



The Person-Affecting Restriction, Comparativism, and the Moral Status of Potential People

Gustaf Arrhenius

ABSTRACT

Traditional ethical theories have paradoxical implications in regards to questions concerning procreation and our moral duties to future people. It has been suggested that the crux of the problem resides in an all too ‘impersonal’ axiology and that the problems of population axiology can be solved by adopting a ‘Person Affecting Restriction’ which in its slogan form states that an outcome can only be better than another if it is better for people. This move has been especially popular in the context of medical ethics where many of the problems of population axiology are actualized. Examples are embryo or egg selection, pre-implantation genetic testing, assisted reproduction programmes, abortion, just to mention a few. I discuss a number of different interpretations of the Restriction and in particular one interpretation which I call Comparativism. According to this view, we should draw a distinction between uniquely and non-uniquely realizable people. The former people only exist in one out of two possible outcomes, whereas the latter exist in both of the compared outcomes. The idea is that we should give more weight to the well-being of non-uniquely realizable people or take it into account in a different way as compared to the well-being of uniquely realizable people. I argue that the different versions of the Person Affecting Restriction and Comparativism either have counterintuitive implications of their own or are compatible with traditional theories such as Utilitarianism.

KEYWORDS

Population Axiology; Future Generations; Human Reproduction; Biomedical Ethics

Introduction

It has been known now for quite a while that traditional ethical theories have very counterintuitive and paradoxical implications in regards to questions concerning procreation and our moral duties to future people. For example, Total Utilitarianism, which tells us to maximize the well-being in the world, seems to imply that we have a moral duty to procreate and that we shall try to have as many offspring as possible. It implies Derek Parfit’s well-known Repugnant Conclusion:

The Repugnant Conclusion: For any outcome where everybody enjoys very high positive welfare, there is an outcome where people have very low positive welfare which is better, other things being equal.¹

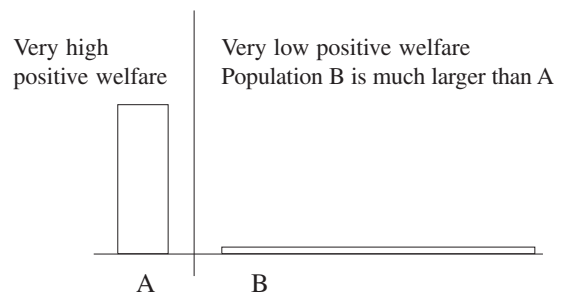


Figure 1.

In figure 1, the width of each block shows the number of people, the height shows their lifetime

AUTHOR INFORMATION: Gustaf Arrhenius, PhD, FD. Department of Philosophy, Stockholm University; Research Program Ethics in Biomedicine, Uppsala University; Centre for Applied Ethics, University of Oxford – E-mail: Gustaf.Arrhenius@philosophy.su.se



welfare. All the lives in the above figure have positive welfare, or, as we also could put it, have lives worth living. People's welfare is much lower in B than in A, since the A people have very high welfare whereas the B people have very low positive welfare. The reason for the very low positive welfare in the B lives could be, to paraphrase Parfit, that there are only enough ecstasies to just outweigh the agonies or that the good things in life are of uniformly poor quality, e.g., working at an assembly line, eating potatoes and listening to Muzak.² However, since there are many more people in B, the total sum of welfare in B is greater than in A. Hence, Total Utilitarianism ranks B as better than A, i.e., we should try to make B come about rather than A.

Notice that problems like these are not just problems for utilitarians or those committed to welfarism, the view that welfare is the only value that matters from the moral point of view, since we can assume that the other things are roughly equal. We can assume that other values and considerations are not decisive for the choice between population A and B in figure 1 (e.g., promises and rights). This is a problem for all moral theories which hold that welfare at least matters when all other things are equal. Although I shall not defend this claim here, this assumption is arguably a minimal adequacy condition for any moral theory.

A number of solutions have been suggested for the Repugnant Conclusion and its cognates. What I shall discuss here is the suggestion that the crux of the problem resides in an all too 'impersonal' axiology and that the problems of population axiology can be solved by a shift to a so-called 'person-affecting' axiology. This move has been especially popular in the context of medical ethics where many of the problems of population axiology are actualized. Examples are embryo or egg selection, pre-implantation genetic testing, assisted reproduction programmes, abortion, just to mention a few.³ More precisely, I shall focus on a view which I call 'Comparativism'. According to this view, we should draw a distinction between uniquely and non-uniquely realizable people. The

former people only exist in one out of two possible outcomes, whereas the latter exist in both of the compared outcomes.⁴ The idea is that we should give more weight to the well-being of non-uniquely realizable people or take it into account in a different way as compared to the well-being of uniquely realizable people.

