politics, philosophy & economics article

© SAGE Publications Ltd

London
Thousand Oaks, CA
and New Delhi

1470-594X
200310 2(3) 342–381
036205

The axiomatic approach to population ethics

Charles Blackorby
University of Warwick, UK

Walter Bossert
University of Montreal, Canada

David Donaldson
University of British Columbia, Canada

abstract This article examines several families of population principles in the light of a set of axioms. In addition to the critical-level utilitarian, number-sensitive critical-level utilitarian, and number-dampened utilitarian families and their generalized counterparts, we consider the restricted number-dampened family and introduce two new ones: the restricted critical-level and restricted number-dependent critical-level families. Subsets of the restricted families have non-negative critical levels, avoid the 'repugnant conclusion' and satisfy the axiom priority for lives worth living, but violate an important independence condition.

keywords population ethics, axiomatic methodology

1. Introduction

Axiomatic approaches to the investigation of principles for social evaluation are intended to identify principles with attractive properties which are expressed as axioms. If no principle can satisfy all of the members of a set of axioms (that is, the axioms are inconsistent), trade-offs must be made and the relative desirability

Charles Blackorby is Professor of Economics, University of Warwick, UK
[email: c.blackorby@warwick.ac.uk]
Walter Bossert is Professor of Economics and CIREQ Research Fellow at the University of Montreal, Canada [email: walter.bossert@umontreal.ca]
David Donaldson is Professor Emeritus of Economics, University of British Columbia, Canada
[email: dvdd@telus.net]

of the axioms should be explored. As Thomson argues, one of the ultimate objectives of the axiomatic approach is to find the dividing line between possibility and impossibility in order to discover what can and cannot be accomplished.[1] In the presence of incompatibilities among axioms, it is therefore inappropriate to declare all principles unacceptable. In practice, social choices cannot be avoided, and they should be based on ethically appropriate principles. Thus, rather than abandoning a line of investigation if an impossibility is encountered, principles which satisfy subsets of the axioms should be identified. Even if no principle can satisfy all desirable axioms, the axiomatic method allows us to evaluate competing principles by finding the axioms that they do satisfy.

Population principles rank alternatives — complete histories of the universe from remote past to distant future — according to their goodness. In this article, we examine population principles by using axioms that are attractive on ethical grounds. Although we consider only human populations here, all the principles discussed can be extended to take account of the interests of sentient non-humans. In support of such a move, Sidgwick argues that we should 'extend our concern to all the beings capable of pleasure and pain whose feelings are affected by our conduct'.[2] Our focus on humans makes our presentation more transparent.[3]

Following standard practice, we normalize utilities so that a lifetime utility of zero represents neutrality: above neutrality, a life, as a whole, is worth living; below neutrality, it is not. For an individual, a neutral life is one which is as good as one in which he or she has no experiences.[4] Because people who do not exist do not have interests or preferences, it does not make sense to say that an individual gains by being brought into existence with a utility level above neutrality. Someone might have an attitude, such as a desire or preference, toward a world in which he or she does not exist, but could not reasonably think that this world would be better or worse for him or her. Similarly, a person who expresses satisfaction with having been born cannot be claiming that existence is better (for him or her) than non-existence. It makes perfect sense, of course, to say that an individual gains or loses by continuing to live because of surviving a life-threatening illness, say. Such a change affects length of life, not existence itself.

Our analysis does not require neutrality to be defined in terms of a life without experiences; all that matters is that a life above neutrality is, from the viewpoint of the individual leading it, worth living according to some criterion. This does not imply that the *ceteris paribus* addition of an individual with a level of lifetime well-being above neutrality is a social improvement, a requirement called 'Pareto plus' by Sikora.[5] We reject this and related axioms because we believe that individuals who do not exist do not have interests.

The axiomatic method was introduced into population ethics by Parfit,[6] who observed that classical utilitarianism leads to the *repugnant conclusion*.[7] A principle implies the repugnant conclusion if every alternative in which each person experiences a utility level above neutrality is ranked as worse than an alternative in which each member of a larger population has a utility level that is above neutrality, but is arbitrarily close to it. In that case, population size can always be used as a substitute for quality of life as long as lives are (possibly barely) worth living. An important implication of Parfit's analysis is that the repugnant conclusion is implied by any population principle that (1) declares the *ceteris paribus* addition of an individual above neutrality to a given population to be a social improvement, (2) ranks any alternative with an equal utility distribution as no worse than any alternative with the same population, the same total utility and an unequal distribution of well-being, and (3) ranks same-population equal-utility alternatives with the common utility levels. Although it does not lead to the repugnant conclusion, average utilitarianism (the main competitor of classical utilitarianism at the time) has other defects, and Parfit's criticism stimulated the search for better population principles. Avoidance of the repugnant conclusion has become an axiom that acceptable principles should satisfy.

Our investigation uses four basic axioms — strong Pareto, continuity, anonymity, and existence of critical levels — and six additional axioms. Three axioms in the latter group ask principles to ignore the utilities or existence of those who are 'unconcerned' when ranking alternatives. For any pair of alternatives, the unconcerned individuals are those who are alive and are equally well off in both. Same-number independence requires the ranking of alternatives with the same population size to be independent of the utilities of the unconcerned; utility independence extends the requirement to different population sizes; and existence independence requires the ranking of any two alternatives to be independent of the existence (and, thus, of the utilities) of the unconcerned.

For any alternative, the critical level of utility is that level which, if assigned to an added person without changing the utility levels of the existing population, creates an alternative which is as good as the original. We employ an axiom that requires critical levels to be non-negative. It ensures that a principle does not sanction the *ceteris paribus* creation of people whose lives are not worth living. Another axiom asks principles to avoid the repugnant conclusion. Our final axiom, which is called 'priority for lives worth living', requires principles to rank all alternatives in which each person who ever lives is above neutrality as better than all alternatives in which each person is below it.

We have selected a set of axioms that we find appealing, but our list is not exhaustive. As a consequence, this article is not a comprehensive survey of all the work that has been done. In addition, it is worth noting that the idea of neutrality is not needed for all of our axioms: only non-negative critical levels, avoidance of the repugnant conclusion and priority for lives worth living depend on it.

Although all the axioms are desirable on a priori grounds, a result proved by Blackorby, Bossert and Donaldson[8] implies that there is no population principle that satisfies the basic axioms, utility (or existence) independence, and priority for lives worth living and, at the same time, avoids the repugnant conclusion. It is possible, however, to find population principles that satisfy all but one of the axioms. We consider the critical-level utilitarian family (Blackorby and Donaldson[9]), the restricted critical-level utilitarian family (introduced in this article), the number-sensitive critical-level utilitarian family (Blackorby, Bossert and Donaldson[10]), the restricted number-sensitive critical-level utilitarian family (introduced in this article), the number-dampened utilitarian family (Ng[11]), and the restricted number-dampened utilitarian family (Hurka[12]). Although all of these principles rank alternatives with the same population size using utilitarianism, they belong to expanded families, the members of which rank same-number alternatives with generalized utilitarianism. Generalized utilitarianism uses a continuous, real-value function which is applied to each person's utility; the resulting transformed utilities are then summed. If the transforming function is strictly concave, generalized utilitarianism is averse to utility inequality and gives priority to the interests of those whose levels of well-being are low.[13]

The second section sets out our basic framework; the third section introduces the axioms; the fourth section presents the principles whose same-number subprinciples are utilitarian and investigates their consistency with the axioms; and the fifth section concludes. The main part of the article offers several examples; mathematical statements of the axioms, principles and theorems are in the Appendix.

2. The framework

Population principles rank alternatives (complete histories of the universe) according to their goodness. Although there are some principles that do not rank all possible alternatives, the principles that we investigate in this article provide *orderings* of alternatives: at-least-as-good-as relations that are complete, reflexive and transitive.[14] Each pair of distinct alternatives is ranked, each alternative is as good as itself, and if alternative x is at least as good as alternative y and y is at least as good as z, x is at least as good as z. Alternative x is as good as alternative y if, and only if, x is at least as good as y and y is at least as good as x; x is better than y if, and only if, x is at least as good as y and it is not the case that y is at least as good as x. Transitivity of at-least-as-good-as implies transitivity of both as-good-as and better-than and, in addition, implies that if x is better than (at least as good as) y and y is at least as good as (better than) z, x is better than z.

It is possible to apply a population principle to single periods of time. The results may be inconsistent with timeless application of the same principle, however, and may recommend killing someone if he or she experiences a low level of well-being in the period in question because of a temporary illness, say.[15] This can happen because a death that occurs just before the period begins is seen, in the period, as a change in population size rather than the shortening of a life. The same observation applies to the application of principles to histories that begin in the present.[16] We therefore focus exclusively on the social evaluation of complete histories using *lifetime* utilities, which we interpret as indicators of lifetime well-being. In that case, killing shortens a life and does not change population size. As a consequence, counter-intuitive recommendations about killing do not arise.

All of the principles considered are welfarist, using only information about the lifetime well-being (utility levels) of individuals who are alive (ever live) in the various alternatives. Although we are concerned with ranking alternatives and therefore do not need to consider uncertainty, all the population principles that we investigate can be extended so that they are capable of ranking actions whose consequences are uncertain.[17] Because welfarist principles use welfare information alone, it is important that they be partnered with a comprehensive account of well-being such as, for example, the one provided by Griffin[18] or that of Sumner.[19]

Throughout the article, we assume that individual well-being is numerically measurable and fully interpersonally comparable, an assumption that permits the largest class of principles. If individual utilities are cardinally measurable and two utility levels (such as neutrality and a utility level that represents a 'satisfactory' life above neutrality) are given particular values such as zero and 100, numerical measurability and full interpersonal comparability results.[20]

3. The axioms

We consider only population principles that satisfy four basic axioms. The first of these is *strong Pareto* and it applies to alternatives with the same population. Consider any pair of alternatives, x and y. If each person is equally well off in both, the alternatives are equally good, and if each person is no worse off and at least one person is better off in x, x is better than y. Strong Pareto is a strengthening of weak Pareto, which requires that if everyone is better off in an alternative x than in another alternative y, then x is better than y.

Our second basic axiom is *continuity*. It requires fixed-population comparisons to be gradual, without sudden jumps from better to worse. Continuity rules out principles such as lexicographic maximin (leximin). Leximin ranks x as better than y if the person with the minimum utility level in x is better off than the person with the minimum utility in y, but when the minima are equal, it ranks x as better if the second-smallest utility level is bigger. If the two second-smallest utilities are the same as well, it compares the third smallest, and so on. The axiom also rules out a principle proposed by Carlson[21] that he calls the 'combined theory'. As he points out himself, his principle also fails to satisfy weak (and, thus, strong) Pareto. The utilitarian and generalized-utilitarian same-number principles are continuous.

Our third basic axiom is an impartiality condition called *anonymity*, and it requires the rankings of alternatives to be independent of the identities of those alive. This requires that any two alternatives in which the same number of people have the same levels of lifetime utility are equally good. Anonymity ensures that individual interests receive equal treatment.

Although we find anonymity attractive, it is not needed to eliminate rules that favour groups whose composition may change over time. Suppose, for example, that alternatives are ranked by giving the utilities of actual people (those who live in all feasible alternatives) a weight of two and the utilities of potential people a weight of one and adding the resulting numbers.[22] In addition, suppose that, in alternative x, the population consists of two people with utilities of 50 and 100 who consider adding a third person. In y, utilities are 50 and 100 for the original population with a utility level of 10 for the added person and, in alternative z, utilities are 30 and 80 for the original population and 60 for the added person. Then, from the point of view of the original population, y is better than x and z is worse. But if y is chosen, the added person becomes an actual person and z is better than y, contradicting the original ranking. This means that a single ordering of the alternatives is not produced.

Welfarist population principles provide a single ordering of utility vectors of all dimensions. Adding anonymity requires that any utility vector and the vector that results from a permutation of its elements are equally good. Alternative x is better than (as good as) alternative y if, and only if, the vector of individual lifetime utilities associated with x is better than (as good as) the vector of utilities associated with y. All principles considered in this article are anonymous.

