

Future Generations, Sufficiency, and Biotechnology

ROGER CRISP

*St Anne's College, Oxford
Oxford Uehiro Centre for Practical Ethics*

ABSTRACT

Many people believe in egalitarian principles of justice, according to which inequality is bad in itself. These principles can of course be applied to the relation between our generation and future generations. I shall argue that equality is not a value, and that what we should care about is giving some priority to promoting the well-being of the worse off. But this priority should be given only up to a certain threshold—the threshold at which people have *enough*. I shall then consider the implications of this sufficientarian version of the priority view for the ethics of future generations, and for the development of biotechnology.

1. POPULATION ETHICS AND UTILITARIANISM

Population ethics is not new. In the late eighteenth and early nineteenth centuries, for example, Thomas Malthus and others were worrying about the implications of unrestricted population growth for existing and future generations. But in recent decades there has been a huge increase in the attention paid by philosophers to population ethics, mainly as a result of the last part of Derek Parfit's epoch-making *Reasons and Persons* (Parfit 1984). In that book, Parfit sets out many problems for various plausible principles governing population size, and it has especially intrigued philosophers that Parfit himself did not offer a solution to these problems.

The standard approach in population ethics is to apply some general principle or

set of principles, thought plausible *a priori*, to the particular question of population. One very common such principle, of course, is utilitarianism, according to which the right action is that which brings about the greatest amount of happiness overall. And the commonest version of utilitarianism is act utilitarianism, which applies the principle directly to acts, construed in terms of a maximisation of *total* rather than *average* utility. As Henry Sidgwick noticed, the difference between these two forms of act utilitarianism is important *only* when issues of population are at stake (Sidgwick 1907: 114–15). In any fixed population, maximising the total happiness will be the same as maximising the average. But total utilitarianism may require us to bring extra people into being so as to increase the total, even if this decreases the average level within the already existing population.

One common objection to the application of the total act utilitarian view to population ethics is Parfit's:

Repugnant Conclusion: For any possible population of at least ten billion people, all with a very high quality of life, there must be some much larger imaginable population whose existence, if other things are equal, would be better even though its members have lives that are barely worth living. (Parfit 1984: 388)

I believe that so-called 'discontinuities' in value can help the total utilitarian to avoid the Repugnant Conclusion. These discontinuities are especially easy to recognise in the case of pain and suffering. Imagine that some evil demon offers you a choice. You must choose between one hundred years of the most agonising torture possible, or one thousand years of a slight itch at the tip of your little finger (the irritation of which never gets any worse, causes any further unpleasant sensations, or has any other negative effect on your well-being). You choose the itch. Then the demon says that after that thousand years is up, you will have to make a similar choice, except that now the time with the itch will be much extended. Indeed the demon keeps extending the length of the itch in an effort to make you choose the agonising torture. It is plausible to claim that *no* amount of itching can be worse than a century of agonising suffering. In the same way, it could be argued, a population of whatever size, whose members have lives barely worth living, could never be more valuable than a population of ten billion people living lives of a very high quality.

But total act utilitarianism is not out of the woods.¹ Consider now the following two outcomes, in which the numbers represent the well-being of members of populations.

	<i>Existing population</i>	<i>Future population</i>
<i>Equality</i>	50	50
<i>Inequality</i>	90	10

For the sake of the argument, assume that (i) well-being can be measured (if only roughly), (ii) each population is of the same size, and (iii) the identities and desert-levels of those in each population respectively are the same in both *Inequality* and *Equality*. According to traditional total act utilitarianism, there is no reason to prefer *Equality* to *Inequality*. Many people find this objectionable, thinking that equality or fairness is relevant to evaluations of population size and relative well-being.

2. EQUALITY, PRIORITY, AND SUFFICIENCY

One apparently straightforward way to deal with this problem for act utilitarianism is to introduce an egalitarian principle, such as:

Egalitarianism: *One outcome is to be preferred to another in so far as (undeserved) inequality is minimised.*

Unfortunately, this principle itself faces a problem which Parfit has called the ‘levelling-down objection’ (Parfit 1998: 10). Consider now the following outcomes:

	<i>Rest of existing population</i>	<i>Future population</i>
<i>LD</i>	9	9
<i>Inequality₂</i>	99	100

Imagine there is some policy decision we can make which will bring down the levels of well-being of both existing and future people, so that these levels are equal.

1. Several of the following arguments are based on Crisp 2003.

According to egalitarianism, there is *something* to be said for doing this. This attachment to an alleged value which not only benefits no one but harms many (and in a sense everyone) seems very hard to justify.

Parfit has suggested that what this sort of example demonstrates is that political egalitarians have not in fact been concerned with the relative positions of individuals. Rather, they have been concerned with the absolute position of the worse off. Parfit advocates what he calls:

The Priority View: *Benefiting people matters more the worse off these people are.*
(Parfit 1998: 12)

On this view, if we consider the levels of well-being of the existing and future populations in our example, we shall prefer *Inequality*₂. The inequality in this outcome does not matter at all; what does matter is improving the position of the worse-off group (i.e. the existing population).

