

XIII*—MORAL THEORY AND GLOBAL POPULATION

by Alan Carter

ABSTRACT Ascertaining the optimum global population raises not just substantive moral problems but also philosophical ones, too. In particular, serious problems arise for utilitarianism. For example, should one attempt to bring about the greatest total happiness or the highest level of average happiness? This article argues that neither approach on its own provides a satisfactory answer, and nor do rights-based or Rawlsian approaches, either. Instead, what is required is a multidimensional approach to moral questions—one which recognises the plurality of our values. Such an approach can be formalised by employing multidimensional indifference-curves. Moreover, whereas classical utilitarianism might be thought to enjoin us to bring about a larger global population, a multidimensional approach clearly suggests a significant reduction in human numbers.

Surely it is absurd to think that we can expand human numbers forever on a finite planet with finite space, finite resources and a finite capacity to absorb our pollution.¹ It would seem, therefore, that at some point in time we will have to limit the world's population. Many environmentalists agree that, because of the devastating environmental impact of human activity, we passed that time quite a while ago. Thus, the question of what the optimum human population should be is a pressing one.

However, it is also a difficult one, raising not only substantive moral issues but challenging philosophical ones as well. This is most apparent in the case of utilitarianism. For when we attempt to bring about, say, the greatest happiness of the greatest number—what hedonistic utilitarianism enjoins us to do—the

1. To illustrate the insane consequences of a policy of unrestrained procreation, Garrett Hardin once calculated that it would take little over 600 years at the then rate of population growth for the entire land surface of the planet to be occupied by human beings standing shoulder to shoulder. Cited in Robert C. Paehlke, *Environmentalism and the Future of Progressive Politics* (New Haven: Yale University Press, 1989), p. 64. And the rate of population growth has not declined significantly since then, as many at the time hoped it would.

* Meeting of the Aristotelian Society, held in Senate House, University of London, on Monday, 7th June, 1999 at 8.15 p.m.

numbers of persons involved in our calculations are usually constant. To ask questions about what number should be brought into existence and made happy raises difficult, new questions that utilitarians had, for a long time, overlooked—questions that simply do not arise when one is discussing the happiness of a constant population.

I

All consequentialist theories (including all varieties of utilitarianism) can take either negative or positive forms. Negative consequentialist theories tell us that our duty is, for example, to diminish some disvalue, rather than, positively, to increase some value. However, if we want to eliminate, say, all future human misery (an obvious example of a disvalue), then we can do that most easily by failing to bring any more people into the world.² But acting, or failing to act, so as to bring about the extinction of the human species clashes with our deeply-held values. Hence, with regard to future populations, negative consequentialist theories, and certainly the most obvious varieties, appear to have profoundly counter-intuitive implications.

Therefore, it is positive consequentialist theories that would seem to be the most promising with regard to the issue of the optimum population. Classical utilitarianism in its hedonistic version is just such a positive consequentialist theory. It enjoins us to bring about the greatest happiness of the greatest number. But this injunction appears to entail a conclusion which many find repugnant, for, as Derek Parfit points out, classical utilitarianism seems to imply that, '[c]ompared with the existence of very many people—say, ten billion—all of whom have a very high quality of life, there must be some much larger number of people whose existence, if other things are equal, would be *better*, even though these people would have lives that are barely worth living.'³

Such a conclusion seems to follow because a society of a finite size—call it 'A'—will, at any particular moment in time, contain

2. For this objection, see Robin Attfield, *The Ethics of Environmental Concern* (Oxford: Basil Blackwell, 1983), p. 116.

3. Derek Parfit, 'Overpopulation and the Quality of Life' in *Applied Ethics*, ed. Peter Singer (Oxford: Oxford University Press, 1986), p. 150.

a finite amount of happiness. But another society—say, *B*—with individuals who are less happy, if it comprises a sufficient number of them, will contain a greater total amount of happiness. But then, society *C*, though comprising people who are even less happy, as long as it comprises enough of them, would be even better than *B*. And this order of ranking would proceed all the way up to society *Z*, which comprises

an enormous population all of whom have lives that are not much above the level where they would cease to be worth living.... There is nothing bad in each of these lives; but there is little happiness, and little else that is good. The people in *Z* never suffer; but all they have is muzak and potatoes. Though there is little happiness in each life in *Z*, because there are so many of these lives *Z* is the outcome in which there would be the greatest total sum of happiness.... (The greatest mass of milk might be in a vast heap of bottles each containing only one drop.)⁴

But most of us would not regard society *Z* as better than society *A*. In fact, most of us would regard it as far worse. Thus, classical utilitarianism, insofar as it entails such a ‘Repugnant Conclusion’, also appears to have profoundly counter-intuitive implications.

In order to avoid this Repugnant Conclusion, an alternative to classical utilitarianism has been proposed—namely, ‘the Average Theory’. Whereas what has come to be called ‘the Total Theory’ enjoins us to maximise the total happiness of a population, the Average Theory enjoins us to maximise its average happiness. And in so doing, the Average Theory seems to avoid the Repugnant Conclusion, for it demands that we bring about and maintain a high level of happiness rather than a large number of people (who might be on a relatively low level of happiness).