Some of our axioms refer to critical levels which are defined as follows. For any alternative, consider another with one additional person alive and suppose that each member of the common population has the same level of well-being in both alternatives. A critical level for the first alternative is a level of utility for the added person that makes the two alternatives equally good. We assume that critical levels exist for all alternatives and call our fourth basic assumption *existence of critical levels*. Although critical levels are defined by using a hypothetical comparison in which the common population is unaffected, any two alternatives with different population sizes can be ranked by using critical levels and the same-number sub-principle for the larger population size.

We assume, without mentioning it explicitly, that all of our four basic assumptions are satisfied. Any principle whose same-number rankings are utilitarian satisfies strong Pareto, continuity and anonymity, but those principles are not the only ones that do. The basic axioms allow priority to be given to the interests of those whose level of well-being is low. That is, they allow for inequality aversion.

The axioms that follow are not the only ones that have been suggested, but they are, in our view, the most important.

3.1. Independence

Suppose that a single parent has a handicapped child whose lifetime utility would be zero (neutrality) without the expenditure of additional resources. Two alternatives are possible. In the first, which we call x, resources are devoted to improving the child's well-being, resulting in utilities of 50 for the child and 60

Table 1

	Parent	First Child	Second Child	Euclid
x	60	50		same in both
y	60	0	60	same in both

for the parent and, in the second, which we call *y*, no additional resources are used to raise the level of well-being of the disabled child, but a second child is born and the same resources are devoted to it, resulting in lifetime utility levels of 60 for the second child and the parent and zero for the first child. The parent and his or her children are not the only people who ever live, however. There is one other, Euclid, who is long dead and has the same utility level in both alternatives. Table 1 summarizes this example.[23]

The parent wants to know which alternative is better. Parfit considers this example and he assumes that utility levels other than those of family members and potential members are irrelevant.[24] That assumption is satisfied if principles such as classical utilitarianism are used to rank the alternatives, but not when some other principles, such as average utilitarianism, are used. The classical-utilitarian ranking of *x* and *y* is independent of Euclid's utility level and even of his existence. But it ranks *y* as better than *x* and this contradicts the moral intuition of many.

If, however, average utilitarianism is used to rank the alternatives, the ranking of *x* and *y* is not independent of Euclid's level of well-being. If his utility level is 40, average utility is 50 in *x* and 40 in *y* and, if his utility level is -100, average utility is 10/3 in *x* and 5 in *y*: average utilitarianism declares *x* to be better if Euclid's life was good and *y* to be better if it was not. In addition, if Euclid's existence is disregarded, average utility is 55 in *x* and 40 in *y*, so *x* is ranked as better in that case.

Independence axioms require the ranking of alternatives to be independent of the utility levels and, in some cases, the existence of people whose well-being or existence is regarded as morally irrelevant. The weakest of these is *same-number independence*. It requires the ranking of any two alternatives with the same population size to be independent of the utilities of individuals who ever live and have the same utility levels in both. Same-number independence is satisfied by all the principles considered in this article.

Utility independence requires the ranking of any two alternatives to be independent of the utility levels of individuals who ever live and have the same utilities in both. It implies same-number independence, but it applies to comparisons in which population sizes are different as well as to those in which they are the same. In the example of Table 1, utility independence requires the ranking of *x* and *y* to be independent of Euclid's utility level, but not necessarily of his existence.

Existence independence requires the ranking of any two alternatives to be independent of the existence of individuals who ever live and have the same utility levels in both. It is the strongest of the independence axioms and implies the other two.[25] It allows population principles to be applied to affected individuals only. At any time, the existence of people whose lives are over in all feasible alternatives (which necessarily have a common past) can be ignored. Thus, principles that satisfy existence independence look to the future, but do not ignore the possibility that the lives of some individuals who were born in the past may continue into the present or future. For that reason, such principles do not make counter-intuitive recommendations about killing.

Which independence axiom is appropriate for population ethics? To answer the question, we consider an example. Suppose that, in the near future, a small group of humans leaves Earth on a spaceship and, after travelling through space for several generations, establishes a colony on a planet that belongs to a distant star. The colonists lose contact with Earth and, in all possible alternatives, the two groups have nothing to do with each other from then on. No decision made by the members of either group affects the other in any way.

Now suppose that the colonists are considering an important social decision and want to know which of the associated alternatives is best. If the population principle satisfies existence independence, the other group can be disregarded: the ranking of the feasible alternatives is independent of its existence and, therefore, of both the number and utility levels of its members. In this case, the population principle can be applied to the colonists alone.

If, however, the population principle satisfies utility independence, but not existence independence, the ranking of the feasible alternatives may depend on the number of people in the other group, even though the number is unaffected by the decisions under consideration. In addition, if the population principle satisfies same-number independence only, the ranking of the feasible alternatives may depend on the utility levels or number of people in the other group, as is the case for average utilitarianism.

We find existence independence ethically attractive because of examples such as the ones discussed above. A case for utility independence alone could be made, however; it would focus on the total number of people who ever live. In that case, the numbers of the long dead, of unaffected independent groups or of unaffected people in the far future could count in social rankings. Consider a two-period example of Carlson.[26] In the first period, either no one lives or a large number of people live and are well off. In the second period, there are also two possibilities: in the first, a few people live with a high quality of life; in the second, a large number live with a lower, but positive utility level. Let x and y be alternatives with no one alive in the past and the two options for period two respectively, and let z and w be alternatives with the second possibility in the past and the two options for period two. Table 2 illustrates this example.

If existence independence is satisfied, the ranking of x and y must be the same

Table 2

	Period 1	Period 2
x		10 people with utility 100 each
y		1000 people with utility 1 each
z	1000 people with utility 100 each	10 people with utility 100 each
w	1000 people with utility 100 each	1000 people with utility 1 each

as the ranking of *z* and *w*. But Carlson claims that it is not unreasonable to rank *y* above *x* and *z* above *w*. If so, existence independence must be rejected. But it does not follow that utility independence must be rejected as well. The weaker axiom allows Carlson's ranking because the period-one population size is different in the two cases. Utility independence does require the rankings to be independent of the utilities of period-one people, however.

Same-number independence has been criticized on the grounds that a principle satisfying it cannot accommodate egalitarian views. Because same-number independence is implied by utility and existence independence, the claim applies to principles that satisfy them as well. Maximin utilities, the principle that declares an alternative *x* at least as good as an alternative *y* if, and only if, the minimum utility in *x* is greater than or equal to the minimum utility in *y*, is completely egalitarian and does violate the independence axioms. However, it can be approximated arbitrarily closely by principles that do satisfy the various forms of independence.[27] If the basic axioms are satisfied, any of the independence axioms imply that same-number principles are generalized utilitarian.

A second possible concern is that the intuitive appeal of examples such as the one involving long-dead people (the Euclid example) and the spaceship example does not necessarily carry over to situations in which the conclusion of existence independence is required for arbitrary groups of people — in particular, groups of contemporaries and groups that do interact. In the presence of anonymity, however, an independence condition restricted to a particular group of people (such as the long dead) is equivalent to the corresponding stronger condition that applies to all groups, including those in which the members of different groups do interact. Thus, the stronger versions of the independence axioms can be defended by combining the weaker ones with anonymity.

3.2. Non-negative critical levels

If an alternative has a negative critical level, then, because of strong Pareto, the addition of someone whose utility is between the critical level and zero to a utility-unaffected population is regarded as good. Thus, principles with negative critical levels sometimes recommend the *ceteris paribus* creation of people whose lives are not worth living.

We therefore adopt an axiom which we call *non-negative critical levels*, and it has an obvious justification.[28] Because its critical levels are all zero, classical utilitarianism passes this test. But average utilitarianism has critical levels that are equal to average utility and, for alternatives with negative average utility, critical levels are negative. Thus, if an existing population of two people has utility levels of 60 and −68, its critical level is −4 and the addition of a person whose utility level is −2 is ranked as desirable.

3.3. The repugnant conclusion

A population principle implies the repugnant conclusion[29] if, and only if, any alternative in which each member of the population has a common positive utility level, no matter how high, is ranked as worse than an alternative in which a sufficiently large population has a common utility level that is above neutrality, but arbitrarily close to it.[30] Such principles may recommend the creation of a large population in which each person is poverty-stricken. As Heyd remarks, 'What is the good in a world swarming with people having lives barely worth living, even if *overall* the aggregation of the "utility" of its members supersedes that of any alternative, smaller world?'[31]

An axiom introduced as weak equality preference in Blackorby, Bossert, Donaldson and Fleurbaey requires any alternative with an equal distribution of utilities to be ranked as no worse than any alternative with the same population and the same total utility.[32] Any principle which satisfies our basic axioms, weak equality preference and has critical levels that are all non-positive implies the repugnant conclusion.[33] To illustrate, consider a principle proposed by Sider which he calls 'geometrism'.[34] It uses a positive constant between zero and one, which we write as k, and ranks alternatives with a weighted sum of utilities: the j^{th}-highest non-negative utility level receives a weight of k^{j-1} and the l^{th}-lowest negative utility receives a weight of k^{l-1}. Critical levels are all zero and the repugnant conclusion is avoided, but because weights on higher positive utilities exceed weights on lower ones, the principle prefers inequality of positive utilities over equality.[35] It follows that, if the repugnant conclusion is to be avoided and a preference for inequality is ruled out, some critical levels must be positive. As an example, average utilitarianism is a principle which has some positive critical levels and does not imply the repugnant conclusion.

We therefore adopt an axiom which we call *avoidance of the repugnant conclusion*. Its most obvious effect in the context of this article is to rule out classical utilitarianism.

3.4. Priority for lives worth living

The axiom of *priority for lives worth living* requires all alternatives in which each person is above neutrality to be ranked as better than all those in which each person is below it. A weaker axiom, suggested by Arrhenius,[36] requires population principles to avoid the 'strong sadistic conclusion' which obtains if, and only if, every alternative in which each person is below neutrality is ranked as better than some alternative in which each person in a sufficiently large population is above neutrality. Any principle that satisfies priority for lives worth living necessarily avoids the strong sadistic conclusion. Although the impossibility result of Blackorby, Bossert and Donaldson[37] mentioned in the introduction (see Theorem 3 in the Appendix) remains true if priority for lives worth living is replaced by the axiom that requires principles to avoid the strong sadistic conclusion, we have chosen to work with our priority axiom because we believe that it has a transparent foundation and captures the intuition that lies behind this and several related axioms.

Arrhenius has also suggested that principles should avoid the 'sadistic conclusion' which obtains if, and only if, when adding people to a utility-unaffected population, the addition of people with negative utility levels can be ranked as better than the addition of a possibly different number of people with positive utility levels.[38] We show in Theorem 1 in the Appendix that any principle whose same-number sub-principles are utilitarian and that ranks no one-person alternative above all those with larger populations cannot avoid both the sadistic and repugnant conclusions. The condition on one-person alternatives is implied by the existence of critical levels. Consequently, all of these principles that avoid the repugnant conclusion necessarily imply the sadistic conclusion. This occurs because avoidance of the sadistic conclusion requires the addition of any number of people at a positive, but arbitrarily small utility level to be ranked as no worse than the addition of a single person at an arbitrarily small, negative utility level.[39]

4. Welfarist population principles

A welfarist population principle provides a single ordering of utility vectors which is used, along with information about well-being, to order alternatives. Many principles have value functions that can be employed to perform the social ranking of alternatives. A value function assigns a number (its value) to each utility vector, and any two alternatives are ranked by comparing the values of the corresponding utility vectors in the two alternatives: if the alternatives have the same value, they are equally good, and if one has a higher value than the other, the former is better than the latter. Classical utilitarianism, for example, declares one utility vector to be at least as good as another if, and only if, the sum of utilities in the first is at least as great as the sum of utilities in the second. Thus, the sum of utilities is a value function which represents the classical-utilitarian ordering of utility vectors. All of the principles discussed in this article have value functions.

The representative utility for a utility vector is that level of well-being which, if assigned to each person, produces a vector which is as good as the original. If a principle's same-number sub-principles are utilitarian, representative utility is average utility. In addition, if a principle has a value function, it can be written in terms of representative utility and population size.[40] The value function must be increasing in representative utility, but its response to population-size increases may be positive, negative or zero for different levels of representative utility.