Consider the following condition:

Neutrality: If there is a one-to-one mapping from outcome A to outcome B such that every person in A has the same welfare as their counterpart in B, then A and B are equally good.

Standard welfarist axiologies, such as the axiological part of Total Utilitarianism, count everybody's welfare equally and thus satisfy Neutrality. A comparativist, however, counts people's welfare differently depending on whether they are uniquely or non-uniquely realizable and thus violate Neutrality. A strict comparativist only counts the welfare of non-uniquely realizable people and completely disregards the welfare of uniquely realizable people. Some of the positions advocated in the literature are not of this kind. Rather, according to these theorists, we should count the welfare of everybody but give more weight to the welfare of non-uniquely realizable people. Accordingly, these views also violate Neutrality. Another group of theorists only counts the positive welfare of non-uniquely realizable people, but counts the negative welfare of all people. In other words, these theorists respect Neutrality in regard to populations with negative welfare. Their reason behind this move is that they try to incorporate an idea called *Asymmetry*: We have no moral reasons to create people with positive welfare, all other things being equal, but we have reasons not to create people with negative welfare, all other things being equal.

In most cases, the motivation behind drawing one or the other of the above distinctions is an idea which goes under the name of the 'Person Affecting Restriction'.⁵ In its slogan form, this view states that an outcome can only be better (or worse) than another if it is better (or worse) for people. From some of the contributions in the literature, one can



get the impression that this restriction is supposed to entail one or another of the above distinctions. How this entailment is supposed to work is by no means clear and depends, of course, on how one understands the Person Affecting Restriction. This is what I shall now turn to.

1. The Person Affecting Restriction

In its slogan form – an outcome can only be better (worse) than another if it is better (worse) for people – the Person Affecting Restriction appears reasonable. It is terribly vague, however, and open to several interpretations. It could be understood as an idea about which kind of objects have moral value, for example, that all moral values are essentially related to the interests of human beings. All moral claims would thus necessarily involve a reference to humans: Outcome A is better than outcome B since people have higher welfare in the former as compared to the latter outcome, or since in the former but not in the latter outcome people's rights are fulfilled, or in the former but not in the latter people have equal opportunities, and so forth. Examples of putative moral claims which are ruled out by this restriction would thus be: Outcome A is better than outcome B since the scenery is beautiful in the former but ugly in the latter outcome, or since the ecosystem is in balance in the former but not in the latter outcome, and so forth. Roughly, this interpretation of the Person Affecting Restriction, which we could call the Human Good Restriction, claims that two outcomes can only differ in value if they differ in regard to some aspect of human goods.⁶ This restriction is pretty reasonable and I think that much of the appeal of the Person Affecting Restriction derives from the Human Good Restriction.⁷ It is, however, clearly insufficient to yield any kind of distinction between the value of people that exist in more than one outcome, on the one hand, and uniquely realizable people, on the other hand.

One can give a stronger interpretation of the Person Affecting Restriction than the Human Good Restriction expresses. One can stress an

individualist aspect of value: All goods belong to, or are located in, individuals. As John Broome says: there are no things “such as pure communal goods, belonging to the community but not to any individual”.⁸ All moral goods are *personal goods* which, roughly, are non-relational goods, belonging to or located in individuals. Another way to put it is to say that personal goods are intrinsic properties of individuals.

Consider the following two outcomes: In A, Krister and Erik are equally happy. In B, they are both happier than in A but Krister is happier than Erik. An egalitarian might argue that B is worse, or at least in one respect worse, than A, since although both Erik and Krister are better off in A than in B, B involves inequality whereas there is perfect equality in A. One might say that B is worse in regard to one aspect of human goods, namely its distribution. “Worse for whom?” some theorists ask rhetorically. Perhaps they endorse a reading of the Person Affecting Restriction, which we could call the Personal Good Restriction, to the effect that an outcome cannot be worse than another, if it is not worse in regard to personal goods.⁹

The egalitarian concern above is grounded in a relational good: What is bad about outcome B is that one person is worse off than another person. Consequently, this concern is ruled out by the Personal Good Restriction. Since B is not worse than A in respect to personal goods, B cannot be worse than A. In other words, if we find this restriction plausible, then we have a reason for rejecting Welfarist Egalitarianism.¹⁰ The Personal Good Restriction does not, however, imply any value distinctions between uniquely and non-uniquely realizable people. It is compatible with such distinctions: one might decide, perhaps on purely intuitive grounds, that only personal goods belonging to non-uniquely realizable people count. It is, however, equally compatible with principles which do not distinguish between uniquely and non-uniquely realizable people. Total Utilitarianism, for example, entails the Personal Good Restriction.