But consider what the Average Theory seems to enjoin us to do in a situation where the human population is extremely small but its average level of happiness is exceedingly high. If the addition of extra people would not make the already existing ones even more happy, we ought not to add to the perilously small population by bringing into existence very happy people who are slightly less happy than the others, for that would lower

4. *Ibid.*, p. 148.

the average happiness. Yet this seems counter-intuitive. But even more counter-intuitive is what the Average Theory appears to enjoin us to do in a situation where everyone lives the most miserable of lives: namely, bring into existence very unhappy people, even if they would rather not have been born because their lives are so miserable, for, by being slightly less miserable than the rest, they lower the average level of unhappiness. On the Average Theory, it would seem that we ought to bring into existence as many such people as we possibly can.⁵ Thus, the Average Theory seems to fare no better than the Total Theory.⁶

Interestingly, there is, in effect, no difference between these two theories when the population remains constant. In their hedonistic versions, the Total Theory enjoins us to bring about the greatest total amount of happiness; the Average Theory enjoins us to bring about the greatest happiness per person—and to find out the average happiness, we must divide the total happiness by the number of people. If the population remains

5. See Attfeld, *The Ethics of Environmental Concern*, *op. cit.*, p. 118.

6. However, some of the difficulties facing consequentialist theories could be avoided if there is an asymmetry between the reasons for bringing happy people into existence and the reasons for not giving birth to unhappy ones. Jan Narveson, for example, argues that we cannot make a person happier by bringing him or her into existence, because it is nonsense to say that a person is happier than he or she was before being conceived. However, it would be wrong to conceive a person whom we knew was going to be unhappy. See Jan Narveson, 'Moral Problems of Population' in *The Monist*, 57 (1973), and 'Utilitarianism and New Generations' in *Mind*, 76 (1967). According to this 'asymmetrist view', if we bring into existence miserable people, then they will be actual people who can justifiably complain; whereas if we do not bring into existence happy people, there will be no persons to object on their own behalf to their not having been born. Thus, we have a duty not to bring unhappy people into existence, but we do not have a duty to conceive happy ones. Consequently, on the asymmetrist view, utilitarianism does not enjoin us to produce as many happy people as possible.

Unfortunately, Narveson's position regards as morally relevant the disbenefits unhappy people might complain about were they to become actual, while disregarding the benefits which happy people would enjoy were they to become actual. And this seems arbitrary. See T. L. S. Sprigge, 'Professor Narveson's Utilitarianism' in *Inquiry*, 11 (1968). We could turn Narveson on his head by arguing that if we bring happy people into existence, they will be actual people who can thank us, but if we do not bring miserable people into existence, they will not be able to thank us for it. Narveson's argument is arbitrary in focusing upon grounds for complaint, while ignoring grounds for gratitude. Hence, it fails to establish that utilitarianism does not enjoin us to create happy people.

Moreover, is it really nonsense, as Narveson insists it is, to claim that a person is happier than he or she was before being conceived? It can certainly be argued that we are able to benefit people by bringing them into existence. See Derek Parfit, *Reasons and Persons* (Oxford: Oxford University Press, 1987), Appendix G.

constant, then whatever increases the total happiness will also increase average happiness, and *vice versa*. The need to distinguish between these theories only arises when one considers the question of how many people we ought to bring into existence. And this is because, in some instances, one theory might recommend an increase in population, while the other might recommend a decrease.

For example, Parfit asks us to imagine a choice between societies *A* and *B*, where *A* has half as many people as *B* but the people in *A* are happier than those in *B*, though not twice as happy. On the Average Theory, society *A* is better because the average level of happiness is higher than in society *B*. Hence, on this theory, it is the smaller population which we ought to prefer. But on the Total Theory, '*B* would be better than *A* because each life in *B* would be *more than half* as happy as each life in *A*. Though the people in *B* would each be less happy than the people in *A*, they *together* would have more happiness—just as two bottles more than half-full hold more than a bottleful.'⁷ So, the Total Theory would tell us to prefer the larger population.

In short, when we consider what population size to strive for, the Total Theory and the Average Theory can lead to different recommendations. But the Average Theory only appears to be required because of the counter-intuitive implications of the Total Theory when it is applied to the problem of ascertaining the optimum population. Unfortunately, as we have seen, the consequences of the Average Theory seem to be just as counter-intuitive.

II

The problem is that both the Average Theory and the Total Theory *do* seem to pick out morally relevant considerations. When we are faced with the choice between a smaller population with a higher average level of happiness and a larger population with a lower average level of happiness, though nevertheless still at a high level, there are features of both alternatives that we like and features that we dislike. In one respect, the larger population with a lower level of happiness is worse insofar as, on average,

7. Parfit, 'Overpopulation and the Quality of Life', *op. cit.*, p. 147.

people in that society are worse off than those in the other society. It is this consideration which the Average Theory emphasises. On the other hand, the larger population is better insofar as any additional life that is worth living is a good. And it is this consideration which the Total Theory can deploy. Moreover, it is because they emphasise different considerations that the Average Theory tells us that society *A* is better than society *B* while the Total Theory tells us the converse—the latter evaluation being easily rejected because it seems to lead inexorably to the Repugnant Conclusion.

Unfortunately, because of what Parfit calls ‘the Mere Addition Paradox’, it appears that *B* cannot be worse than *A*, and this seems to make the Repugnant Conclusion unavoidable. Parfit asks us to consider a situation in which two societies are initially separated by an uncrossable ocean, thus preventing any contact between them. Society *A*—with its small population and very high average level of happiness—is one of these societies. The other society has a lower average level of happiness—a level which is, nevertheless, still high enough for its members to have lives that are well worth living. Parfit labels the situation comprising both societies ‘*A+*’. And the only differences between situation *A+* and one in which society *A* exists on its own is, first, the addition of a group of people who, while being worse-off than those in *A*, have lives that are well worth living, and, second, an inequality in *A+* which does not obtain in *A*. However, this inequality is not based on exploitation, for the two societies have no contact with each other. Rather, it is a natural inequality.

Now, can we justifiably claim that *A+* is worse than society *A* on its own? To answer in the affirmative is to imply that it would have been better if the additional group had never been born. However, as Parfit remarks, while there is an inequality in *A+*, it ‘does not seem to justify the view that the extra group should never have existed. Why are they such a blot on the universe?’⁸ Clearly, then, *A+* is not worse than *A*. But as *A+* has the lower level of average happiness, Parfit’s example shows the Average Theory to be erroneous. Note, however, that his

8. *Ibid.*, p. 152.

argument succeeds by exploiting the morally relevant consideration which underpins the Total Theory.