Carlson, Carter, and, implicitly, Parfit have suggested that value functions for principles whose same-number sub-principles are utilitarian should be expressible in terms of average utility and total utility.[41] Because population size is equal to total utility divided by average utility, any value function can be written in terms of average utility and population size or in terms of average and total utility. Carter suggests, however, that the value function should be increasing in both average and total utility. We show in the Appendix (Theorem 2) that any principle with this property has some negative critical levels. Consequently, if value functions are written in terms of average and total utility, they should not be increasing in both. For simplicity of presentation, we work with the representative-utility, population-size representation.[42]

All the principles discussed below are members of larger families whose same-number sub-principles are generalized utilitarian. Such principles employ transformed utilities which result from the application of a continuous and increasing function to individual utilities. Value functions for the generalized counterparts of the population principles in this section are presented in the Appendix.

4.1. Classical utilitarianism

The value function for classical utilitarianism (CU) is the sum of utilities. As a consequence, if average utility is constant, increases in population size are good if average utility is positive and bad if average utility is negative.

Classical utilitarianism is illustrated in Figure 1. The dotted lines join points of equal value and we refer to the resulting curves as iso-value curves. Points on iso-value curve I are better than points on II, which are better than points on III, which are better than points on IV. The four curves join average-utility, population-size pairs which are as good as utility vectors in which one person has a utility level of 60, 30, zero and −30 respectively.

Because the addition of an individual with a utility level of zero to a utility-unaffected population does not change total utility, all critical levels are equal to zero. Classical utilitarianism satisfies existence independence (and, therefore, utility and same-number independence), has non-negative critical levels and satisfies priority for lives worth living. As is well known, however, it leads to the repugnant conclusion. The repugnant conclusion is implied because the iso-value curve for any average-utility, population-size pair with positive average utility approaches the population-size axis as numbers increase. This is true of iso-value curves I and II in Figure 1. As a consequence, for any utility vector in which each person experiences the same positive utility level, it is possible to find a larger population size such that an arbitrarily small average-utility level paired with that population size is better.

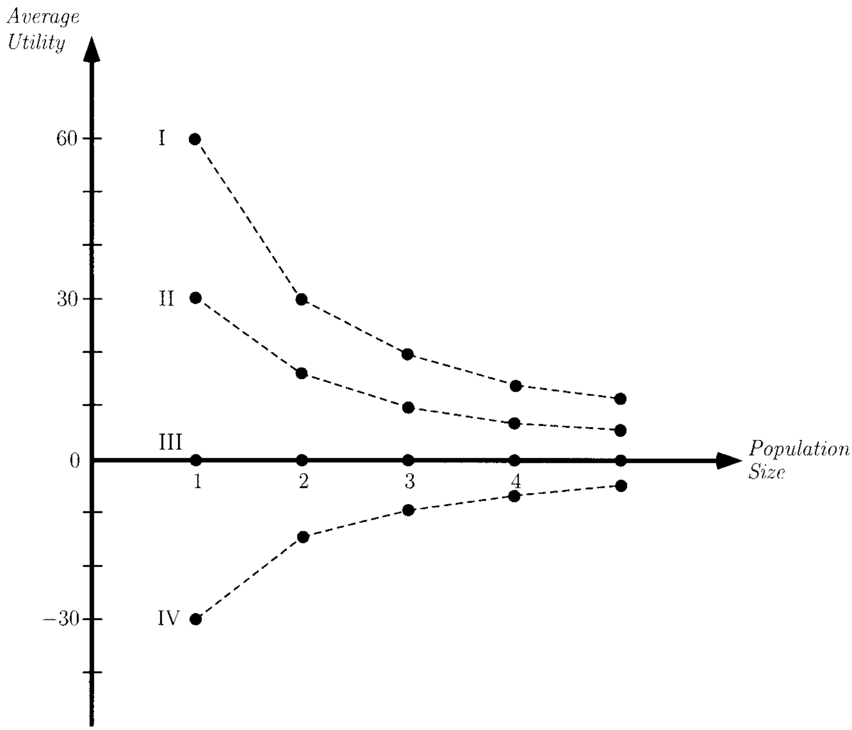

Figure 1 **Classical utilitarianism**

4.2. Critical-level utilitarianism

Critical-level utilitarianism (CLU) is a family of population principles, one for each value of a *fixed* level of utility which is the critical level for every alternative.[43] If the critical level is zero, classical utilitarianism results. The CLU value function can be computed by subtracting the critical level from average utility and multiplying by population size or by subtracting the critical level from the utility level of each person and adding the resulting numbers.

Critical-level utilitarianism with a critical level of 30 is illustrated in Figure 2. The four iso-value curves are constructed in the same way as the iso-value curves of Figure 1. If average utility is constant, increases in population size are good if average utility is above the critical level and bad if average utility is below the critical level. Any alternative with average utility above the critical level is ranked as better than any alternative with average utility below it.

Iso-value curves for average-utility, population-size pairs with average utility above the critical level do not drop below iso-value curve II. If the critical level is positive, therefore, CLU avoids the repugnant conclusion.

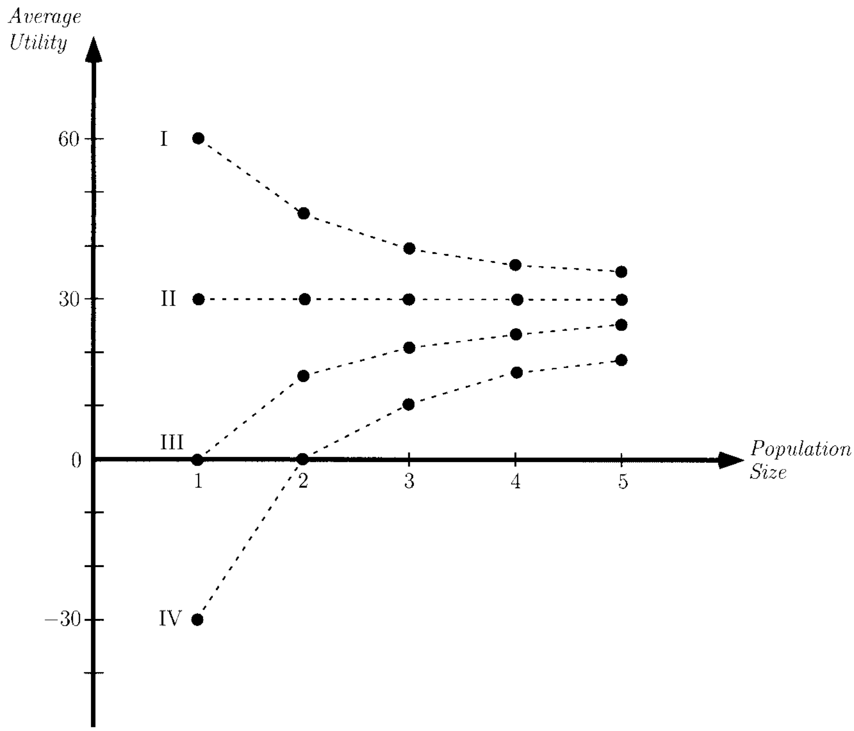

Figure 2 **Critical-level utilitarianism**

If avoidance of the repugnant conclusion is regarded as desirable, a critical-level utilitarian principle with a positive critical level should be chosen. Any such principle violates priority for lives worth living, however. This can be seen in Figure 2 by looking at iso-value curve IV, which crosses the population-size axis from below and stays above it. Any alternative in which one person is alive with a utility level of −30 is ranked as better than any alternative in which each of four people has a utility level of 10. A similar comparison can be found for any alternative in which population size is arbitrary and each person's utility level is negative.

In the example of Table 1, the one-child alternative *x* is better than the two-child alternative *y* if, and only if, the critical level is greater than 10. This is a reflection of a general consequence of a positive critical level: the principle is more conservative than classical utilitarianism about population expansion. In addition, it is easy to check that the ranking of the two alternatives is independent of the existence of Euclid and, therefore, of his utility level. If CLU with a critical level of 20 is applied to the family alone, values are 70 for *x* and 60 for *y*.

Critical-level utilitarianism satisfies existence independence and, therefore, utility and same-number independence. In addition, the critical-level generalized-utilitarian principles are the only ones that satisfy our basic axioms and existence independence.[44] Existence independence implies that critical levels are the same for all alternatives. To see this, suppose that a single individual is added to a utility-unaffected population at the critical level. In that case, the two alternatives are equally good by definition. Because the utility levels of the original population are the same in both alternatives, existence independence requires the ranking of the two alternatives to be independent of both utilities and population size. Consequently, the critical level for the original alternative must be the critical level for all alternatives.

4.3. Restricted critical-level utilitarianism

Critical-level utilitarian principles with positive critical levels can be modified so that all members of the resulting family avoid the repugnant conclusion and satisfy priority for lives worth living. The positive critical level for a CLU principle becomes the critical-level parameter for the corresponding restricted principle. The value function is given by the CLU value function if average utility is greater than the critical-level parameter, by the ratio of average utility and that parameter less one if average utility is positive and no greater than the parameter, and by total utility less one if average utility is non-positive. It is illustrated for a parameter value of 30 in Figure 3 (iso-value curves I to IV are defined as before and iso-value curve V is added). We call the resulting family restricted critical-level utilitarianism (RCLU). It ranks alternatives whose average utilities are greater than 30 using CLU (iso-value curve I), alternatives whose average utilities are positive and no greater than 30 using average utilitarianism (iso-value curves II and V), and alternatives whose average utilities are non-positive with classical utilitarianism (iso-value curves III and IV). In addition, alternatives in the first group are ranked as better than those in the second which, in turn, are ranked as better than those in the third.

Suppose average utility is constant. If it is above the critical-level parameter, population increases are good; if it is non-negative and no greater than the parameter, population increases are neither good nor bad; and if it is negative, population increases are bad.

Critical levels are equal to the critical-level parameter for alternatives whose average utility is above it, average utility for alternatives whose average utility is positive and no greater than the parameter, and zero for alternatives whose average utility is non-positive. Consequently, all critical levels are non-negative.

Because the iso-value curves for average-utility, population-size pairs with average utility above the parameter do not approach the population-size axis (iso-value curves such as I are bounded below by II), the repugnant conclusion is avoided. In addition, because the iso-value curves for average-utility, population-size pairs with negative average utilities, such as IV, do not cross the population-size axis the priority for lives worth living is satisfied.

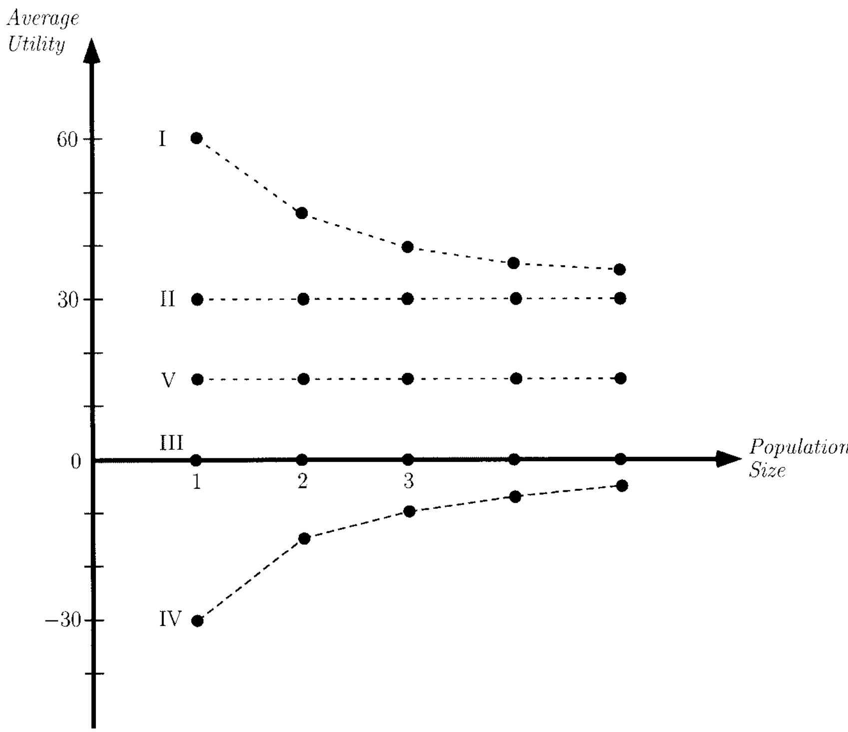

Figure 3 **Restricted critical-level utilitarianism**

These principles satisfy neither utility nor existence independence. Consider, again, the example of the disabled child summarized in Table 1 and suppose that RCLU with a critical-level parameter equal to 20 is used to rank *x* and *y*. If Euclid's utility level is 100, average utility is 70 in *x* and 55 in *y*, which are both greater than 20. Consequently, values are 150 for *x* and 140 for *y*, and the alternative with one child is better. Now suppose that Euclid's utility level is −140. Then average utilities are −10 for *x* and −5 for *y* and values are −31 for *x* and −21 for *y*, and the two-child alternative is better. Consequently, utility independence is not satisfied and, because existence independence implies utility independence, it is also not satisfied.