The next step to take is to stress the individualist aspect of value even more by claiming that



morality is essentially *person comparative*: If an outcome is better (worse) than another, then it is *better for* (worse for) at least one person. We shall formulate this view with a little bit more content:

The Person Affecting Restriction

- (a) If outcome A is better (worse, equally as good) than (as) B, then A is better (worse, equally as good) than (as) B for at least one individual.
- (b) If outcome A is better (worse) than B for someone, but worse (better) for no one, then A is better (worse) than B.


Henceforth, this is the principle I shall refer to as the Person Affecting Restriction. In cases involving only the same people, this view is not very controversial. In cases involving uniquely realizable people, however, this restriction is ambiguous. An outcome A is better than B for Peter if Peter has, for example, higher welfare in A as compared to B. But what if Peter exists in outcome A but not in outcome B? Is outcome A then better than outcome B *for Peter*? This is the crux of the matter. Depending on the answer to this question, different versions of the Person Affecting Restrictions result.

2. Strict Comparativism

One possible answer to the question whether existence can be better or worse for a person is to claim that non-existence is neither better, nor worse, nor equally good as existence: non-existence and existence are incomparable in value for a person. This yields a strict version of the Comparativist view I mentioned earlier: we should disregard the welfare of uniquely realizable people, that is, people that only exist in one of the compared outcomes. Sometimes, this appears to be David Heyd's view. He argues that the welfare of future possible people has "no direct moral significance and cannot be decided in ethical terms".¹¹ Furthermore, he holds that "... the very comparison of the welfare of two possible children is based on the fallacious notion of an abstract,

impersonal quantity of happiness in the world which should be maximized".¹² He thinks that we can solve the problems in population axiology "... by simply rejecting the logical legitimacy of comparisons between the welfare of a possible population A and a possible population B (when they consist of *different* people)".¹³

This version of the Person Affecting Restriction, taken as a population axiology, is inconsistent. Consider the following case:



The x people and y people exist in outcome A (see figure 2), the y people and z people exist in B, and the z people and x people exist in C. Assume that all of these people have positive welfare, but that the y people are better off in B as compared to A, the z people are better off in C as compared to B, and the x people are better off in A as compared to C.¹⁴ Since the x people do not exist in B, B is neither worse nor better than A for them. Similarly, since the z people do not exist in A, A is neither worse nor better than B for them. However, B is better than A for the y people. Consequently, B is better than A according to the second clause of the Person Affecting Restriction. The same reasoning yields that C is better than B, and A is better than C. But if B is better than A, and C is better than B, then transitivity yields that C is better than A. Consequently, C is both better and worse than A.

Perhaps an adherent of this version of the Person Affecting Restriction could argue that we should abandon the transitivity of the relation 'is



better than'.¹⁵ Apart from the counterintuitive implications of this move, it would not help much since there are other problems ahead. Consider the following case:

The Energy Policy Case: A country is facing a choice between implementing a certain energy policy (alternative A) or not (alternative B). Were this country to implement this policy, then there would be a marginal increase in the welfare of the present people of this country (the x people). On the other hand, this increase would be greatly outweighed by the misery the waste from this energy system will cause in the lives of people in the future (the y people). The existence of these future people is contingent upon the implementation of this energy policy. If the country does not implement this energy policy, other people will exist in the future with very good lives (the z people). The advantages and disadvantages of other effects of this policy balance out.




Figure 3

Most of us, I guess, would consider outcome B clearly superior to outcome A and, since the cost to present people is marginal, we ought to realize B rather than A.

According to the strict comparativist version of the Person Affecting Restriction, A is incomparable in value to B for all the y people and z people, since they are uniquely realizable people. Consequently, outcome A cannot be equally as good as, better, worse for the y people and z people, as compared to B, according to the first clause of the

Person Affecting Restriction. Outcome A is slightly better for the x people, however, and consequently, this version of the Person Affecting Restriction ranks A as better than B and yields the wrong answer to the Energy Policy Case.