Parfit then asks us to imagine that, as a result of environmental changes, situation $A+$ turns into one where the average level of happiness in society A falls, while the average level of happiness in the other group rises, until they are both equal. Moreover, the average level of happiness that they both attain is greater than the average for the two societies in $A+$. Most of us would regard this new situation as better than $A+$, both because of the increase in utility and because of the greater equality. The new situation is better in utilitarian terms because the sum of the benefits minus the losses—the net sum—is greater than in $A+$. And in egalitarian terms the new situation is better, too, for it was the worst-off who received the benefits.

Hence, as Parfit observes, we could only insist that the new situation was not better if we were ‘to claim that the loss to the best-off people in $A+$ matters more than the greater gain to the equally numerous worst-off people. This seems to commit us to the *Elitist* view that what matters most is the condition of the best-off people’⁹—a view which most of us would reject. Consequently, most of us would prefer the new situation to $A+$. However, note: as the average level of happiness has clearly risen, this part of Parfit’s argument seems mainly to succeed by exploiting not only egalitarian values but also the morally relevant consideration which the Average Theory picks out.

Finally, Parfit asks us to imagine that the two societies unite and become one. There is no reason to think that this would be a worse situation, since the population size, the average level of happiness and its distribution would be unchanged. But it turns out that this new society has the same population and average level of happiness as society B . Hence, B must be better than $A+$, while $A+$ is not worse than A . Thus, B cannot be worse than A . But, Parfit insists, B does strike many of us as worse than A —which is all very paradoxical. Moreover, if B is better than A , as this Mere Addition Paradox seems to demonstrate, then we are led inevitably to the Repugnant Conclusion. How, then, can this paradox, and with it the Repugnant Conclusion, be avoided?

9. *Ibid.*, p. 153.

In attempting to avoid the Repugnant Conclusion, Parfit asks us to compare two possible futures: a person could enjoy a 'Century of Ecstasy' or, alternatively, 'live for ever, with a life that would always be barely worth living. Though there would be nothing bad in this life, the only good things would be muzak and potatoes.'¹⁰ Parfit thinks that he is not alone in preferring a Century of Ecstasy to such a 'Drab Eternity', and he prefers it because it would be qualitatively better. As he writes: 'Though each day of the Drab Eternity would have the same value for me, *no* amount of this value could be as good for me as the Century of Ecstasy.'¹¹ In short, although the quantity of pleasure in the Drab Eternity is high, it lacks the quality of pleasure enjoyed in the Century of Ecstasy. And this emphasis on quality enables an appeal not to elitism but to *perfectionism*.

In a word, Parfit does not want to claim that the move from *A* + to *B* is undesirable because, although it improves the condition of the worst-off people, it worsens the condition of the best-off, for that would be an elitist view. But he does want to argue that the move from *A* to *Z* is bad if moving from *A* to *B* costs us Mozart's music, from *B* to *C* costs us Haydn's, and so on. What is so bad about the transition from *A* to *Z* is the disappearance of the best things in life, which the benefits to those made better off do not, in Parfit's opinion, make up for. Hence, by appealing to perfectionism, Parfit can resist not only the move from *A* to *Z* but also the very first move—from *A* to *B*. But does such an appeal have any plausibility? Parfit thinks it has: 'When we are most concerned about overpopulation, our concern is only partly about the value that each life will have for the person whose life it is. We are also concerned about the disappearance from the world of the kinds of experience and activity which do most to make life worth living.'¹² And as Parfit asks rhetorically: 'without Perfectionism how can we avoid the Repugnant Conclusion?'¹³

III

Perhaps, before indicating how the Repugnant Conclusion might be avoided without appealing to perfectionism, I should mention

10. *Ibid.*, p. 160.

11. *Ibid.*, p. 161.

12. *Ibid.*, p. 163.

13. *Ibid.*, p. 164.

one possible implication of Parfit's own position. Peter Singer, in a famous article on our obligations to aid the starving in other countries, expresses his justifiable outrage at the relatively little done in the early 1970s to help eight million refugees in Bengal, who lacked food, shelter and medical care. The meagre response to this real-life situation is summed up by the aid from Australia (one of the more 'generous' countries), which donated approximately one-twelfth of the amount spent on Sydney's then new Opera House.¹⁴ Now, if, in order to preserve Mozart's music, it is necessary to build opera houses so that musicians can develop the performance skills required to keep Mozart's operas alive, then Parfit's perfectionism might require the Sydney Opera House to be valued higher than the lives of eight million refugees. But I suspect that many people would find any such conclusion just as repugnant as Parfit views society Z.

It is, perhaps, also worth noting that perfectionism might not, in fact, be required in order to provide a reason for preserving great works of art. If one were to add to the Average Theory something like John Stuart Mill's notion of 'higher pleasures', then it might be argued that the average level of happiness could only be raised above a very low level if the works of art necessary for providing such pleasures were produced and preserved. Thus, if, as Parfit insists, our concern is also 'about the disappearance from the world of the kinds of experience and activity which do most to make life worth living', then perfectionism could be considered to be redundant, for the Average Theory appears to be quite capable of incorporating such a concern.