4.4. Number-sensitive, critical-level utilitarianism

The number-sensitive, critical-level utilitarian (NCLU) family of principles allows critical levels to depend on population size but not on utility levels, and includes the critical-level utilitarian family as a special case. We write the critical level for population size *n* as c_n. If the null alternative is included, its critical

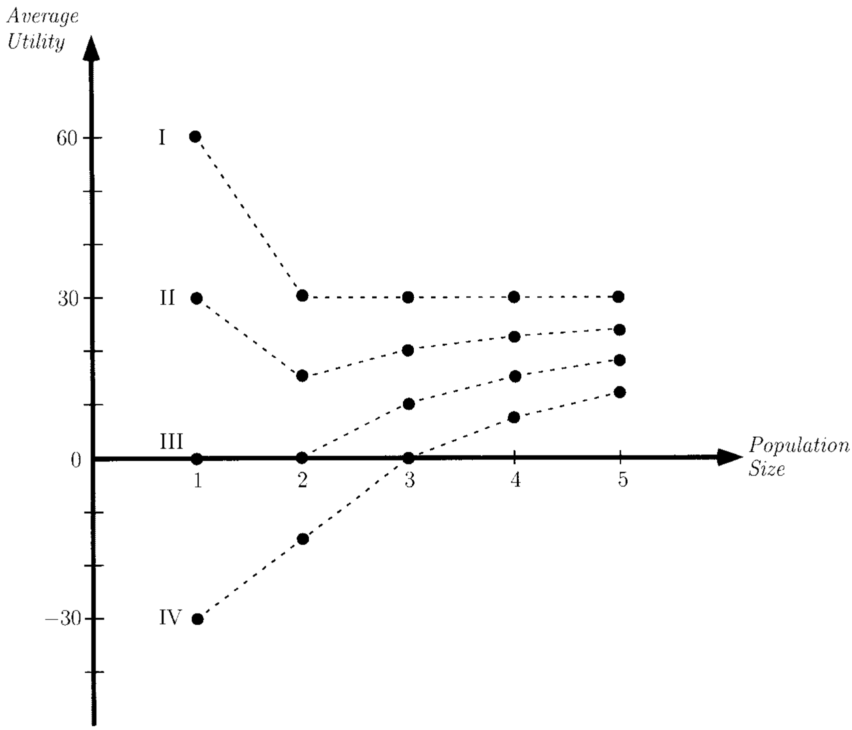

Figure 4 **Number-sensitive, critical-level utilitarianism**

level is c_0 and, if not, c_0 can be chosen arbitrarily (it makes no difference to the rankings). The NCLU value function for an alternative with population size n can be found by subtracting the average of $c_0, \ldots, c_{n-1}$ from average utility and multiplying by the population size. Equivalently, it can be found by adding the first utility level less c_0, the second less c_1, and so on.

Number-sensitive critical-level utilitarianism is illustrated in Figure 4. Critical levels are zero for population size one and 30 for population sizes greater than one. If average utility is constant, population increases are bad for average-utility, population-size pairs that lie below iso-value curve III (and have negative average utility), but, for pairs that are above curve III, the goodness or badness of population-size increases depends on how large the population is. In the figure, NCLU coincides with CU for population sizes one and two and, as population size becomes large, it approximates CLU with a critical level of 30.

If critical levels are all non-negative, the repugnant conclusion is avoided if, and only if, there is a sequence of population sizes such that the sequence of corresponding critical levels does not approach zero.[45]

If critical levels are not all the same, utility independence is satisfied, but existence independence is not. To see that, consider the disabled-child example summarized in Table 1 and suppose that critical levels are equal to zero for population sizes one to three and 90 for population sizes above three. Without loss of generality, c_0 may be chosen to be zero. Writing Euclid's utility level as t, values are $110 + t$ for x and $120 + t$ for y, and y is better than x for all values of t. Suppose, now, that we discover that Euclid had an identical twin brother whose lifetime utility level was also equal to t. In that case, values are $110 + 2t$ for x and $30 + 2t$ for y, so x is better. Although the ranking of the two alternatives is independent of the utilities of the Euclids, it is not independent of their number.

The number-sensitive critical-level generalized-utilitarian principles are the only ones that satisfy our basic axioms and utility independence.[46] Avoidance of the repugnant conclusion requires the same conditions on critical levels as those for NCLU.

All members of the number-sensitive critical-level utilitarian family that avoid the repugnant conclusion fail to satisfy the priority for lives worth living. This can be seen in Figure 4 by noting that iso-value curve IV crosses the population-size axis. We show, in the following subsection, that it is possible to modify these principles so that they avoid the repugnant conclusion and satisfy priority for lives worth living, but, in that case, neither existence nor utility independence are satisfied.

4.5. Restricted number-sensitive, critical-level utilitarianism

The restricted number-sensitive critical-level utilitarian (RNCLU) family of principles is a modification of the number-sensitive subfamily with non-negative, non-decreasing critical levels and at least one positive critical level. It uses the critical levels for NCLU as parameters and we write $\bar{c}_n$ as the average of $c_0, \ldots, c_{n-1}$, where c_0 is non-negative. The value function is equal to the value of the corresponding NCLU value function if average utility is greater than $\bar{c}_n$, equal to average utility divided by $\bar{c}_n$ less one if average utility is positive and no greater than $\bar{c}_n$, and equal to total utility less one if average utility is non-positive. RNCLU with $c_0 = c_1 = c_2 = 0$ and $c_n = 30$ for all population sizes greater than two is illustrated in Figure 5. Note that there are no average-utility, population-size pairs on iso-value curve V for population sizes one and two. This occurs when some of the critical-level parameters are zero and some are positive, and does not occur when all are positive. All average-utility, population-size pairs with an average utility greater than $\bar{c}_n$ are better than all pairs with a positive average utility that is no greater than $\bar{c}_n$ and these are, in turn, better than all pairs in which average utility is non-positive. RNCLU ranks alternatives with average utilities above $\bar{c}_n$ with the corresponding NCLU principle and alternatives with non-positive average utilities with CU.

The critical level for an alternative with population size n is c_n if average utility is greater than $\bar{c}_n$, positive and no greater than $\bar{c}_n$ if average utility is pos-

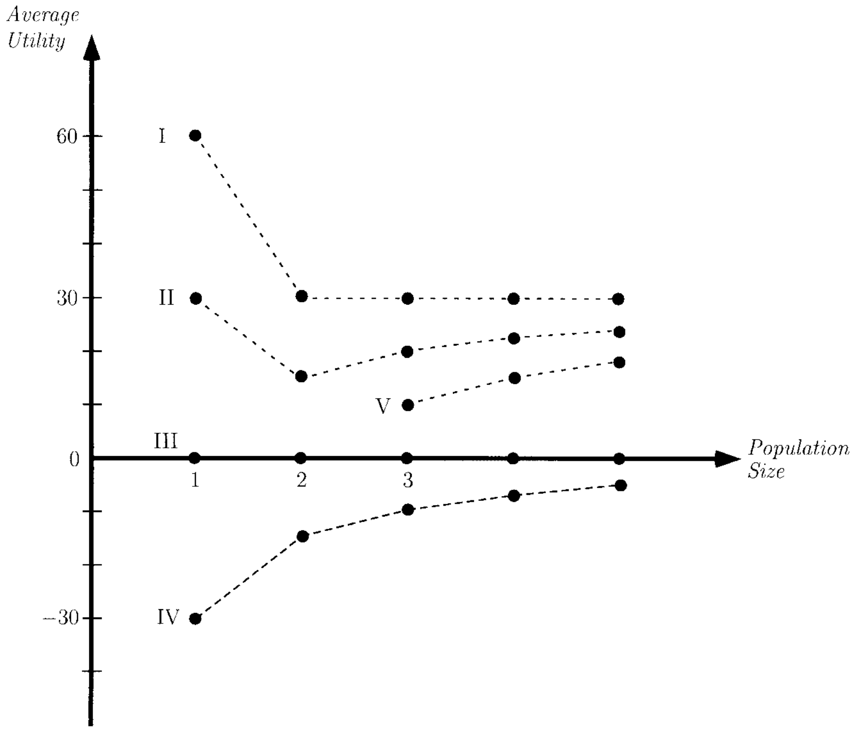

Figure 5 **Restricted number-sensitive critical-level utilitarianism**

itive and no greater than $\bar{c}_n$, and zero if average utility is non-positive. Consequently, all of the RNCLU principles satisfy non-negative critical levels.

The repugnant conclusion is avoided if, and only if, the corresponding NCLU principle avoids it. This is guaranteed because the critical-level parameters are non-decreasing and at least one is positive. All of the restricted number-sensitive, critical-level utilitarian principles satisfy priority for lives worth living because all alternatives with a positive average utility are ranked as better than all those with a negative average utility. An illustration is provided by iso-value curve IV in Figure 5: it does not cross the population-size axis.

Restricted number-sensitive principles satisfy neither existence nor utility independence. An example is provided by the one discussed in connection with the RCLU family because those principles also belong to the RNCLU family.

4.6. Average utilitarianism

Average utilitarianism (AU) ranks alternatives with a value function which is equal to average utility. It is illustrated in Figure 6. The flat iso-value curves indi-

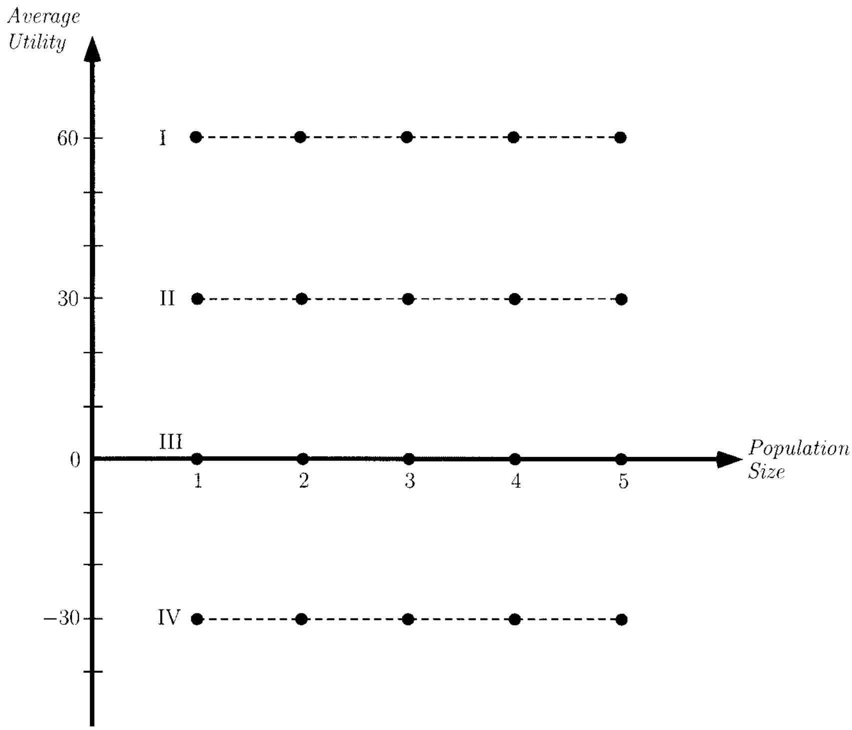

Figure 6 **Average utilitarianism**

cate that, if average utility is constant, the principle is indifferent to changes in population size. As a consequence, the principle makes some stark trade-offs: an alternative with a population of any size in which each person is equally well off is ranked as worse than an alternative in which a single person enjoys a trivially higher utility level.

Because the addition of a person whose utility level is equal to the average utility of an unaffected population does not change average utility, the critical level for any alternative is average utility. Consequently, critical levels for alternatives with negative average utilities are negative and the axiom non-negative critical levels is not satisfied.