But Parfit's statement of the view is incomplete. How exactly does benefiting the worse off matter? One obvious interpretation of the view is:

The Absolute Priority View: *When benefiting others, the worst off individual (or individuals) is (or are) to be given absolute priority over the better off.*

Consider the following distributions, where WP is the worst-off person in the existing population, and each group contains one thousand people:

	WP	Existing population	Future population
<i>Status Quo</i>	8.9	9.1	9.1
<i>Absolute Priority</i>	9	9.1	9.1
<i>Expanded Concern</i>	8.9	100	100

The absolute priority view, in this case, favours moving from *Status Quo* to *Absolute Priority* rather than *Expanded Concern*. Because the absolute priority view is an 'innumerate' maximin principle, it will, like Rawls's 'difference principle', allow

the smallest benefit to the smallest number of worst off to trump any benefit, however large, to any but the worst off, even the next worst off. And this, it may be thought, is almost as absurd as levelling down.

What is required, then, is a principle that allows us to give priority to the worse off, but in giving priority to take into account the size of benefits at stake and the numbers of people who will benefit. The most plausible such principle, which can avoid the conclusion that large numbers of tiny benefits to the very well off can trump some smaller number of huge benefits to the very badly off, is:

The Number-weighted Priority View: Benefiting people matters more the worse off those people are, the more of those people there are, and the larger the benefits in question. But the number of beneficiaries matters less the better off they are.

But now consider an example in which you can offer fine wine to different groups of well-off individuals, now or in the future:

	<i>10 Existing people</i>	<i>10,000 Future people</i>
<i>Status Quo</i>	80	90
<i>Lafite 1982</i>	81	90
<i>Latour 1982</i>	80	91

Here it seems somewhat absurd to think that the existing people should be given any priority over future people. Indeed, what the wine example brings out is that, once recipients are at a certain level, any prioritarian concern for them disappears entirely. This implies that any version of the unrestricted priority view must fail: when people reach a certain level, even if they are worse off than others, benefiting them does *not*, in itself, matter more. What is required is an account that incorporates a threshold above which priority does not count, but below which it does—and we may assume that it will be priority that takes into account both size of benefits and numbers of recipients, so as to avoid the problems of the absolute priority view, as well as how badly off those below the threshold are. What is to happen above the threshold? The placing of the threshold might well be understood as a tempering of act-utilitarian accounts of distribution, so above the threshold goods and bads should be distributed so as to maximise well-being impartially.

Where is the threshold? This is a very difficult question. One answer that seems quite plausible to me is a life of high value that lasts eighty or so years. What seems to be important, on whatever model one adopts, is that the threshold is tied to the notion of a lack. Where the individual in question has *enough*, special concern seems to give out—though of course their well-being will play its part in the overall good. This gives us:

The Sufficiency Principle: *Special concern for any being B is appropriate up to the point at which B has a level of well-being such that B can live a life which is sufficiently good.*

So it might be thought that the principle of justice between generations that we should adopt is based on the idea that, rather than maximising overall well-being in the future, we should see it as a reasonable aim (perhaps one among others) that future people (and of course existing people too) have lives which are sufficiently good.

3. NON-IDENTITY AND PERSON-AFFECTINGNESS

Now, however, we have to remove one of the theoretical assumptions on which we have been so far relying. We have been assuming that those who exist in different possible futures are the same individuals. But of course they will not be. The very identity of future people depends on which decisions we now make. So as long as the lives of future people are better than nothing, they cannot complain if past generations have used up resources and damaged the environment. For if they hadn't done that, then these future people wouldn't exist. Some others would exist instead.

To see this, consider a case adapted from Parfit (1984: 362).

Resources. *A global community has to decide how to use its scarce resources. It has two main choices. Policy (1) will continue with rapid depletion of resources. In about three centuries, this will cause the level of well-being of those then existing to be significantly lower than current levels. But these individuals will still have lives of positive value to them. Policy (2) will involve serious conservation of resources*

and research into alternatives. This will slightly decrease levels of well-being over the next three centuries. But levels of well-being after that time will be significantly higher than they would otherwise have been.

Imagine that this community chooses policy (1). In three centuries, the individuals then alive will not be able to complain about that earlier decision. For if policy (2) had been chosen, *they* would not have existed. For the choice of such a policy has major social consequences which affect the very identities of the people who are born.

Utilitarianism has no problem with non-identity. It will advocate choosing policy (2) because all that matters is that the total level of well-being in the history of the world be as great as possible. What about prioritarianism? If we view prioritarianism as a conception of just distribution according to which we should allocate goods and bads in such a way that no one has anything to ‘complain’ of, it will be hard to apply the principle in relation to future generations. Imagine that we choose a policy of resource-depletion which has the result that one group of individuals in future do significantly worse than some other group. The worse off appear to have no complaint against us, since there is no way that *they* could have been better off. If we had chosen a different policy, they would not have existed. It may of course be the case that they have such complaints against the better off in their world. And we may decide that one reason against our choosing a policy of depletion here is that it will result in an outcome in which one group of people have a claim of justice against another group. But this seems to be moving in the direction of a non-person-affecting or impersonal version of prioritarianism, since any duty we have to avoid causing such outcomes is not plausibly understood as a duty to any particular people.