3. *Asymmetrical Comparativism*

A popular answer to the question whether existence can be better or worse for a person is to claim that a life with positive welfare is neither better nor worse, nor equally good as non-existence for a person but a life with negative welfare is worse for a person than non-existence. This answer will yield the right answer in the Energy Policy Case but it will still yield non-transitive orderings in cases like the one depicted in figure 2 above. Moreover, it will fall foul of a version of the Energy Policy Case: assume that the y people have lives barely worth living. Since outcome A is better for the x people and not worse for the y people whereas outcome B is worse for the x people but not better for the z people, outcome A is ranked better than B by this version of the Person Affecting Restriction. Again, a great loss for future people is outweighed by a marginal gain for present people.

In addition, it is not clear that it is possible to uphold this kind of asymmetry. The standard argument for why existence is not better for a person than non-existence is well expressed by John Broome:

...[I]t cannot ever be *true* that it is better for a person that she lives than that she should never have lived at all. If it were better for a person that she lives than that she should never have lived at all, then if she had never lived at all, that would have been worse for her than if she had lived. But if she had never lived at all, there would have been no her for it to be worse for, so it could not have been worse for her.¹⁶

This argument works, however, equally well against the idea that existence could be worse for someone than non-existence: if it were worse for



a person that she exists than that she should never have existed, then it would have been better for her if she had never existed. If she had never existed, then there would have been no her for it to be better for, so it could not have been better for her. Thus, it cannot be true that it could be worse for a person to exist than not to exist. In other words, it does not look possible to uphold an asymmetry here.

4. *Soft Comparativism*

This leaves us with the option of the soft comparativist: we should count the welfare of everybody but give more weight to the welfare of non-uniquely realizable people. Julian Savulescu has presented me with an interesting version of this theory.¹⁷ He suggests that it makes an outcome worse if people are worse off than they otherwise could have been. Another way to put it is to say that such people have a legitimate complaint or grievance and this makes the outcome worse. In addition to the well-being of everybody, Savulescu's proposal takes the badness of legitimate complaints, or comparative harms, in terms of well-being into account. Consider the following case from Parfit:

Two Medical Programmes: If a pregnant woman has medical condition J, which a simple treatment could cure, this will cause the child she is carrying to have a certain handicap. If a woman has condition K when she conceives a child, this child would get the same handicap. Condition K cannot be treated but disappears after two months. There are two medical programmes: Pregnancy Testing for J (PTJ) and Preconception testing for K (PCTK). In PTJ, women would be tested during pregnancy and those found to have condition J would be treated. It is predicted that if we implement PTJ, 1000 children that would otherwise have been handicapped will be born without the handicap. In PCTK, women would be tested when they intend to become pregnant, and those found to

have K would be advised to postpone conception for at least two months. It is predicted that if we implement PCTK, a 1000 children will be born without the handicap rather than a 1000 (different) handicapped children. We only have funds for one of the medical programmes. Which one should we choose?¹⁸

Since both programmes would reduce the number of handicapped children by 1000, many would, like Parfit, consider these programmes equally good. If we choose to implement PCTK, however, there will be 1000 children with a handicap that they would not have had if we had chosen to implement PTJ instead – these people are non-uniquely realizable and will exist irrespective of our choice. These people can therefore be said to have been harmed and thus have a legitimate complaint. If we choose to implement PTJ, there will be a 1000 children with a mild handicap but since these children owe their existence to our choice — they are all uniquely realizable people — they cannot be said to have been harmed or made worse off and thus do not have a legitimate complaint. Consequently, although the effect on people's welfare is the same for both programmes, PCTK is worse in one respect since it will cause people to be worse off than they could have been and thus there will be people who can legitimately complain.

This is an interesting idea but there are problems ahead. Consider figure 2 again. All the outcomes in figure 2 are equally good in respect to the amounts of people's well-being. However, since the *y* people are worse off in A as compared to B, the *y* people would have a complaint if we choose A. In this respect, A is worse than B. Consequently, all things considered, A is worse than B. The same reasoning yields that B is worse than C, and C is worse than A. But if A is worse than B, and B is worse than C, then transitivity yields that A is worse than C. Consequently, A is both better and worse than C.