Iso-value curves for alternatives in which average utility is positive do not approach the population-size axis, and this means that the repugnant conclusion is avoided. In addition, all alternatives with positive average utility are ranked as better than all those with negative average utility and, as a consequence, priority for lives worth living is satisfied. The discussion of average utilitarianism following the disabled-child example summarized in Table 1 demonstrates that AU satisfies neither utility nor existence independence.

4.7. Number-dampened utilitarianism

The number-dampened utilitarian (NDU) family[47] has both classical and average utilitarianism as members. Its value function is equal to average utility multiplied by a positive-valued function of population size. If the function is equal to population size or any positive multiple, the principle is CU and, if the function is equal to any positive constant, AU results.[48]

It is possible for an NDU principle to approximate CU for 'small' population sizes and AU for 'large' ones, a property originally suggested by Hurka.[49] Such a case is illustrated in Figure 7. For that principle, the function takes on the values 1.0, 2.0, 2.6 and 3.0 for population sizes one, two, three and four or more, respectively. For population sizes one and two, the value function coincides with that of CU and alternatives with population sizes greater than three are ranked by AU.

It is shown in the Appendix that critical levels for NDU are equal to a multiple of average utility and that the multiple can depend on population size. In the example of Figure 7, the ratios of critical levels to average utility are 0.00, 0.31, 0.47 and 1.00 for population sizes one, two, three and four or more.

A subfamily specializes NDU in a way that is parallel to the way that constant critical levels specialize number-sensitive critical-level utilitarianism. In that subfamily, the ratio of critical levels to average utilities is a positive constant between zero and one. A second subfamily also uses a positive constant k less than one and makes the function that multiplies average utility equal to the sum $1 + k + \ldots + k^{n-1}$.[50] Because critical levels for the NDU principles are equal to a multiple of average utility, they have some negative critical levels unless the ratio is equal to zero. In that case, however, the principle is CU: all other members of the family have some negative critical levels.

Some members of the NDU family, such as CU, imply the repugnant conclusion and others, such as AU, do not. The repugnant conclusion is avoided if, and only if, the multiplying function does not increase without limit as population size increases. All NDU principles for which the ratio of critical levels to average utilities is a positive constant between zero and one lead to the repugnant conclusion. If the ratio is non-constant, the requirement that it be non-decreasing is consistent with Hurka's suggestion and Carlson's intuition (discussed in subsection 3.1 above). In order to avoid the repugnant conclusion, any NDU principle for which the ratio of critical levels to average utilities is non-decreasing and between zero and one must approximate average utilitarianism as population size becomes large. That is true when the multiplying function is $1 + k + \ldots + k^{n-1}$. In addition, the sum itself approaches the finite number $1/(1 - k)$ as population size becomes large and the repugnant conclusion is avoided.

Because every NDU principle ranks all alternatives with positive average

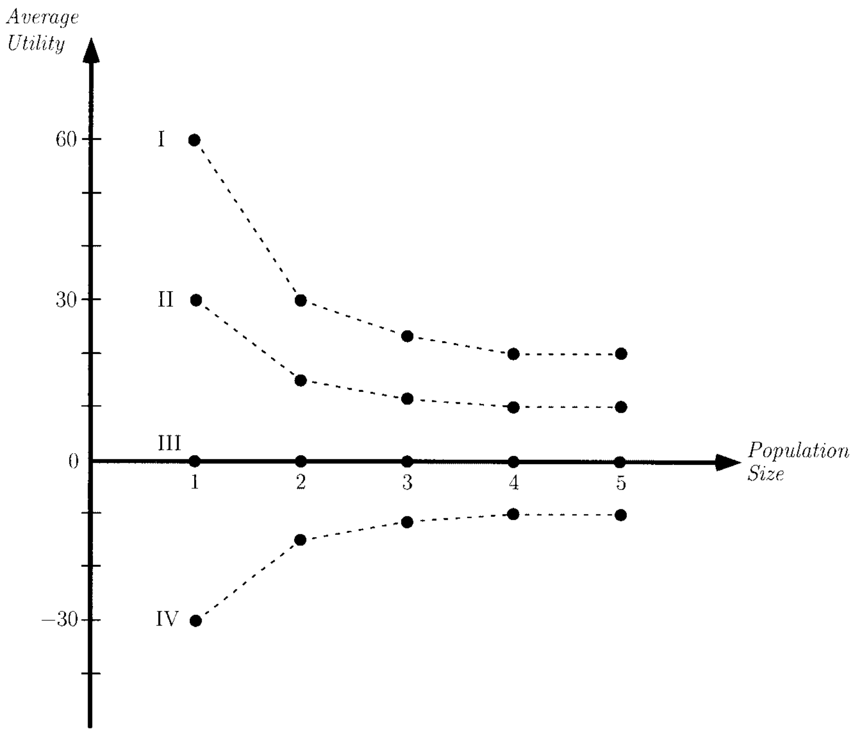

Figure 7 **Number-dampened utilitarianism**

utilities as better than all alternatives with negative average utilities, all of the NDU principles satisfy priority for lives worth living. Every member of the number-dampened utilitarian family satisfies same-number independence, but none of them, other than classical utilitarianism, satisfies either utility or existence independence.

4.8. Restricted number-dampened utilitarianism

Suggested by Hurka,[51] the restricted number-dampened utilitarian (RNDU) family of principles provides a partial solution to one of the most important defects of the number-dampened family, namely, that all those principles, other than CU, have negative critical levels.

The value function for the restricted principles coincides with the NDU value function when average utility is positive and with the CU value function when average utility is non-positive. The restricted version of the example of Figure 7 is illustrated in Figure 8. Above the population-size axis, the iso-value curves are the same for both principles. But below the population-size axis, iso-value curves

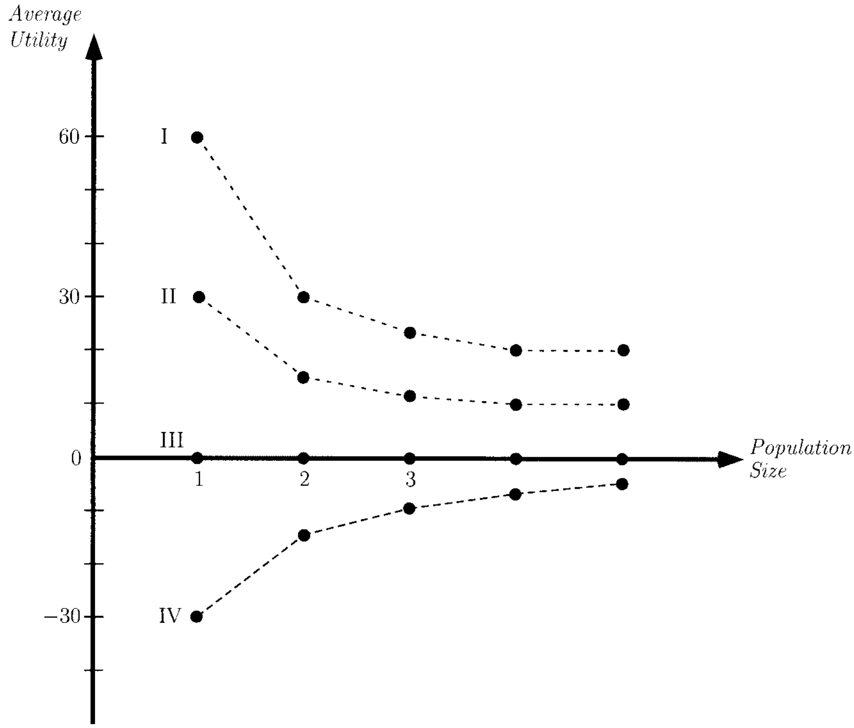

Figure 8 **Restricted number-dampened utilitarianism**

for the restricted principle approach the population-size axis for large population sizes, reflecting the fact that the value function coincides with the CU value function for negative average utilities.

Because of this, critical levels for alternatives with non-positive average utilities are zero and, hence, non-negative. Critical levels for alternatives with positive average utilities are not necessarily non-negative, however. Suppose that the function that multiplies average utility takes on the values one and four for population sizes one and two and consider an alternative in which a single person has a utility level of four. Then, the critical level is minus two. For any NDU principle, the ratios of critical levels to average utilities are non-negative if, and only if, the ratio of the multiplying function to population size does not increase as population size increases. This condition, applied to RNDU principles, is necessary and sufficient for non-negative critical levels.

Several special cases of RNDU have, however, non-negative critical levels for all alternatives. The first of these is restricted average utilitarianism (RAU). Its value function is equal to average utility when average utility is positive and total utility when average utility is non-positive. Its critical levels are equal to average utility for alternatives with positive average utility and zero for alternatives whose average utility is non-positive. The restricted version of the NDU family for which the ratio of critical levels to average utilities is a constant between zero and one also has non-negative critical levels: they are equal to the constant multiplied by average utility when average utility is positive and zero when average utility is non-positive. In addition, when the multiplying function is $1 + k + \ldots + k^{n-1}$, critical levels for the restricted principle are positive for all alternatives with positive average utility and zero otherwise.

The repugnant conclusion refers to alternatives with positive average utilities only. Because value functions for the restricted number-dampened utilitarian principles coincide with the value functions for their unrestricted counterparts when average utility is positive, the conditions for avoidance of the repugnant conclusion are the same for the restricted and unrestricted families.

As is the case for the NDU principles, all members of the RNDU family rank all alternatives with positive average utilities as better than all those with negative average utilities. Consequently, they all satisfy the priority for lives worth living. Of all the RNDU principles, only classical utilitarianism satisfies utility independence. This follows from the discussion of independence and the NDU principles.

5. Concluding remarks

Table 3 summarizes the results of the previous section. Because all of the principles considered satisfy same-number independence, that axiom is omitted. The table is divided into two parts: the first lists families of principles that can be regarded as generalizations of CU and the second lists families that can be thought of as generalizations of AU. The first group consists of all the unrestricted and restricted critical-level families and the second consists of all the unrestricted and restricted number-dampened families. In the second group, we have included restricted average utilitarianism which is a member of the restricted number-dampened family and three cases of number-dampened principles and their restricted counterparts. The first group of number-dampened principles consists of the whole of the NDU family; the second selects members of the family in which the ratio of critical levels to average utilities is equal to a constant between zero and one; and the third selects NDU principles for which the ratio of critical levels to average utilities is positive, non-decreasing and approaches one as population size becomes large. The last group approximates AU as population size increases. There is no principle in the list that satisfies all our axioms and it is worth asking whether this is a general result. The answer is 'yes' and the impossibility theorem can be found in Blackorby, Bossert and Donaldson.[52] We state a slightly weaker result in the Appendix (Theorem 3). There is no population principle that satisfies our basic axioms, utility independence, avoidance of the

Table 3

	Utility independence	Existence independence	Non-negative critical levels	Avoidance of the repugnant conclusion	Priority for lives worth living
CU[1]	•	•	•		•
CLU[2]	•	•	•	•	
RCLU[3]			•	•	•
NCLU[4]	•		•	•	
RNCLU[5]			•	•	•
AU[6]				•	•
RAU[7]			•	•	•
NDU[8]					•
RNDU[9]					•
NDU[10]					•
RNDU[11]			•		•
NDU[12]				•	•
RNDU[13]			•	•	•

1. Classical Utilitarianism
2. Critical-Level Utilitarianism:positive critical level
3. Restricted Critical-Level Utilitarianism:positive critical-level parameter
4. Number-Sensitive, Critical-Level Utilitarianism:non-negative,non-decreasing critical levels; some positive
5. Restricted Number-Sensitive, Critical-Level Utilitarianism:restricted version of 4
6. Average Utilitarianism
7. Restricted Average Utilitarianism
8. Number-Dampened Utilitarianism:general case
9. Restricted Number-Dampened Utilitarianism:general case
10. Number-Dampened Utilitarianism:ratio of critical level to average utility is a positive constant less than one
11. Restricted Number-Dampened Utilitarianism:restricted version of 10
12. Number-Dampened Utilitarianism:ratio of critical level to average utility is positive,non-decreasing and approaches one as numbers increase
13. Restricted Number-Dampened Utilitarianism:restricted version of 12

repugnant conclusion and priority for lives worth living. Because existence independence implies utility independence and requires constant critical levels, it is also true that there is no population principle that satisfies the basic axioms, existence independence, avoidance of the repugnant conclusion and priority for lives worth living. In comparing principles, therefore, we should bear in mind that no principle can be completely satisfactory.