This notwithstanding, the Average Theory, as well as the Total Theory, does seem to have counter-intuitive implications, as we have seen. And this is possibly why Parfit eschews interpreting the Average Theory in a way that would make his perfectionist alternative redundant. But is it the case that only perfectionism can avoid the Repugnant Conclusion? When we consider the real world, rather than imaginary counter-examples, the Total Theory does not, as a matter of fact, have all the counter-intuitive implications which it is usually thought to have. Parfit seems to have demonstrated that utilitarianism implies that we ought to continue to increase the world's population up to a point where

14. See Peter Singer, 'Famine, Affluence and Morality' in *Philosophy and Public Affairs*, Vol. 1, No. 3 (1972), p. 230.

everyone's life is marginally worth living—the Repugnant Conclusion. But, as a matter of fact, the Total Theory does not require us to increase the global human population, because, as Robin Attfield quite rightly points out,

[w]ith 800 million people already living in absolute poverty, additions to the population are very unlikely in most cases to increase the total [amount] of happiness, or...years of worthwhile life. If the additional people are not malnourished or miserable themselves, their arrival would in any case be likely to lengthen the world-wide food queue; and if they have more purchasing power than those least able to procure food, they will prevent yet more people in poor countries from being able to afford what they need.¹⁵

Moreover, whose happiness it is and when that happiness occurs are not relevant considerations for utilitarians. Consequently, as Attfield insightfully argues, if the Total Theory requires us to bring into existence as many happy people as possible, that doesn't mean that they have to be brought into existence now, or even in the near future:

The Total [T]heory, even if it enjoins a larger number of lives than some of its rivals, does not require them to be simultaneous. Much the likeliest way to maximise the number of worthwhile lives is to guarantee a population level sustainable into the indefinite future; and to maximise population in the short term might precisely exclude this outcome, through the exhaustion of resources, the ruination of fertile land, pollution, losses to the natural gene-pool or the breakdown of whole ecosystems on which humans ultimately depend. Thus the theory commends no higher population than the maximum sustainable....¹⁶

In short, if we consider the real world, a constant increase in population is not what the Total Theory demands, and hence it would not, in fact, result in the Repugnant Conclusion. And as Attfield remarks: 'even though the Conclusion might be yielded in some possible (ampler) world, that does not seem to be a conclusive reason for rejecting the Total [T]heory. For the alternative theories fare far worse.'¹⁷ Consequently, as Attfield regards the

15. Attfield, *The Ethics of Environmental Concern*, *op. cit.*, p. 127.

16. *Ibid.*, p. 128.

17. *Ibid.*, p. 129.

Total Theory to be the one held by classical utilitarianism, the problem of an optimum population does not, in his view, require the abandonment of utilitarianism or its revision, as many have assumed.

However, Attfield doesn't seem to have noticed that the problem of ascertaining the optimum population remains. For, within the bounds of sustainability, should we bring about the greatest total happiness or maximise the average happiness? Moreover, Attfield's insight that happiness can be maximised across time allows the possibility of a *diachronic* version of the Repugnant Conclusion. Environmentalists who insist that we ought to reduce our consumption and emit less pollution might be attacked by their opponents for demanding that we change to a lifestyle which leads to the abandonment of the good things in life merely so that countless billions of future people can enjoy the equivalent of muzak and potatoes. Thus, a perfectionist might well prefer our generation enjoying a Century of Ecstasy to humanity facing a Drab Eternity as a result of everyone being forced to live a sustainable lifestyle. So, rather than agreeing to the conservation of resources and the avoidance of unnecessary pollution, which might entail a reduction in the present generation's enjoyment, a perfectionist might, instead, declare: 'Forget our grandchildren. Let's party! Even better, let's party to Mozart!'

Furthermore, even if it is the case that the alternative theories currently subscribed to fare far worse, this would not entail that classical utilitarianism is acceptable as it stands. Perhaps we have not, as of yet, developed an adequate moral theory? And Parfit's Mere Addition Paradox does appear to show that there is *something* inadequate about both the Total Theory and the Average Theory as they stand. But it seems to do so by making use of the very values which each theory prioritizes. In the move from the situation comprising only society *A* to situation *A* +, Parfit relies on the value we ascribe to a total increase in happiness; whereas in the move to society *B*, Parfit seems to exploit the value we ascribe to an increase in the average level of happiness as well as the value we ascribe to a more egalitarian distribution. But insofar as the paradox relies on the good that each theory focuses upon, it would dissipate if we were to stress the bad that each picks out. For example, if, instead of claiming that each extra

life has value—hence we cannot regard it as a blot on the universe—we were to claim that each extra life has value only if it doesn't reduce the average happiness to an unacceptably low level, then the Repugnant Conclusion could be avoided.

In addition, the value which egalitarianism stresses, and which Parfit is content to employ, seems to tell against his own perfectionist alternative (as the case of the Sydney Opera House would appear to evince). However, John Rawls has argued that there are times when an egalitarian distribution ought to be sacrificed, but for reasons very different to Parfit's perfectionism. Quite simply, as egalitarians surely value everyone's welfare, then if an inequality would be to everyone's benefit, we ought to allow it.¹⁸ Perhaps, therefore, what we require is some mixed theory—one which takes all relevant values into consideration.

In short, each of the moral theories discussed so far is inadequate on its own. Classical utilitarianism certainly attempts to bring about something valuable. But supporters of the Average Theory rightly criticise classical utilitarians for being prepared to lower the average level of happiness to an unacceptable level, while advocates of human rights justifiably criticise them for not having safeguards at the very core of their theory which would protect individuals against their possible sacrifice for the pleasures of the majority.¹⁹ The Average Theory also attempts to bring about something valuable. But classical utilitarians rightly criticise its supporters for being prepared to lower the total quantity of happiness to an unacceptable level, while advocates of human rights justifiably criticise them on the same grounds that they criticise classical utilitarians. And egalitarian theories equally attempt to bring about something valuable. But they have been criticised for being prepared to sacrifice everyone's well-being for some rigidly egalitarian ideal.

Why, then, is each theory inadequate? I suggest: because in focusing upon only one value it ignores all others. *Each theory is*

18. For a concise summary of this argument, see John Rawls, 'Justice as Fairness' in *Philosophy, Politics and Society: Second Series*, ed. P. Laslett and W. G. Runciman (Oxford: Basil Blackwell, 1972).