There is, however, another way of explicating Savulescu's idea which does not imply non-transitive orderings. When determining the value of an out-



come we should consider both people's well-being and whether they are harmed in the sense of being worse off than they could have been. The value of an outcome is determined by the value of the total well-being in the outcome reduced by a factor that reflects whether people are harmed in the sense of being worse off than they could have been. Here is an example. Assume that we represent well-being on a numerical scale and that the total well-being of the best-off people in figure 2 is 10 units and the total well-being of the worst-off people is 5 units. The value of outcome A would then be 10 minus some factor h that represents the fact that the y people are worse off than they could have been. Intuitively, this factor should correspond to how much worse off the y people are in A as compared to B. Similarly, the value of outcome B and C would be 10 minus h . Consequently, on this view all the outcomes in figure 2 are ranked as equally good which seems to be the intuitively correct answer. However, in regard to the two medical programmes, this version of soft Comparativism would pick PTJ since the two programmes are equally good in regard to people's welfare but PCTK is worse in one respect since it will cause people to be worse off than they otherwise could have been.

Although Savulescu's theory neatly captures some intuitions regarding the two medical programmes and avoids the threat of non-transitivity, I do not think it will be of much help in regard to other problems in population axiology. A difficulty shared by all versions of Soft Comparativism is that all the problems afflicting neutral theories will reappear in the specification of the method for aggregating people's welfare: summing implies Parfit's Repugnant Conclusion; averaging implies that it can be better to create miserable rather than happy people, and so forth. For instance, assume, as we did above, that the value of an outcome is determined by the value of the total well-being in the outcome reduced by a factor that reflects whether people are harmed. In all cases involving only uniquely realizable people, this version of Soft Comparativism determines the ranking by the total sum of

people's welfare since such cases do not involve any harm. Consequently, like Total Utilitarianism, it will imply the Repugnant Conclusion in respect to future populations where there is no overlap of individuals in the compared populations. Indeed, it implies the Repugnant Conclusion even in cases that involve such overlaps and that involve great losses in the welfare of non-uniquely realizable people. Assume that k is a positive finite number that represents the weight given to harm of an individual due to the fact that she is worse off than she could have been. For any population of n non-uniquely persons with very high welfare u_1 , there is a mixed population of $n+m$ uniquely and non-uniquely realizable people with very low positive welfare u_2 such that $nku_1 < (n+m)u_2$, namely a mixed population consisting of $(n+m) > nku_1/u_2$ people with welfare u_2 . In this respect, Soft Comparativism does not constitute any kind of advance towards a satisfactory theory of the moral status of potential people as compared to neutral welfarist axiologies.

Soft Comparativism also has implications which some people might consider counterintuitive. Consider the following version of the Energy Policy Case. Assume that the total difference in well-being for the x people in the two outcomes equals the difference in well-being for the y people and z people. In other words, A and B now involve the same number of people and the same total sum of well-being. A very reasonable and modest egalitarian consideration implies that B is better than A since they are equally good in regards to the total (and average) well-being but there is perfect equality in B whereas there is inequality in A. Soft Comparativism, however, implies that A is better than B since the x people would be harmed if we were to choose outcome B rather than A since they then would be worse off than they otherwise could have been, whereas the y people and z people cannot be harmed in this way since they are uniquely realizable. Consequently, there is a tension between Soft Comparativism and a modest egalitarian consideration.

It gets worse, however. Consider the following three outcomes:




Figure 4

There is the same number of people in all three outcomes in the figure above. Everyone is better off in A as compared to B, and everyone is better off in B as compared to C. Again, the x people would be harmed if we choose outcome B since they would be worse off than they otherwise could have been, i.e., if we had chosen outcome A instead. Let h represent the total value of the harm done to the x people if we were to choose outcome B. Let d represent the total difference in well-being between the x people in B and the y people in C. The difference in value between outcome B and C will then be d minus h . Consequently, if h is greater than d , then Soft Comparativism will rank C as better than B although everyone is better off in B. I find this very counterintuitive.

5. Personal Good Restriction Regained

The last answer I will consider to the question whether existence can be better or worse for a person is the claim that existence with positive welfare is better for a person than non-existence and that existence with negative welfare is worse for a person than non-existence. As my discussion above indicates, I'm sceptical that such statements really make sense when the 'better for' relation is understood in ordinary welfarist terms, that is, in terms of what is better or worse for the individual concerned rather than in terms of what we ought to do or what makes the world better or worse.