We regard negative critical levels as an unacceptable property of population principles and use it to eliminate some of the subfamilies in Table 3. None of the families in the first group are eliminated, but all of the unrestricted families in the second are. In addition, the general case of restricted NDU principles is eliminated for the same reason. As a consequence, we focus on restricted average utilitarianism and the second and third restricted number-dampened families.

Many investigators, among them Heyd and Parfit,[53] make a strong case for avoidance of the repugnant conclusion. If we accept that view, we may eliminate classical utilitarianism and all but the last of the restricted versions of number-dampened utilitarianism. The families that remain fall into two groups: those that satisfy utility or existence independence, but violate priority for lives worth living, and those that satisfy neither utility nor existence independence, but satisfy priority for lives worth living. We consider the two groups in turn.

The number-sensitive critical-level utilitarian family, which includes the critical-level utilitarian family, is the only one that can satisfy utility independence and all axioms other than the priority for lives worth living. Of the members of this family that do not imply the repugnant conclusion, the most attractive are those whose critical levels do not decrease as population size increases.[54] In that case, avoidance of the repugnant conclusion requires critical levels to be positive after some population size is reached. If there is more than one critical level, however, existence independence is not satisfied: it requires constant critical levels. Therefore, if existence independence is thought to be desirable, members of the critical-level family with positive critical levels are the only satisfactory principles.

The families in the second group avoid the repugnant conclusion, satisfy priority for lives worth living, but necessarily violate both utility and existence independence. Subfamilies in Table 3 that have the requisite characteristics are restricted critical-level utilitarianism, restricted number-dependent critical-level utilitarianism, restricted average utilitarianism and restricted number-dampened utilitarianism in which the ratio of critical levels to average utility is positive, non-decreasing and approaches one as numbers increase. Of these, restricted average utilitarianism retains the stark trade-offs of AU for alternatives with positive average utilities.

The third subfamily of the restricted NDU principles has its own problems, however. All members of this subfamily, other than RAU, must have ratios of critical levels to average utilities that are different for some population sizes. As a consequence, some moral significance must be attached to certain absolute numbers. If the principle approximates CU at small population sizes and AU at large ones, some numerical meaning for 'small' and 'large' must be found so that the 'speed' of the transition between the two limiting cases can be chosen. It is, however, very difficult to imagine how this can be done without reference to the carrying capacity of the universe. If that occurs, then value becomes confounded with constraints. The same consideration applies to the number-sensitive critical-level utilitarian principles.

Both the restricted and unrestricted critical-level principles require the choice of a single parameter that must be positive if the repugnant conclusion is to be avoided. In both cases, that parameter places a 'floor' on the trade-off between average utility and numbers: alternatives with average utility above the parameter are better than all others. Feldman[55] defends critical-level utilitarianism with the critical level set equal to 'some modest level of happiness that people deserve merely in virtue of being people' and adds that if a person has a positive utility level that is below the critical level, 'it does not make the world better' and 'may make the world worse'.[56] These comments are consistent with both the unrestricted and restricted critical-level principles. Feldman allows for different levels of desert for different people and they can be based on non-welfare characteristics such as industriousness. Even with multiple critical levels, however, the resulting principles rank same-population alternatives with utilitarianism: the critical levels play no role. In addition, such principles prefer the *ceteris paribus* addition of a person with a low critical level to a person with a high critical level, provided that they are expected to have the same lifetime utility level. Consequently, it may be better to regard the critical level as an ethical parameter that is independent of desert.

It is true, of course, that we do not have to choose principles whose same-number sub-principles are utilitarian. All of the principles discussed above are members of larger families whose same-number sub-principles are generalized utilitarian. The axioms we have employed do not rule out giving priority to the interests of low-utility individuals and the resulting inequality aversion that generalized-utilitarian families can represent. Carlson's combined principle[57] is unsatisfactory because it is not continuous and satisfies neither weak nor strong Pareto, but it is possible to find an acceptable principle that accords with his intuitions. His principle gives priority to the well-being of individuals who are below neutrality. This can be accomplished, for example, with a generalized-utilitarian same-number sub-principle where transformed utility is equal to utility itself for non-negative utilities and equal to utility times two for negative utilities. Although Carlson rejects existence independence, the example he discusses (see subsection 3.1) is consistent with utility independence. If utility independence is accepted, then a number-sensitive, critical-level, generalized-utilitarian principle would correspond to Carlson's idea. To be consistent with his intuition, critical levels would have to be non-decreasing and non-constant. If utility independence is rejected, a restricted number-sensitive, generalized-utilitarian principle or a restricted number-dampened, generalized principle with a suitable choice of the multiplying function would be appropriate.[58]

Appendix

Let $\mathcal{Z}_{++}$ be the set of positive integers and let $\mathcal{R}$ be the set of real numbers. For $n \in \mathcal{Z}_{++}$, $\mathcal{Z}^n_{++}$ is the n-fold Cartesian product of $\mathcal{Z}_{++}$ and $\mathcal{R}^n$ is the n-fold Cartesian product of $\mathcal{R}$. The positive (non-negative; negative) orthant of $\mathcal{R}^n$ is denoted by $\mathcal{R}^n_{++}$ ($\mathcal{R}^n_{+}$; $\mathcal{R}^n_{--}$). For $n \in \mathcal{Z}_{++}$, 1_n is the vector consisting of n ones. Individuals are indexed by the positive integers. Therefore, there is a countable number of potential individuals; however, all actual populations considered here are finite. For all $n \in \mathcal{Z}_{++}$, let $\mathbf{Z}^n \subset \mathcal{Z}^n_{++}$ be the set of all $\pi \in \mathcal{Z}^n_{++}$ such that $\pi_i \neq \pi_j$ for all distinct $i, j \in \{1, \ldots, n\}$. We use the notation $\Omega = \cup_{n \in \mathcal{Z}_{++}} \mathcal{R}^n$. Letting $u_\emptyset$ denote the utility vector associated with the alternative in which no one is alive, we define $\Omega_\emptyset = \Omega \cup \{u_\emptyset\}$. In addition, $\Omega_{++} = \cup_{n \in \mathcal{Z}_{++}} \mathcal{R}^n_{++}$ and $\Omega_{--} = {}_{n \in \mathcal{Z}_{++}} \mathcal{R}^n_{--}$.

For each alternative, the associated population size is denoted by $n \in \mathcal{Z}_{++}$ and the associated vector of individual identities by $\pi = (\pi_1, \ldots, \pi_n) \in \mathbf{Z}_n$. Utilities for those alive are $u = (u_1, \ldots, u_n) \in \mathcal{R}^n$ where, for all $i \in \{1, \ldots, n\}$, u_i is individual π_i's lifetime utility in the alternative in question. Thus, $(\pi, u) = ((1001, 4, 27), (23, 12, -6))$ describes an alternative in which individuals 1001, 4 and 27 are alive and have lifetime-utility levels of 23, 12 and −6 respectively.

A welfarist population principle is represented by an ordering $\overset{*}{R}$ on $\cup_{n \in \mathcal{Z}_{++}} (\mathbf{Z}^n \times \mathcal{R}^n)$ where, for all $(\pi, u), (\rho, v) \in \cup_{n \in \mathcal{Z}_{++}} (\mathbf{Z}^n \times \mathcal{R}^n)$, $(\pi, u)\overset{*}{R}(\rho, v)$ means that (π, u) is at least as good as (ρ, v). The better-than relation and the as-good-as relation corresponding to $\overset{*}{R}$ are denoted by $\overset{*}{P}$ and $\overset{*}{I}$ respectively. Because $\overset{*}{R}$ ranks named utility vectors, for all $n \in \mathcal{Z}_{++}$ and for all $(\pi, u), (\rho, v) \in \mathbf{Z}^n \times \mathcal{R}^n$, if ρ is a permutation of π and v is the same permutation of u, then $(\pi, u)\overset{*}{I}(\rho, v)$.

Our basic axioms are strong Pareto, continuity, anonymity and existence of critical levels. These conditions are defined as follows.

Strong Pareto

For all $n \in \mathcal{Z}_{++}$, for all $\pi \in \mathbf{Z}^n$ and for all $u, v \in \mathcal{R}^n$, if $u_i \geq v_i$ for all $i \in \{1, \ldots, n\}$ with at least one strict inequality, then $(\pi, u)\overset{*}{P}(\pi, v)$.

Continuity

For all $n \in \mathcal{Z}_{++}$, for all $\pi \in \mathbf{Z}^n$ and for $u \in \mathcal{R}^n$, the sets $\{v \in \mathcal{R}^n \mid (\pi, v)\overset{*}{R}(\pi, u)\}$ and $\{v \in \mathcal{R}^n \mid (\pi, u)\overset{*}{R}(\pi, v)\}$ are closed.

Anonymity

For all $n \in \mathcal{Z}_{++}$, for all $\pi, \rho \in \mathbf{Z}^n$ and for all $u \in \mathcal{R}^n$, $(\pi, u)\overset{*}{I}(\rho, u)$.

Existence of critical levels

For all $n \in \mathcal{Z}_{++}$ and for all $(\pi, u) \in \mathbf{Z}^n \times \mathcal{R}^n$, there exist $j \in \mathcal{Z}_{++} \setminus \{\pi_1, \ldots, \pi_n\}$ and $c \in \mathcal{R}$ such that $((\pi, j), (u, c))\overset{*}{I}(\pi, u)$.

If $\overset{*}{R}$ satisfies anonymity, it is isomorphic to an ordering R with better-than and as-good-as relations P and I. This is the case because individual identities are irrelevant for anonymous social evaluation. That is, for all $(\pi, u), (\rho, v) \in \cup_{n \in \mathcal{Z}_{++}} (\mathbf{Z}^n \times \mathcal{R}^n)$, $(\pi, u)\overset{*}{R}(\rho, v)$ if, and only if, uRv. Furthermore, the same-number restrictions of R must be anonymous: if v is a permutation of u, uIv. Strong Pareto and continuity imply that these restrictions are strictly monotonic and continuous. From now on, we use R instead of $\overset{*}{R}$ for simplicity.

Same-number independence requires that, for a fixed population size, the relative ranking of any two utility vectors is independent of the utilities of those individuals whose utility levels are the same in both.

Same-number independence

For all $n, m \in \mathcal{Z}_{++}$, for all $u, v \in \mathcal{R}^n$ and for all $w, s \in R^m$:

$$(u, w)R(v, w) \Leftrightarrow (u, s)R(v, s). \tag{1}$$

A stronger axiom extends same-number independence to different-number comparisons and requires the social ranking to be independent of the utilities of unconcerned individuals, but not necessarily of their existence.

Utility independence

For all $u, v \in \Omega_{\emptyset}$, for all $r \in \mathcal{Z}_{++}$ and for all $w, s \in \mathcal{R}^r$:

$$(u, w)R(v, w) \Leftrightarrow (u, s)R(v, s). \tag{2}$$

An extended version of the axiom also applies to different-number comparisons. It requires the social ranking to be independent of the existence of the unconcerned. Thus, the ranking is independent of their utilities and their number.

Existence independence

For all $u, v, w \in \Omega$:

$$(u, w)R(v, w) \Leftrightarrow uRv. \tag{3}$$

That same-number independence is implied by utility independence follows from the definitions. It is also true that existence independence implies utility independence. Let u, v, w, and s be defined as above. Then, using existence independence twice:

$$(u, w)R(v, w) \Leftrightarrow (u, w, s)R(v, w, s) \Leftrightarrow (u, s)R(v, s). \tag{4}$$

The requirement that critical levels be non-negative is defined as follows.

Non-negative critical levels

For all $u \in \Omega$ and for all $c \in \mathcal{R}$, if $(u, c)Iu$, then $c \geq 0$.

A population principle implies the repugnant conclusion if, and only if, for any population size $n \in \mathcal{Z}_{++}$, any positive utility level ξ and any utility level $\epsilon \in (0, \xi)$, there exists a population size $m > n$ such that an m-person alternative in which every individual experiences utility level ϵ is ranked as better than an n-person society in which every individual's utility level is ξ. The axiom that requires the repugnant conclusion to be avoided is defined as follows.

Avoidance of the repugnant conclusion

There exist $n \in \mathcal{Z}_{++}$, $\xi \in \mathcal{R}_{++}$ and $\epsilon \in (0, \xi)$ such that, for all $m > n$, $\xi 1_n R \epsilon 1_m$.