How, then, should we understand prioritarianism in relation to future generations? In a suggestive paper on the topic, Nils Holtug defines the view as follows (Holtug 2007: 10–12):

Prioritarianism. An outcome is (non-instrumentally) better, the larger a sum of weighted individual benefits it contains, where benefits are weighted such that they gain a greater value, the worse off the individual to whom they accrue.

As Holtug points out, prioritarianism so understood supports the Repugnant Conclusion even more strongly than utilitarianism, since the benefits in the world

with the much larger population go to people who are very badly off, while those in the world of ten billion people go to those who are very well off. But there is something peculiar about the view understood in this way. The cases which lead us initially to adopt prioritarianism tend to involve a choice about which existing person or persons will receive a benefit. We have some special concern for the worse off, and this leads us to give them some degree of priority over the better off. And if we give the benefit to the worse off, the better off will in these standard cases retain whatever benefits they have. But in the case of the Repugnant Conclusion, we do not have existing people who are respectively worse off and better off. We have to decide which people to bring into being. And it is not clear why we should decide that it matters more to bring into being people who will be badly off in preference to people who will be less badly off. If we bring the badly off into existence, there will be no less badly off or 'better off'.

Impersonal prioritarianism, then, should not be understood in terms of weighting benefits more the lower the absolute level of well-being of their recipient. The prioritarian is indeed concerned especially with the well-being of the worse off, but her aim is, in the right way, to raise the levels here. So, when applied to the Repugnant Conclusion, prioritarianism may in fact favour the population of ten billion, because the 'worse off' (or the 'worst off') in that outcome do so much better than their analogues in the much larger population.

The best distributive principle governing future generations, therefore, will be a form of threshold-based prioritarianism or sufficientarianism, which will maximise well-being overall unless there are individuals in an outcome falling below the sufficiency threshold. If there are, benefiting these individual will be given special weight or priority.

4. BIOTECHNOLOGY, HUMAN ENHANCEMENT, AND JUSTICE

Many existing people have complaints based on sufficientarian justice against others. One obvious example is those malnourished children in developing countries who die early as a result of easily preventable disease. But there are other ways to increase well-being than preventing disease. And this brings us to biotechnology and its potential for positively enhancing human well-being, in a non-therapeutic but still beneficial way.

I have admitted that where the threshold of sufficientarian concern lies is unclear.

But biotechnology, along with other measures, may place us in a position in which most future people have lives at a level of well-being equivalent to the *most well-off* individuals now alive. Consider, for example, memory enhancement. The first to benefit greatly from memory-enhancing drugs may well be those suffering from conditions such as Alzheimer's. But it may also be that these drugs will be developed to the point at which all human beings have memories as good as, or even better than, those with the best memories living today. On any plausible view of human well-being, it is easy to see how the possession of such a memory could increase well-being: more people could enjoy more activities, and to a greater degree; more preferences could be satisfied; greater accomplishments would be more widely available; and so on.

Or consider human relationships. Many human beings are made very unhappy through their inability to engage well in close personal relationships with others. This again seems true on any plausible account of well-being: good relationships enhance felt contentedness and enjoyment; they satisfy deep desires and needs; and they are thought by many to serve as independent constituents of well-being in themselves. It may be that germ-line genetic engineering will enable future human beings to be more co-operative, understanding, and patient with one another than many existing people. Again this could bring many people above the threshold of sufficientarian concern who would have otherwise been below it. The correct principle governing our actions in relation to future generations, then, appears to support further appropriate research and development of biotechnology.²

2. For comments on an earlier version of this paper, I am most grateful to the audience at the 10th Uehiro-Carnegie-Oxford Conference, held in Tokyo in May 2012. For very helpful advice and discussion, I wish to thank John Broome, Tom Douglas, and Guy Kahane.

REFERENCES

- Crisp, R. 2003: 'Equality, priority, and compassion', *Ethics* 113/4: 745–63.
- Holtug, N. 2007: 'On giving priority to possible future people', in Rønnow-Rasmussen et al. Available at <http://www.fil.lu.se/hommageawlodek/site/papper/HoltugNils.pdf> (accessed 19th March 2013).
- Mason, A. (ed.) 1998: *Ideals of Equality* (Oxford: Blackwell).
- Parfit, D. 1984: *Reasons and Persons* (Oxford: Clarendon Press).
- 1998: 'Equality and Priority', in Mason 1998: 1–20.
- Rønnow-Rasmussen, T., Petersson, B., Josefsson, J. and Egonsson, D. (eds.) 2007: *Hommage à Wlodek: Philosophical Papers Dedicated to Wlodek Rabinowicz*. Available at <http://www.fil.lu.se/hommageawlodek/index.htm> (accessed 19th March 2013).
- Sidgwick, H. 1907: *The Methods of Ethics* (London: Macmillan).