There are some ways of analyzing the 'better for' relation that would make sense of this idea, however. For example, one could consider a state X as better for a person than state Y if this is what a benevolent impartial observer would choose for them. If I had to choose between bringing someone into existence with negative welfare or not bringing him or her into existence at all, I would of course choose the latter. Consequently, given this understanding of the 'better for' relation, one can claim that it is worse for a person to exist with negative welfare than not to exist at all, and that it is better for a person to exist with positive welfare than not to exist at all, without implying any absurdities. This answer to the question whether existence can be better or worse for a person yields a version of the Person Affecting Restriction which does not have any of the disagreeable implications of the versions discussed above. It does not have any force, however, and does not imply any value distinctions between uniquely and non-uniquely realizable people. Actually, it is just a restatement of the Personal Good Restriction discussed earlier.¹⁹ Consequently, this version of the Person Affecting Restriction is entailed by neutral theories such as Total Utilitarianism.

Conclusion

As we have seen, the different versions of the Person Affecting Restriction and Comparativism either have counterintuitive implications or are compatible with neutral theories such as Utilitarianism. The negative conclusion is that an appeal to the Person Affecting Restriction will not deliver what its proponents have hoped and will not help us solve the problems in population axiology. The positive conclusion, albeit not much of a comfort, is that we can reject the pejorative rhetoric of an 'impersonal' ethics as unfounded since we can stick to Neutrality and count everybody's welfare equally but still couch our principles in person-affecting terms.²⁰



Notes

- ¹ See D. Parfit, *Reasons and Persons* (Oxford: Clarendon Press, 1984), 388. My formulation is more general than Parfit's and he does not demand that the people with very high welfare are equally well off. Although it is through Parfit's writings that this implication of Total Utilitarianism has become widely discussed, it was already noted by Henry Sidgwick, *The Methods of Ethics* (London: MacMillan, [1967] 1907), 415, before the turn of the century. For other early sources of the Repugnant Conclusion, see C.D. Broad, *Five Types of Ethical Theory* (London: Routledge and Kegan Paul, [1979] 1930), 249-50; J.M.E. McTaggart, *The Nature of Existence* (Cambridge: Cambridge University Press, 1927), 452-3; and J. Narveson, "Utilitarianism and New Generations," *Mind* 78 (1967): 62-72.
- ² See Parfit (1984), 388, and D. Parfit, "Overpopulation and the Quality of Life," in *Applied Ethics*, ed. Peter Singer (Oxford: Oxford University Press, 1986), 148. In G. Arrhenius, "Future Generations: A Challenge for Moral Theory," Ph.D. diss. (Uppsala, 2000), I discuss different interpretations of the Repugnant Conclusion in some detail.
- ³ For a good survey of medical practices in which the problems of population axiology are actualized, see T. Hope, "Physicians' Duties and the Non-Identity Problem." Mimeo (Oxford University, 2003).
- ⁴ I've taken the term 'uniquely realizable person' from K. Bykvist, *Changing Preferences: A Study in Preferentialism*, Ph.D. diss. (Uppsala University, 1998). There are a number of related but different views that appear in the literature (albeit seldom made explicit) that most often are conflated with Comparativism and the Person Affecting Restriction (see below). Comparativism should be distinguished from 'Presentism' which draw a distinction between presently existing people and non-existing people; 'Necessitarianism' which distinguish between people that exist or will exist irrespective of how we act and people whose existence is contingent on our choices; and 'Actualism' which differentiate people that have existed, exist or who are going to exist in the actual world, on the one hand, and people who have not, do not, and will not exist, on the other. These distinctions do not amount to the same thing but there are relations among them. A presently existing person is also a necessary and actual person but not the other way around since necessary and actual people may be located in the past and the future. A necessary person is also an actual person but a future actual person may be contingent on our choice. A uniquely realizable person is also a contingent person, but a contingent person is not necessarily uniquely realizable in respect to all pairs of outcomes in a choice situation since she can exist, for instance, in two out of three outcomes. Assume, for example that a couple is deliberating whether to have a child and, as a matter of fact, they do decide to have the child (but they could have chosen otherwise). Then this child is both an actual person and a uniquely realizable person. I discuss these distinctions at length in G. Arrhenius, "Future Generations: A Challenge for Moral Theory," and in G. Arrhenius, "The Moral Status of Potential People" Mimeo (Stockholm University, 2003).
- ⁵ Temkin claims that this restriction, which he dubs 'the Slogan', is presupposed in many arguments in moral philosophy, political theory, and welfare economics. See: L.S. Temkin, *Inequality* (Oxford: Oxford University Press, 1993); L.S. Temkin, "Harmful Goods, Harmless Bads," in *Value, Welfare and Morality*, ed. R.G. Frey and C.W. Morris (Cambridge: Cambridge University Press, 1993). The term "Person Affecting Restriction", introduced by Glover might be misleading since many theorists would, sensibly I think, lessen the restriction to also include other sentient beings. See J. Glover, *Causing Death and Saving Lives* (Penguin Books, 1977), 66; Narveson (1967); N. Holtug, "In Defence of the Slogan," in *Preference and Value: Preferentialism in Ethics*, ed. W. Rabinowicz (Lund University, 1996). Below, I shall only discuss applications of the Person Affecting Restriction on human populations. Consequently, whenever I claim that a certain interpretation of the Person Affecting Restriction is reasonable, this claim only holds for human populations.
- ⁶ Perhaps it is this restriction which is at stake in Moore's criticism of Sidgwick at the turn of the century. It can be seen as a denial of Moore's idea in *Principia Ethica* that an unpopulated beautiful world is intrinsically better than an unpopulated ugly world, and a reaffirmation of Sidgwick's view that all moral goods must be of 'Human Existence'. See G.E. Moore, *Ethica* (Cambridge: Cambridge University Press, 1903), section 50; H. Sidgwick (1907), Book I, chapter IX, section 4.
- ⁷ See however endnote 3.
- ⁸ J. Broome, *Ethics out of Economics* (Cambridge: Cambridge University Press, 1999), 84; J. Broome, *Weighing Goods: Equality, Uncertainty, and Time* (Oxford: Basil Blackwell, 1991), chapters 8 and 9.