A population principle implies the sadistic conclusion[59] if, and only if, when adding people to a utility-unaffected population, it can be better to add people with negative utilities rather than a possibly different number of people with positive welfare.

Avoidance of the sadistic conclusion

For all $u \in \Omega$, for all $v \in \Omega_{++}$ and for all $w \in \Omega_{--}$, $(u, v)R(u, w)$.

A population principle implies the strong sadistic conclusion[60] if, and only if, any alternative in which each person experiences a negative utility level is ranked as better than some alternative with a population in which each member experiences a positive utility level. The strong sadistic conclusion is avoided by all principles that satisfy the priority for lives worth living.

Priority for lives worth living

For all $u \in \Omega_{++}$ and for all $v \in \Omega_{--}$, uPv.

For population size n and an n-person utility vector u, we write average utility as:

$$\mu = \mu^n(u) = \frac{1}{n}\sum_{i=1}^{n} u_i. \tag{5}$$

According to same-number utilitarianism, two utility vectors of the same population size are ranked by their average or total utilities. For the population principles considered here (which have same-number utilitarian sub-principles), the value function V can be written as a function W of population size n and average utility μ. That is, utility vector u is at least as good as utility vector v if, and only if, the value of W calculated at the population size and average utility of u is greater than or equal to the value of W calculated at the population-size, average-utility pair corresponding to v. Formally, for all $n, m \in \mathcal{Z}_{++}$, for all $u \in \mathcal{R}^n$ and for all $v \in R^m$:

$$uRv \Leftrightarrow V(u) \geq V(v) \Leftrightarrow W(n, \mu^n(u)) \geq W(m, \mu^m(v)). \tag{6}$$

We now present value functions for the population principles we investigate.

Classical utilitarianism uses total utility as the criterion to rank utility vectors, and we obtain the value function:

$$V_{CU}(u) = W_{CU}(n, \mu) = n\mu = \sum_{i=1}^{n} u_i. \tag{7}$$

For classical utilitarianism, all critical levels are equal to zero, the utility level that represents neutrality.

Critical-level utilitarianism is a family of principles, one for each value of a fixed utility level which is the critical level for every alternative. The CLU value functions are given by:

$$V_{CLU}(u) = W_{CLU}(n, \mu) = n(\mu - \alpha) = \sum_{i=1}^{n} (u_i - \alpha) \tag{8}$$

where $\alpha \in \mathcal{R}$ is the critical level for the particular principle represented by the function. If the critical level α is zero, classical utilitarianism results and, therefore, CU is a member of the CLU family. A CLU principle avoids the repugnant conclusion if, and only if, α is positive.

The restricted critical-level utilitarian family of principles is introduced in this article.

Its value functions are given by:

$$V_{RCLU}(u) = W_{RCLU}(n, \mu) = \begin{cases} n(\mu - \alpha) = \sum_{i=1}^{n}(u_i - \alpha) & \text{if } \mu > \alpha, \\ \mu/\alpha - 1 = \sum_{i=1}^{n} u_i/(n\alpha) - 1 & \text{if } 0 < \mu \leq \alpha, \\ n\mu - 1 = \sum_{i=1}^{n} u_i - 1 & \text{if } \mu \leq 0, \end{cases} \tag{9}$$

where α is positive. This value function is equal to the value function for CLU for all average utilities that are greater than α, equal to the percentage shortfall of average utility from α when average utility is positive and less than or equal to α, and equal to total utility less one when average utility is non-positive. Consequently, all alternatives whose average utility is above the critical-level parameter are better than all whose average utility is positive and not greater than it, and these alternatives are, in turn, better than all whose average utilities are non-positive. Critical levels are equal to α for all alternatives in the first set, average utility for those in the second, and zero for those in the third. All RCLU principles avoid the repugnant conclusion and satisfy the priority for lives worth living.

The number-sensitive critical-level utilitarian family is a generalization of the CLU family. Its critical levels are independent of utility levels, but not necessarily independent of population size. We write the critical level for population size n as c_n. Because the null alternative is not considered, c_0 is an arbitrary real number (the number chosen makes no difference to rankings in this case). If the null alternative were included, c_0 would be its critical level. The value functions for the NCLU principles can be written as:

$$V_{NCLU}(u) = W_{NCLU}(n, \mu) = n(\mu - \bar{c}_n) = \sum_{i=1}^{n}(u_i - c_{i-1}) \tag{10}$$

where

$$\bar{c}_n = \frac{1}{n}\sum_{i=1}^{n} c_{i-1}. \tag{11}$$

The average of c_0 and the critical levels for population sizes 1 to $n - 1$ is $\bar{c}_n$. CLU results from making c_0 and all the critical levels equal to the same real number, so that $\bar{c}_n$ is equal to α, the fixed critical level.

Members of the restricted number-sensitive critical-level utilitarian family, the second family introduced in this article, are represented by value functions that are generalizations of those for RCLU and are derived from those for NCLU. Formula

(11) defines $\bar{c}_n$ and we assume that $c_0 \geq 0$, that the c_n are non-decreasing, and that at least one c_n is positive. The value functions can be written as:

$$V_{RNCLU}(u) = W_{RNCLU}(n, \mu) = \begin{cases} n(\mu - \bar{c}_n) = \sum_{i=1}^{n}(u_i - c_{i-1}) & \text{if } \mu > \bar{c}_n, \\ \mu/\bar{c}_n - 1 = \sum_{i=1}^{n} u_i / \sum_{i=1}^{n} c_{i-1} - 1 & \text{if } 0 < \mu \leq \bar{c}_n, \\ n\mu - 1 = \sum_{i=1}^{n} u_i - 1 & \text{if } \mu \leq 0. \end{cases} \quad (12)$$

It is possible to have $\bar{c}_n = 0$ for some n and, in that case, the middle branch of (12) does not apply. Although the ordering of population-size, average-utility pairs represented by the NCLU value functions is independent of the choice of c_0 when the null alternative is not included, that is not true of RNCLU: the point at which the value function switches between the first and second branches is determined by $\bar{c}_n$, which depends on c_0. If $c_n = \alpha > 0$ for all values of n (including zero), the principle is restricted critical-level utilitarian. All critical levels exist for the RNCLU principles. For alternatives with average utility above $\bar{c}_n$, the critical level is c_n; for alternatives with a positive average utility that is no greater than $\bar{c}_n$, the critical level can be found by multiplying $c_n/\bar{c}_n$ by average utility; and for alternatives with non-positive average utility, the critical level is zero. All RNCLU principles avoid the repugnant conclusion and satisfy the priority for lives worth living.

The value function for average utilitarianism is average utility, that is:

$$V_{AU}(u) = W_{AU}(n, \mu) = \mu = \frac{1}{n}\sum_{i=1}^{n} u_i. \quad (13)$$

Critical levels exist for all utility vectors and are equal to average utility.

The number-dampened utilitarian family[61] includes both average and classical utilitarianism as members. Its value functions can be written as:

$$V_{NDU}(u) = W_{NDU}(n, \mu) = f(n)\mu = \frac{f(n)}{n}\sum_{i=1}^{n} u_i, \quad (14)$$

where f is a positive-valued function of population size. If $f(n) = n$ or any positive multiple, CU results and if $f(n)$ is independent of n, AU results. Critical levels for NDU are equal to multiples of average utility and the multiple can depend on population size.

The critical level c for an alternative with population size n and average utility μ satisfies:

$$f(n)\mu = f(n+1)\frac{n\mu + c}{n+1}. \quad (15)$$

Consequently:

$$c = h(n)\mu = \left[\frac{f(n)(n+1) - f(n+1)n}{f(n+1)}\right]\mu. \quad (16)$$

The function h must satisfy $h(n) > -n$ for all n.[62]

Restricted number-dampened utilitarianism[63] uses the value function:

$$V_{RNDU}(u) = W_{RNDU}(n, \mu) = \begin{cases} f(n)\mu = f(n) \sum_{i=1}^{n} u_i/n & \text{if } \mu > 0, \\ n\mu = \sum_{i=1}^{n} u_i & \text{if } \mu \leq 0. \end{cases} \quad (17)$$

These principles coincide with number-dampened utilitarianism for positive average utilities and with classical utilitarianism for non-positive average-utility levels. Critical levels are given by equation (16) for alternatives with positive average utilities and are equal to zero for all others. All RNDU principles satisfy the priority for lives worth living, but not all avoid the repugnant conclusion.

A subfamily of the RNDU family uses the single parameter $k \in (0, 1)$, and:

$$f_k(n) = \sum_{j=1}^{n} k^{j-1}. \quad (18)$$

Because $\lim_{n\to\infty} \sum_{j=1}^{n} k^{j-1} = 1/(1-k)$, f_k is bounded above and the principle avoids the repugnant conclusion.[64] The ratio of critical levels to average utility for this principle is given by the function h_k, where:

$$h_k(n) = \frac{\sum_{j=1}^{n} k^{j-1} - nk^n}{\sum_{j=1}^{n+1} k^{j-1}} \quad (19)$$

which is positive for all values of n and increasing with a limit of 1. Because $h_k(1) = (1-k)/(1+k) > 0$, CU is not approximated at small population sizes.

Carlson's combined principle[65] uses an NDU value function applied to non-negative utilities and Sider's geometrism applied to negative utilities. For any $u \in \Omega$, let $\hat{u} \in \mathcal{R}_{+}^{a}$ be a sub-vector of all non-negative utilities and let $\check{u} \in \mathcal{R}_{--}^{b}$ be the sub-vector of all negative utilities arranged in non-decreasing order. Then the value function for the combined theory is:

$$\frac{1}{a} \sum_{i=1}^{a} k^{i-1} \sum_{i=1}^{a} \hat{u}_j + \sum_{j=1}^{b} k^{j-1} \check{u}_j, \quad (20)$$

where $k \in (0, 1)$. If $a = 0$, the first term of the sum in Formula (20) is replaced with zero. This principle satisfies neither weak nor strong Pareto. Consider the two-person alternatives x and y. In x, utility levels are 27 and -1 and in y, utility levels are 28 and 4. If $k = 1/2$, values are $27 - 1 = 26$ for x and $3/2 \times 16 = 24$ for y; so x is better than y, although both people are better off in y.

Principles whose same-number sub-principles are generalized utilitarian use the sum of transformed utilities to rank same-number alternatives. In each case, the transformation g:$\mathcal{R}\to\mathcal{R}$ is continuous and increasing and, without loss of generality, is normalized so that $g(0) = 0$. Transformed utilities result from the application of the transforming function to individual utilities. If the transformation is strictly concave, priority is given to the interests of low-utility individuals and the principle is strictly inequality averse.[66]

Critical-level generalized utilitarianism (CLGU) is a family of principles that contains classical generalized utilitarianism as a special case. The critical level is an arbitrary constant $\alpha \in R$, not necessarily equal to zero. The critical-level generalized-utilitarian value functions can be written as:

$$V_{CLGU}(u) = \sum_{i=1}^{n} \left[g(u_i) - g(\alpha)\right]. \tag{21}$$

Classical generalized utilitarianism results if $\alpha = 0$. Value functions for principles in the restricted critical-level, generalized-utilitarian family are given by:

$$V_{RCLGU}(u) = \begin{cases} \sum_{i=1}^{n} \left[g(u_i) - g(\alpha)\right] & \text{if } \sum_{i=1}^{n} g(u_i) > ng(\alpha), \\ \sum_{i=1}^{n} g(u_i)/\left[ng(\alpha)\right] - 1 & \text{if } 0 < \sum_{i=1}^{n} g(u_i) \leq ng(\alpha), \\ \sum_{i=1}^{n} g(u_i) - 1 & \text{if } \sum_{i=1}^{n} g(u_i) \leq 0. \end{cases} \tag{22}$$

Value functions for principles in the number-sensitive, critical-level, generalized-utilitarian family can be written as:

$$V_{NCLGU}(u) = \sum_{i=1}^{n} \left[g(u_i) - g(c_{i-1})\right] \tag{23}$$

where the critical levels are defined as in the number-sensitive, critical-level utilitarian case. The corresponding restricted principles have value functions that are given by:

$$V_{RNCLGU}(u) = \begin{cases} \sum_{i=1}^{n} \left[g(u_i) - g(c_{i-1})\right] & \text{if } \sum_{i=1}^{n} g(u_i) > \sum_{i=1}^{n} g(c_{i-1}), \\ \sum_{i=1}^{n} g(u_i)/\sum_{i=1}^{n} g(c_{i-1}) - 1 & \text{if } 0 < \sum_{i=1}^{n} g(u_i) \leq \sum_{i=1}^{n} g(c_{i-1}), \\ \sum_{i=1}^{n} g(u_i) - 1 & \text{if } \sum_{i=1}^{n} g(u_i) \leq 0. \end{cases} \tag{24}$$

It is possible to have $\sum_{i=1}^{n} g(c_i - 1) = 0$ for some n and, in that case, the middle branch of Formula (24) does not apply.