19. Classical utilitarians can respond that happiness would not, in fact, be increased by sacrificing individuals, because everyone would feel insecure in a society that was prepared to act in such a way. But this reply indicates that any safeguards rely on contingent factors, rather than the theory recognising the inherent value of each individual.

designed to maximise the value it prioritises, and, in so doing, it has counter-intuitive implications. In maximising only one value, it generates consequences which offend the other values we hold. What seems to be required, therefore, is some approach to moral theory which respects all our values simultaneously. One attempt at such an approach is Rawls'.

IV

Rawls' theory of justice attempts to combine a theory of the right with a theory of the good. As there is a loss of liberty which no increase in wealth could compensate for, Rawls accords lexical priority to political liberties over 'the Difference Principle', which allows inequalities attached to offices, for which there is a fair equal opportunity of attaining, just so long as those inequalities are to the benefit of the worst-off.

Many students of Rawls have found it perplexing that he should refer to two principles of justice, when it seems clear that there are three which are lexically ordered: political liberty takes precedence over fair equality of opportunity, which takes precedence over the Difference Principle. However, there are, in fact, *four principles* which are lexically ordered. Rawls insists that his 'two' principles only apply to societies which are above a certain level of income and wealth.²⁰ Thus, a certain level of income and wealth is lexically prior to certain liberties, which are lexically prior to a certain degree of equal opportunity, which is lexically prior to further income and wealth. But when viewed in this light, Rawls' theory, with its lexical ordering, is far messier than it appears at face value, for the liberties 'prioritised' in his theory, while described as lexically prior to income and wealth, are also lexically subordinate. Hence, it is not surprising that Rawls should reduce the number of his stated principles.

It could be argued, therefore, that Rawls' lexical ordering is merely a rough and ready way of approximating to the relationship which obtains between the various values we hold.²¹ It might

20. See John Rawls, *A Theory of Justice* (Massachusetts: Harvard University Press, 1971), p. 542.

21. Indeed, Rawls admits: 'While it seems clear that, in general, a lexical order cannot be strictly correct, it may be an illuminating approximation under certain special though significant conditions.' *Ibid.*, p. 45.

also be argued that this relationship is better represented by a, now neglected, suggestion in Brian Barry's early work—a neglect resulting from Rawls' attempt to provide a rational grounding for the relationship between our various values. In *Political Argument*, Barry maintains that 'one can sensibly speak of rational choices on the basis of principles which are not all reducible to a single one provided only that the (actual or hypothetical) choices made show a consistent pattern of preference.'²² This parallels the way in which many economists understand rational consumer choice. One is consistent in one's choices if one chooses, say, four grapes and three potatoes whenever that choice is available and when one chooses five grapes and three potatoes in preference to four grapes and three potatoes. Barry, however, seeks to extend this notion of rational choice to political principles. As he writes:

Whatever may be the case with grapes and potatoes this idea seems to me eminently suitable for application to political principles. Suppose that we imagine there to be only two very general principles which we may call 'equity' and 'efficiency'.... Then for each person who evaluates in terms of these principles we can draw up a set of indifference curves showing along each line different combinations of the two between which he [or she] would be indifferent.²³

In short, equity and efficiency can, within certain limits, be substituted for each other. Put another way, one can be indifferent between a certain loss in efficiency and a specific gain in equity, and *vice versa*. Thus, in Barry's view: 'The problem of someone making an evaluation can...be regarded as the problem of deciding what mixture of principles more or less implemented out of all the mixtures which are available would be, in his [or her] own opinion, best.'²⁴

This way of conceiving the relation between our principles might allow us to make better sense of Rawls' lexical ordering. Were we to produce an indifference-curve for income and liberty, a level of income which fell beneath that indifference-curve

22. Brian Barry, *Political Argument* (London: Routledge, 1965), p. 4.

23. *Ibid.*, p. 5.

24. *Ibid.*, p. 6.

would indicate that no gain in liberty could compensate for such poverty. But above that level of poverty, we could also model Rawls' prioritisation of liberty over further gains in income by means of an indifference-curve. If income were measured on the X-axis and liberty on the Y-axis, then the result would be, when a certain level of poverty had been passed, a near horizontal indifference-curve. Moreover, a level of liberty which was too low to fall on the indifference-curve would indicate that no increase in income could compensate for it.²⁵

Now, it could be argued that the truth in Rawls' theory merely consists in the fact that there are levels of income and liberty which are so low that no gain in the other can compensate for them. But if this is so, then neither liberty in general nor income in general are accurately characterisable as lexically prior. Rather, it might be claimed, it is simply the case that there are levels of each which no increase in the other, no matter how great, can amend.

Rawls agrees that there are losses of liberty which no amount of wealth could outweigh. However, he dismisses the view that, in the society he advocates, gains in income and wealth might be rationally preferred to some losses in liberty:

in a well-ordered society...its members take little interest in their relative position as such.... [T]hey are not much affected by envy and jealousy, and for the most part they do what seems best to them as judged by their own plan of life without being dismayed by the greater amenities and enjoyment of others. Thus there are no strong psychological propensities prompting them to curtail their liberty for the sake of greater absolute or relative economic welfare.²⁶

Meanwhile, back in the real world, the problem of a growing human population is pressing. And while Rawls' lexical ordering coheres with our finding a certain lack of liberty and/or a certain level of poverty unacceptable, it fails to acknowledge the degree to which many people are, and in other circumstances (particularly where some people would find themselves having to work

25. Interestingly, while commenting on the employment of indifference-curves as a means of relating together political principles, Rawls acknowledges that 'there is nothing intrinsically irrational about this intuitionist doctrine. Indeed, it may be true.' Rawls, *A Theory of Justice*, *op. cit.*, p. 39.

26. *Ibid.*, pp. 543-4.