- ⁹ I have taken the term 'personal good' from J. Broome (1991), chapter 8. The Personal Good Restriction is not, however, equivalent to his principle of personal good.
- ¹⁰ *Ibid.*, 180-181, suggests a way of understanding the goodness of equality that turns it into a personal good.
- ¹¹ D. Heyd, "Procreation and Value: Can Ethics Deal With Futurity Problems?" *Philosophia* 18 (1988): 151-170, 157. See also D. Heyd, *Genethics: Moral Issues in the Creation of People* (Berkeley: University of California Press, 1992).
- ¹² He also argues against Asymmetry by claiming that it "is inconsistent with a person-affecting theory as it presupposes the comparability of non-existence with life of a certain quality".
- ¹³ Heyd (1988), 159-161 (emphasis in original). The logic of Heyd's reasoning is not completely clear to me. He claims that his view is "grounded in an 'anthropocentric' conception of value according to which value is necessarily related to human interests, welfare, expectations, desires and wishes – that is to say to human volitions" (164). How this "volitional concept of value" is supposed to generate the conclusion that "[e]xcluding the welfare and interest of future merely possible person ... is a necessary consequence of a coherent person-regarding theory of value" (161) is not spelled out in a clear fashion by Heyd. As I pointed out above, I'm sceptically inclined towards the validity of such deductions.
- ¹⁴ Temkin (1987), 168-169, uses a similar example to illustrate the intransitivity of the Person Affecting Restriction.
- ¹⁵ Another option is to claim that the only thing we can say about this case is that B is better than A for the y people, C is better than B for the z people, and so forth, and that we cannot say anything at all about the all things considered ranking of these outcomes. In other words, extensive incomparability would appear in all cases involving uniquely realizable people. Apart from counterintuitive implications of this move (it seems reasonable to claim that the outcomes above are equally good and it seems crazy to claim that the outcomes involved in the Repugnant Conclusion, and the Energy Policy Case discussed below, are incomparable), it would not be very helpful in the context of medical ethics and other practical contexts where we have to make a choice.
- ¹⁶ Broome (1999), chapter 10, 168 (emphasis in original). See also Narveson, 67: "If you ask, 'whose happiness has been increased as a result of his being born?', the answer is that nobody's has. Remember that the question we must ask about *him* is not whether he is happy but whether he is happier as a result of being born. And if put this way, we see that again we have a piece of nonsense on our hands if we suppose the answer is either 'yes' or 'no'. For if it is, then with whom, or with what, are we comparing his new state of bliss? Is the child, perhaps, happier than he used to be before he was born? Or happier than his alter ego? Obviously, there can be no sensible answer here." (emphasis in original) See also Parfit (1984), 395, 489, and Heyd (1988).
- ¹⁷ See J. Savulescu et al., "Behavioural Genetics: Why Eugenic Selection is Preferable to Enhancement," (Mimeo, Oxford University, 2003) where it is stated that "[a]ccording to a person affecting view of harm, a person is harmed by an act if she is made worse off than she would otherwise have been if that act had not been performed". Similar ideas are put forward in Hope (2003) and Meyer (2003).
- ¹⁸ Parfit (1984), 367. I have changed the wording of the example.
- ¹⁹ With one qualification: It involves a weak dominance condition which we did not include in the description of the Personal Good Restriction.
- ²⁰ I would like to thank Krister Bykvist, John Broome, Erik Carlson, Axel Gosseries, David Heyd, Tony Hope, Lukas Meyer, Wlodek Rabinowicz and especially Julian Savulescu for very fruitful discussions and comments. Earlier versions of this paper were presented at the Faculty of Philosophy, Oxford University; Department of Philosophy, University of Calgary, and at the EU Workshop *Reproductive Rights*, Universität Bremen, 2003. I would like to thank the participants at these occasions for their stimulating criticism. Thanks also to the Faculty of Philosophy and the Centre for Practical Ethics, Oxford University, for hosting me during some of the time when this paper were written. Financial support from the Swedish Foundation for International Cooperation in Research and Higher Education (STINT) is gratefully acknowledged.



