* in these circumstances and under this aspect in order that the requirement be met.' (*Modal Thinking*, Oxford: Blackwell, 1975, p. 140.) I prefer my language ('what accords with the regularity or rule', 'what meets the standard', 'what realizes the aim or ideal') to his ('requirement', 'owing'). Consider: 'if you wanted to break in, you ought to have brought a jemmy.' Breaking into a house may be someone's aim, but it is not a 'requirement'; a jemmy may be helpful but it need not be necessary ('required'). Whether a jemmy is helpful or necessary, it is strained to think of it as 'owing'.

6. Sinnott-Armstrong, *op. cit.*, pp. 111–26.

7. On the differences between material implication in propositional logic and the if-then relation in ordinary speech, it is still worth reading P. F. Strawson, *An Introduction to Logical Theory*, London: Methuen, 1952.

8. *Value Judgement: Improving our Ethical Beliefs*, Oxford: Clarendon Press, 1996, ch. II. This section is a revised version of *op. cit.*, ch. VI.

9. See John Updike, *Self-Consciousness*, Harmondsworth: Penguin, 1990, p. 201.

10. For the distinction between 'crude' and 'educated' self-interest, see my *op. cit.* ch. V. sect. 3.

11. This is Shelly Kagan's answer in his *The Limits of Morality*, Oxford: Clarendon Press, 1989, ch. 8.

12. See her *The Sovereignty of Good*, London: Routledge and Kegan Paul, 1970. See also a profile of her in *The Independent*, London, 29 April 1989.

13. See profile, *op. cit.*

17

14. Thomas Nagel, *The View from Nowhere*, New York: Oxford University Press, 1986, ch. 8 sect. 5; ch. 9 sects. 1–3; *Equality and Partiality*, New York: Oxford University Press, 1991, ch. 2. I develop my objections more fully in *Value Judgement*, *op. cit.*, ch. II.

15. John Rawls, *A Theory of Justice*, Oxford: Clarendon Press, 1972, pp. 432–3.

16. Bernard Williams, *A Critique of Utilitarianism*, in J. Smart and B. Williams, *Utilitarianism: For and Against*, Cambridge: Cambridge University Press, 1973, p. 116.

17. For further discussion see my *Value Judgement*, *op. cit.*, ch. II.

18. To the question, Can we meet the more strenuous demands of equality?, G. A. Cohen, in *Rescuing Justice and Equality*, Cambridge, MA: Harvard University Press, 2008, pp. 153, 170–80, 250–54, says, Yes. I want now to give reasons to say, No.

19. I am grateful to Larry Temkin for pressing me on this point.

20. L. N. Tolstoy, *The Lion and the Honeycombe: The Religious Writings of Tolstoy*, A. N. Wilson (ed.), London: Collins, 1987, pp. 63–76.

21. Matthew 5:48.

The following lectures have been published in individual pamphlet form and may be obtained from the Department of Philosophy at a price of $5.00 plus $1.00 for handling ($6.00 per lecture).

1978   "Moral Relativism"
       By Philippa Foot, Senior Research Fellow, Somerville College, Oxford; and Professor of Philosophy, University of California, Los Angeles.

1979   "The Idea of the Obscene"
       By Joel Feinberg, Professor of Philosophy, University of Arizona.

1980   "Goods Beyond Price and Other Apparent Anachronisms"
       By Warner Wick, Professor of Philosophy, University of Chicago.

1981   "Morality, Property and Slavery"
       By Alan Donagan, Professor of Philosophy, University of Chicago.

1982   "Expressing Evaluations"
       By Donald Davidson, Professor of Philosophy, University of California, Berkeley.

1983   "How Not to Solve Ethical Problems"
       By Hilary Putnam, Professor of Philosophy, Harvard University.

1984   "Is Patriotism a Virtue?"
       By Alasdair MacIntyre, W. Alton Jones Professor of Philosophy, Vanderbilt University.

1985   "How Free Does the Will Need to Be?"
       By Bernard Williams, Provost of King's College Cambridge.

1986   "Moral Agent and Impartial Spectator"
       By Gilbert Harmon, Professor of Philosophy, Princeton University.

1987   "Projection and Truth in Ethics"
       By John McDowell, Professor of Philosophy, University of Pittsburgh.

1988   "The Politics of Imagination"
       By Arthur Danto, Professor of Philosophy, Columbia University.

1989   "Constituting Democracy"
       By David Gauthier, Distinguished Service Professor of Philosophy, University of Pittsburgh.

1990   "Justice and the Good Life"
       By Ronald Dworkin, Professor of Jurisprudence, Oxford University, New York University Law School.

1991   "Equality or Priority?"
       By Derek Parfit, Professor of Philosophy, All Souls College, Oxford, Harvard University.

1992   "Objectivity and Position"
       By Amartya Sen, Lamont University Professor and Professor of Economics and Philosophy, Harvard University.

1994   "Families, Nations, and Strangers"
       By Samuel Scheffler, Professor of Philosophy, University of California, Berkeley.

1995   "Wittgenstein on Practice and the Myth of the Giving"
       By Susan Hurley, Professor of Political and Ethical Theory, University of Warwick.

(*) Pamphlet out of print.