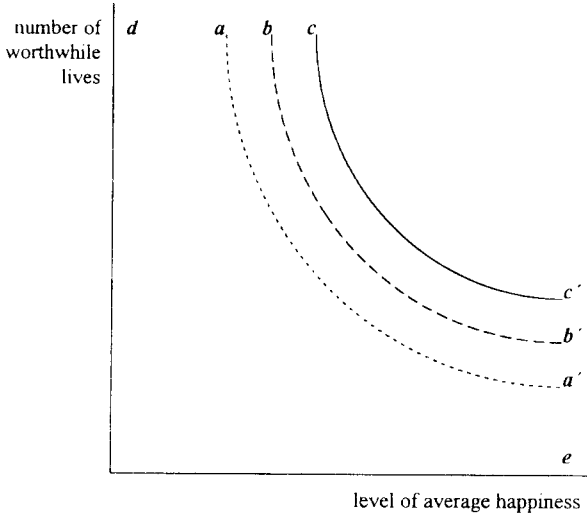


Figure 1. A Utilitarian Indifference-Map.

for others) probably still would be, prepared to trade liberty for income. In the real world outside of Rawls' mythical well-ordered society, such a trade-off is made all the time. Employees give up all kinds of liberties in order to gain income—for example, the freedom of speech with respect to industrial secrets or, if one is an employee of the British state, the signing of the Official Secrets Act. It might be argued, therefore, that indifference-maps better represent how our values actually combine together than Rawls' lexical ordering.

V

If indifference-curves are employed in order to represent the relationship between our various values, then it soon becomes apparent why both the Total Theory and the Average Theory are inadequate. Each theory focuses upon one value and attempts to maximise it. The Total Theory seeks to maximise the number of worthwhile lives—something we value—but it does so by ignoring the level of average happiness—which we also value. The Average Theory is the exact converse. One possible relationship between this pair of values is mapped out in Figure 1, where an

individual is indifferent between specific numbers of worthwhile lives and certain measures of their average happiness which fall along curve aa' . An individual is also indifferent between the points on curve bb' , but would prefer any point on bb' to any point on aa' , and any point on cc' to any point on bb' .

In short, it can be argued that there is a smaller total number of future worthwhile lives that an increase in average happiness would compensate for, and a greater number of future worthwhile lives that would off-set a decrease in average happiness. The latter would certainly seem to be the case, for we would not consider acceptable such a small number of people at an incredibly high level of happiness that the extinction of our species was threatened. Point e represents such an unacceptably small number of people living at an unimaginably high level of happiness—perhaps one generation enjoying a Century of Ecstasy. But we also consider a very large number of people living at a very low level of happiness—the Repugnant Conclusion—unacceptable. This is represented by point d .

However, if there are acceptable trade-offs between average happiness and the total number of people, then they would fall upon curves similar to aa' , bb' and cc' , with the latter indifference-curve being the most preferable because it is furthest from the origin. And the most significant feature of such curves is that *none touches an axis*. What this suggests is that Parfit is mistaken in thinking that in order to avoid the Repugnant Conclusion—in order to avoid moving from society A to society Z —we must resist the move from society A to society B . That I would not agree to a million grapes and no potatoes does not entail that I would not swap my five grapes and three potatoes for one hundred grapes and two potatoes. While society Z certainly falls beneath any of our indifference-curves, society B might not.

But the Total Theory and the Average Theory are not inadequate only because they leave out of account the value which the other theory seeks to maximise. There are other values which they downplay. Consequently, what we seem to require is something more than a two-dimensional indifference-curve. Many of us would not be happy with a large number of people enjoying an extremely high average level of happiness which involved a majority living miserable lives. Hence, we also might need to factor-in distribution patterns. Such a three-dimensional indifference-curve is represented in Figure 2. But we also value freedom,

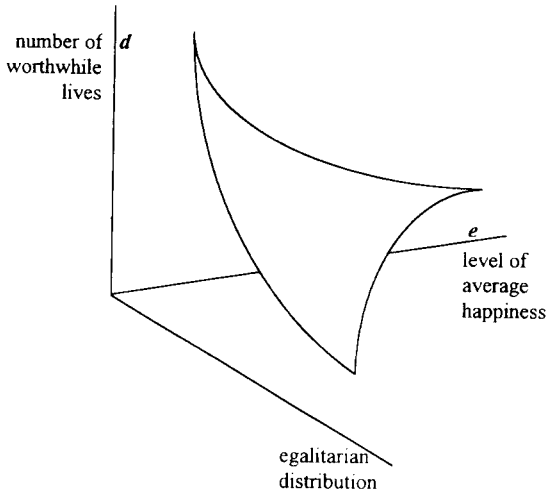


Figure 2. A Three-Dimensional Indifference Map.

and that might not be regarded as an element of well-being. Hence, we might need to factor-in liberty, as well. We might also require an axis which measures rights violations (inversely). In short, all of this suggests that *any practicable approximation to the moral ideal is an outcome the evaluation of which falls on a multidimensional indifference-curve relating together all our seriously held values, and located as far as possible from the origin.*²⁷

What, then, might such an ideal be like? It is an ideal for persons, because they matter. But if persons matter, so do future persons. And what do we want for all persons? The good life (or the primary social goods needed for them to live out their own conception of it). But if the good life applies to everyone, then the only acceptable versions of it are those that do not prevent future generations from sharing in it. And which values do we need to accommodate within our vision of the good life? Minimally, it would seem that we need to include the right to live a good life and a high average happiness for those, comprising a large number, living it. In short, the moral ideal is a vision of the

27. I apologise for not providing a pictorial representation of a multidimensional indifference-curve—it would be rather difficult to draw! However, mathematicians are not unfamiliar with multidimensionality.

good life. But if there are to be a large number of people across time, and if future persons matter, then it must be a *sustainable good life*.