Bibliography

- Arrhenius, G. "Future Generations: A Challenge for Moral Theory." Ph.D. diss., Uppsala, 2000.
- Arrhenius, G. "The Moral Status of Potential People." Mimeo, Stockholm University, 2003.
- Broad, C.D. *Five Types of Ethical Theory*. London: Routledge and Kegan Paul, [1979] 1930.
- Broome, J. *Weighing Goods: Equality, Uncertainty, and Time*. Oxford: Basil Blackwell, 1991.
- Broome, J. *Ethics out of Economics*. Cambridge: Cambridge University Press, 1999.
- Bykvist, K. *Changing Preferences: A Study in Preferentialism*. Ph.D. diss., Uppsala University, 1998.
- Glover, J. *Causing Death and Saving Lives*. Penguin Books, 1977.
- Heyd, D. "Procreation and Value: Can Ethics Deal With Futurity Problems?" *Philosophia* 18 (1988): 151-170.
- Heyd, D. *Genethics: Moral Issues in the Creation of People*. Berkeley: University of California Press, 1992.
- Holtug, N. "In Defence of the Slogan." In *Preference and Value: Preferentialism in Ethics*. Edited by W. Rabinowicz. Lund University, 1996.
- Hope, T. "Physicians' Duties and the Non-Identity Problem." Mimeo, Oxford University, 2003.
- McTaggart, J.M.E. *The Nature of Existence*. Cambridge: Cambridge University Press, 1927.
- Meyer, L.H. "Intergenerational Justice." In *The Stanford Encyclopedia of Philosophy*. Summer 2003 Edition. Edited by Edward N. Zalta. URL=<<http://plato.stanford.edu/archives/sum2003/entries/justice-intergenerational>>
- Moore, G.E. *Principia Ethica*. Cambridge: Cambridge University Press, 1903.
- Narveson, J. "Utilitarianism and New Generations." *Mind* 78 (1967): 62-72.
- Parfit, D. *Reasons and Persons*. Oxford: Clarendon Press, 1984.
- Parfit, D. "Overpopulation and the Quality of Life." In *Applied Ethics*. Edited by Peter Singer. Oxford: Oxford University Press, 1986.
- Savulescu, J., M. Hemsley, A. Newson, and B. Foddy. "Behavioural Genetics: Why Eugenic Selection is Preferable to Enhancement." Mimeo, Oxford University, 2003.
- Sidgwick, H. *The Methods of Ethics*. 7th ed. London: MacMillan, [1967] 1907.
- Temkin, L.S. "Intransitivity and the Mere Addition Paradox." *Philosophy and Public Affairs* 16/2 (1987): 138-187.
- Temkin, L.S. *Inequality*. Oxford: Oxford University Press, 1993.
- Temkin, L.S. "Harmful Goods, Harmless Bads." In *Value, Welfare and Morality*. Edited by R.G. Frey and C.W. Morris. Cambridge: Cambridge University Press, 1993.