Consequently, given the plurality of our values, *the moral ideal can be argued to consist in everyone having the primary social goods necessary for living a sustainable life, with maximum liberty, that is shared by very many people (quantified diachronically), all of whom are recognised as having inherent value, and which is so happy that those living it flourish*—flourishing requiring a level of happiness which includes higher pleasures. Such an ideal combines what is of value in utilitarianism, Kantianism, Aristotelianism and Rawlsianism. The first three approaches are attractive insofar as they seek to maximise a moral value. But each of these theories, in isolation, is repellent insofar as, in seeking to maximise one value, it fails to recognise other values. Each gives rise to outcomes the evaluation of which falls beneath our indifference-curves. Rawlsianism, on the other hand, is attractive in seeking to combine several values. But it seems to fail to combine them satisfactorily.

What we appear to require, therefore, is a moral theory which recognises the plurality of our values. But it is difficult to see how a moral theory could successfully combine utilitarianism, Kantianism and Aristotelianism as theories. So, it must combine the values the individual theories seek to maximise at the level of those values themselves. And if it cannot maximise all values simultaneously, then they must be traded-off. One way that we can represent such value trade-offs is by means of indifference-curves representing the values our present moral theories seek to maximise.

In a word, given the plurality of our values, it would seem that anything short of a multidimensional approach to morality is bound to be inadequate. And while most moral theories maximise one of our values, the theoretical approach sketched here differs from (and goes beyond) them in formalizing the underlying structure which relates our multiple values shaping our intuitive responses. In so doing, it also goes beyond traditional versions of intuitionism, which merely take our intuitive responses as given. By seeking to uncover the relations between our various values, the approach offered here draws our attention to the way that values can be traded off. In particular, it

focuses our attention upon their being tradeable only within certain boundaries—whereas traditional versions of intuitionism tend to ignore this underlying structure shaping our responses. Hence, traditional versions of intuitionism are often unable to reach certain substantive conclusions—conclusions which can only be obtained by paying attention to the limits set by this underlying structure. Put another way, while traditional versions of intuitionism seem to restrict our ability to engage in moral argument (and thus, perhaps, limit our ability to reach agreement), the theoretical approach adopted here increases our ability to argue for substantive moral conclusions—as I now show.

VI

In order to establish that a multidimensional approach to morality along the lines suggested here does, indeed, enable arguments to be developed which generate substantive moral conclusions, I return to the question of the optimum global population.

We have seen that the problem of plural values is particularly apparent when trying to ascertain the optimum global population, and that our subscribing to several values seems to pose insurmountable difficulties for the moral theories which currently prevail. In response to these difficulties, I have suggested that a morally acceptable outcome can be construed as one which falls on a multidimensional indifference-curve relating all our seriously held values. Thus, the axes might measure: (1) rights violations (inversely); (2) liberty; (3) distribution; (4) the number of worthwhile lives; and (5) the average level of happiness. Is there an optimum population which our various values endorse?

(1) Many of us believe that the best possible world would contain, *inter alia*, no rights violations. Rights derive their moral force from the inherent value of all persons. However, which human persons matter?²⁸ There is reason to think that it is both those who presently exist and those who will exist.²⁹ But then,

28. For reasons of simplicity, I shall not address the question of whether or not there are non-human persons who possess rights. See, however, Tom Regan, 'The Case for Animal Rights' in *In Defence of Animals*, ed. Peter Singer (Oxford: Blackwell, 1985).

29. For example, see Mary Anne Warren, 'Do Potential People have Moral Rights?' in *Obligations to Future Generations*, eds. R. I. Sikora and B. Barry (Philadelphia: Temple University Press, 1978).

surely the exercise of certain mooted rights by our generation could violate similar or more basic rights of future generations?³⁰ For example, the mooted right to procreate, if exercised without restraint, could conceivably lead to such numbers of people so polluting the world (perhaps by contaminating the environment to such an extent that future generations were made infertile) that future generations would be prevented from procreating. This might be regarded as a violation of their right to procreate.

On the other hand, it might not. For it could be argued instead that, while future generations would still retain the formal right to procreate, they would merely lack the power to do anything with their right (just as the poor supposedly have the right to stay at the Hilton, but lack the financial power to exercise that right effectively). Thus, strict deontologists holding a rights-based theory might continue to insist that we have a right to procreate. Moreover, they might further insist that such a right is not defeasible, and that it is never permissible to violate a person's rights.³¹

In their favour, rights-based theories, like utilitarian theories, attempt to bring about something valuable. They seek to maximise the inviolability of every single person who has inherent value. However, insofar as the unconstrained exercise of an ostensible right to procreate might lead to so much pollution that future generations suffered unimaginably, then the harm caused to future generations by our otherwise rightful actions could be argued to place clear constraints on any such mooted right. Utilitarians of all persuasions are surely justified in criticising certain advocates of deontological theories for being prepared to sacrifice the happiness of millions for the sake of one individual. Any deontological theory which takes an individual's duty or the respecting of a right to extremes can be argued to conflict with our values. We all, perhaps with regret, take decisions which

30. Regarding the rights of future generations, see Joel Feinberg, 'The Rights of Animals and Unborn Generations' *Philosophy and Environmental Crisis*, ed. W. T. Blackstone (Athens, Georgia: University of Georgia Press, 1974).

31. For example, Nozick implausibly regards rights as side-constraints such that it is never permissible to violate a right, not even to prevent greater rights violations. See Robert Nozick, *Anarchy, State and Utopia* (New York: Basic Books, 1974), p. 29. For one critique of Nozick, see Alan Carter, *The Philosophical Foundations of Property Rights* (Hemel Hempstead: Prentice-Hall, 1989), Ch. 4.

benefit the majority at the expense of some individuals. Who would vote for a political party that wished to devote all resources (beyond those which would keep the majority barely alive) to research into cures for diseases that affected one person in a billion? All of us thus seem prepared, within limits, to sacrifice the rights, even the lives, of some small minority in certain circumstances. And compared to the billions who would live after us if we were to safeguard the environment, we are a tiny minority. Not to trade-off some of our ostensible rights for the welfare of countless future people would strike many of us as monstrous. It seems quite clear, therefore, that the exercise of our supposed rights to a degree that would lead to all future generations facing untold misery falls beneath any reasonable indifference-curve.

(2) Discourse about liberty often has much in common with discourse about rights. For many of us, the best possible world would contain, *inter alia*, maximum liberty for all. However, the un(self-)restrained exercise of one person's liberty can easily infringe that of others. And if we do not restrain ourselves now, then we are likely to restrict the liberty of future generations. Again, a trade-off seems to be required. But far more important is the possible trade-off that can be made between liberty and well-being. And any trade-off between the well-being of all future generations and our own freedom which left us with all of the freedom and them with a nightmarish existence would appear to fall beneath any reasonable indifference-curve.

(3) Many of us believe that the best possible world would be an egalitarian one. But if all persons matter, then (with obvious exceptions, such as advances in knowledge) an egalitarian distribution should not just be within a generation but also between generations. Put another way, if future people matter, shouldn't egalitarians, on pain of inconsistency, be concerned with equality across generations? But that would surely require us to be far less profligate in our consumption of resources. Moreover, equality across generations would also require us not to give birth to so many people that it would be impossible for them to live as we do.

However, against equality it can be argued, as Rawls does, that we should be prepared to sacrifice equal distributions if an inequality would be to everyone's benefit. In other words, equal

distributions can be traded for increased income. But it is improbable in the extreme that an increased income for the affluent of today would benefit future generations, who would most likely face, in consequence, a severe shortage of resources and a heavily polluted environment. In which case, the present affluence of the industrialised countries at the expense of both the poor in the Third World today and all future generations would not fall on any reasonable indifference-curve. And it is even more improbable that everyone today enjoying even more offspring would be to the benefit of all individuals within future generations, when many of them might turn out to prefer not to have been born. Hence, our enjoying unrestrained procreation would fall beneath any indifference-curve mapping reasonable trade-offs between equality and income.³²

(4) Many of us value every extra worthwhile life. But we can best maximise the number of worthwhile lives diachronically rather than synchronically. However, if environmentalists are right about the seriousness of our impact on planetary life-support systems, then because of the potential environmental damage that is likely to be caused by very large numbers of humans living at any one time, it would seem that we can, practicably, only maximise human numbers in total by having a relatively small human population at any particular time—a population which is small enough to live sustainably (i.e., without causing irreparable damage to the biosphere). Of course, few people, if any, would find it acceptable for the population to be so small that humanity was threatened with extinction. But, in present circumstances, extinction seems far more likely if we try to maintain a large population.

(5) Many of us who value every extra worthwhile life also believe that the best possible world contains a very high level of average happiness. But if all people count, the average happiness across time counts. But the average happiness across time appears to be threatened by the affluent of today—because of the misery the consequences of their affluence are likely to cause

32. It might be asked: What has family size to do with income? Well, I might choose to have a large family in order to increase my income (for example, by having sons and daughters who could work on my farm, say, or support me in my old age). But in a world of finite resources, such a policy is highly unlikely to be to the benefit of all generations.

future generations. It is also threatened by our giving birth to so many people that the most likely result will be that there will remain insufficient resources to go round and that their activities will cause such a polluted environment that countless billions will suffer horribly.

Unfortunately, in valuing both an increase in the number of worthwhile lives and a high average level of happiness, we face precisely those difficulties encountered by the Total Theory and the Average Theory—difficulties which seem to require a trade-off between our values. But, this said, have we any clear idea of what size of population we ought to strive for? One thing seems clear: we ought to live sustainable lifestyles. But the question remains whether we should maximise the number who, at any one time, live worthwhile lives in a sustainable manner or whether we should maximise their average happiness, instead.

An indifference-curve (such as Figure 1) relating average happiness to the total number of worthwhile lives might help us answer this question. As there would be a sufficiently large number of people over time were we to enjoy sustainable lifestyles, then a far smaller human population at any one time than exists at present would not fall beneath a reasonable indifference-curve. Thus, it can be argued that any trade-off between the total number of worthwhile lives in the *near future* and the average happiness of those living them should be weighted in favour of the average happiness of those living sustainably. In a word, it can be argued that the optimum population is not as many humans as can be fitted into a sustainable lifestyle (as Attfield's acceptance of the Total Theory might suggest), but a far smaller population, living a sustainable lifestyle, which is secure from extinction (thus guaranteeing a very large population across time), while enjoying (also across time) the minimum of rights violations, the maximum liberty, and the most equal distribution of benefits insofar as they are all compatible with the greatest average happiness.

In conclusion, considerations of how to minimise rights violations, maximise liberty, equalise distribution, maximise the number of worthwhile lives and maximise the average level of happiness offer, in isolation, some reasons for claiming that one of the most urgent moral requirements is to reduce significantly

the size of the present global population.³³ But when due attention is paid to how our values are structured, it becomes clear that the limitations we place on the various acceptable trade-offs (which can be mapped using multidimensional indifference-curves) provide the basis for a more complete argument for the immorality of unrestrained procreation on a finite planet. Moreover, a consideration of how our various values are structured helps us towards a clearer picture of the size of the optimum human population.

Heythrop College
University of London
Kensington Square
London W8 5HQ
E-mail: a.carter@heythrop.ac.uk
Website: www.heythrop.ac.uk/carthome.htm

33. However, I should add that, in my view, it is extremely doubtful that a smaller population living sustainable lifestyles can be achieved by either authoritarian, Marxist or reformist means. See Alan Carter, *A Radical Green Political Theory* (London: Routledge, 1999).