Population principles in the number-dampened, generalized-utilitarian family have value functions that can be written as:

$$V_{NDGU}(u) = \frac{f(n)}{n} \sum_{i=1}^{n} g(u_i) \tag{25}$$

and value functions for the restricted number-dampened, generalized-utilitarian principles are given by:

$$V_{RNDGU}(u) = \begin{cases} f(n)\sum_{i=1}^{n} g(u_i)/n & \text{if } \sum_{i=1}^{n} g(u_i) > 0, \\ \sum_{i=1}^{n} g(u_i) & \text{if } \sum_{i=1}^{n} g(u_i) \leq 0. \end{cases} \tag{26}$$

Suppose that a population principle has same-number sub-principles that are utilitarian and, in addition, for any one-person alternative, there exists an alternative with at least two individuals that is at least as good. That is, the principle has the following property.

Minimal one-person trade-off

For all $\xi \in \mathcal{R}$, there exists $u \in \Omega \setminus \mathcal{R}^1$ such that $uR\xi 1_1$.

This condition is implied by the existence of critical levels, as is easy to verify. We now obtain the following impossibility result.

Theorem 1

There exists no population principle that has utilitarian same-number sub-principles and satisfies avoidance of the repugnant conclusion, avoidance of the sadistic conclusion and minimal one-person trade-off.

Proof

Suppose, by way of contradiction, that R has utilitarian same-number sub-principles and satisfies avoidance of the repugnant conclusion, avoidance of the sadistic conclusion and minimal one-person trade-off. For any $n \in \mathcal{Z}_{++} \setminus \{1\}$, $\xi \in \mathcal{R}_{++}$, $\epsilon \in (0, \xi)$ and $\delta \in (0, \epsilon)$, let $\nu = (n\xi + \delta)/(n-1)$. By same-number utilitarianism:

$$(\nu \mathbf{1}_{n-1}, -\delta) I \xi \mathbf{1}_n. \tag{27}$$

By avoidance of the sadistic conclusion:

$$(\nu \mathbf{1}_{n-1}, \delta \mathbf{1}_m) R (\nu \mathbf{1}_{n-1}, -\delta) \tag{28}$$

for all $m \in \mathcal{Z}_{++}$. Because $\delta < \epsilon$, there exists $\bar{m} \in \mathcal{Z}_{++} \setminus \{1\}$ such that

$$\frac{(n-1)\nu + \bar{m}\delta}{n - 1 + \bar{m}} < \varepsilon. \tag{29}$$

By same-number utilitarianism:

$$\varepsilon \mathbf{1}_{n-1+\bar{m}} P (\nu \mathbf{1}_{n-1}, \delta \mathbf{1}_{\bar{m}}). \tag{30}$$

Formulas (30), (28) and (27) and the transitivity of R imply:

$$\varepsilon \mathbf{1}_{n-1+\bar{m}} P \xi \mathbf{1}_n. \tag{31}$$

It remains to be shown that Formula (31) holds for $n = 1$ as well. By minimal one-person trade-off, there exists $\hat{n} \in \mathcal{Z}_{++} \setminus \{1\}$ and $\hat{u} \in R^{\hat{n}}$ such that:

$$\hat{u}R\xi\mathbf{1}_1. \tag{32}$$

Choose $\omega \in R_{++}$ such that $\omega \geq \xi$ and $\omega > \hat{u}_i$ for all $i \in \{1, \ldots, \hat{n}\}$. By same-number utilitarianism:

$$\omega\mathbf{1}_{\hat{n}}P\hat{u}. \tag{33}$$

By the argument used to obtain Formula (31), there exists $\hat{m} \in Z_{++}$ such that:

$$\varepsilon\mathbf{1}_{\hat{n}-1+\hat{m}}P\omega\mathbf{1}_{\hat{n}}. \tag{34}$$

Formulas (34), (33) and (32) and the transitivity of R imply $\epsilon 1_{\hat{n}-1+\hat{m}}P\xi 1_1$. This, in turn, together with Formula (31), implies the repugnant conclusion — a contradiction.

Instead of the value function W, a value function $\mathring{W}$ which depends on average utility μ and total utility $\tau = n\mu$ can be employed to represent population principles. Provided that average utility is non-zero, $\mathring{W}(\tau, \mu) = W(\tau/\mu, \mu)$. It has been suggested that the value function should be increasing in both total utility and average utility.[67] If this monotonicity requirement is combined with our basic axioms of strong Pareto, anonymity and the existence of critical levels, then some critical levels must be negative. Therefore, we obtain the following impossibility result.

Theorem 2

There exists no population principle that satisfies strong Pareto, anonymity, the existence of critical levels and non-negative critical levels if the associated value function $\mathring{W}$ exists and is increasing in total utility and in average utility.

Proof

By way of contradiction, suppose that R is an ordering satisfying the axioms in the theorem statement (R can be used rather than R because of anonymity). Consider the utility vector $\xi 1_n$ with $n \in Z_{++}$ and $\xi < 0$. Since the value function $\mathring{W}$ is increasing in average utility, we have $(\xi 1_n, 0)P\xi 1_n$ because, in moving from $\xi 1_n$ to $(\xi 1_n, 0)$, average utility increases and total utility is unchanged. Furthermore, we have $\xi 1_n P \xi 1_{n+1}$ because average utility is unchanged and total utility decreases when moving from $\xi 1_n$ to $\xi 1_{n+1}$. By strong Pareto, the critical level c for $\xi 1_n$ must satisfy $\xi < c < 0$, contradicting non-negative critical levels.

We conclude with another impossibility result. There is no population principle that satisfies our basic axioms, utility independence, avoidance of the repugnant conclusion and the priority for lives worth living.

Theorem 3

There exists no population principle that satisfies strong Pareto, continuity, anonymity, the existence of critical levels, utility independence, avoidance of the repugnant conclusion and the priority for lives worth living.

Theorem 3 follows from Theorem 10 of Blackorby, Bossert and Donaldson.[68] The original theorem is slightly stronger. Instead of the existence of critical levels, it uses an axiom that requires the existence, for each population size, of only one utility vector with a critical level. In addition, instead of the priority for lives worth living, it uses the weaker axiom avoidance of the sadistic conclusion.

notes

We thank Thomas Hurka who suggested the restricted number-dampened family in his comments on a paper presented at the Utilitarianism 2000 conference of the International Society for Utilitarian Studies held at Wake Forest University, March 2000. His suggestion inspired the discovery of the restricted number-sensitive, critical-level family. We also thank Gustaf Arrhenius, John Broome, Erik Carlson, John Weymark, Catherine Wilson, an associate editor and two referees for helpful comments and criticisms. Financial support through grants from the Social Sciences and Humanities Research Council of Canada is gratefully acknowledged.

1. W. Thomson, 'On the Axiomatic Method and its Recent Applications to Game Theory and Resource Application', *Social Choice and Welfare* 18 (2001): 327–86.
2. H. Sidgwick, *The Methods of Ethics* (New York: Dover, 1966 [1907]), p. 414.
3. Readers who are interested in the extension of population principles to non-humans are referred to C. Blackorby and D. Donaldson, 'Pigs and Guinea Pigs: A Note on the Ethics of Animal Exploitation', *Economic Journal* 102 (1992): 1345–69; P. Singer, *Practical Ethics* (Cambridge: Cambridge University Press, 1979), Chs 3 and 5.
4. See J. Broome, 'Goodness is Reducible to Betterness: The Evil of Death is the Value of Life' in *The Good and the Economical: Ethical Choices in Economics and Management*, edited by P. Koslowski (Berlin: Springer-Verlag, 1993); J. Broome, 'Weighing Lives' (unpublished manuscript, 1999), Ch. 8; D. Heyd, *Genethics: Moral Issues in the Creation of People* (Berkeley: University of California Press, 1992), Ch. 1; J. McMahan, 'Wrongful Life: Paradoxes in the Morality of Causing People to Exist' (unpublished manuscript, 1996); D. Parfit, *Reasons and Persons* (Oxford: Oxford University Press, 1984), Appendix G.
5. R. Sikora, 'Is It Wrong to Prevent the Existence of Future Generations?', in *Obligations to Future Generations*, edited by R. Sikora and B. Barry (Philadelphia, PA: Temple, 1978), pp. 112–66.
6. D. Parfit, 'On Doing the Best for our Children', in *Ethics and Population*, edited by M. Bayles (Cambridge: Schenkman, 1976); D. Parfit, 'Future Generations, Further Problems', *Philosophy and Public Affairs* 11 (1982): 113–72; D. Parfit, *Reasons and Persons*, Ch. 19.
7. The axiomatic investigation of principles for ranking alternatives in a fixed-population environment began with K. Arrow, *Social Choice and Individual Values*, 2nd edn. (New York: Wiley, 1963 [1951]) and, although there is a small body of work that is concerned with population questions, social-choice theory has, for the most part, ignored those issues.
8. C. Blackorby, W. Bossert and D. Donaldson, 'Population Principles with Number-Dependent Critical Levels', *Journal of Public Economic Theory* 4 (2002): 347–68.
9. C. Blackorby and D. Donaldson, 'Social Criteria for Evaluating Population Change', *Journal of Public Economics* 25 (1984): 13–33.

10. Blackorby and Donaldson, ‘Population Principles’.
11. Y-K. Ng, ‘Social Criteria for Evaluating Population Change: An Alternative to the Blackorby-Donaldson Criterion’, *Journal of Public Economics* 29 (1986): 375–81.
12. T. Hurka, ‘Comment on “Population Principles with Number-Dependent Critical Levels”’, unpublished manuscript.
13. C. Blackorby, W. Bossert and D. Donaldson, ‘Utilitarianism and the Theory of Justice’, in *Handbook of Social Choice and Welfare*, edited by K. Arrow, A. Sen and K. Suzumura (Amsterdam: Elsevier, 2002), pp. 543–96; Blackorby and Donaldson, ‘Population Principles’; D. Parfit, ‘Equality or Priority?’, *Ratio* 10 (1997): 202–21; J. Broome, ‘Equality or Priority: A Useful Distinction’ in *Fairness and Goodness in Health*, edited by D. Wikler and C. Murray (Geneva: World Health Organization, forthcoming); M. Fleurbaey, ‘Equality Versus Priority: How Relevant is the Distinction?’, in *Fairness and Goodness in Health*, edited by D. Wikler and C. Murray (Geneva: World Health Organization, forthcoming).
14. For discussions of principles that provide incomplete rankings, see C. Blackorby, W. Bossert and D. Donaldson, ‘Quasi-Orderings and Population Ethics’, *Social Choice and Welfare* 13 (1996): 129–50; Broome, ‘Weighing Lives’, Ch. 10.
15. See C. Blackorby, W. Bossert and D. Donaldson, ‘Intertemporally Consistent Population Principles: Classical Utilitarian Principles’, in *Social Choice Reexamined*, Vol. 2, edited by K. Arrow, A. Sen and K. Suzumura (London: Macmillan, 1996), pp. 137–62; C. Blackorby, W. Bossert and D. Donaldson, ‘Critical-Level Utilitarianism and the Population Ethics Dilemma’, *Economics and Philosophy* 13 (1997): 197–230.
16. For a discussion in an inter-temporal framework, see C. Blackorby, W. Bossert and D. Donaldson, ‘Intertemporal Population Ethics: Critical-Level Utilitarian Principles’, *Econometrica* 63 (1995): 1303–20.
17. See C. Blackorby, W. Bossert and D. Donaldson, ‘Uncertainty and Critical-Level Population Principles’, *Journal of Population Economics* 11 (1998): 1–20; C. Blackorby, W. Bossert and D. Donaldson, ‘Utilitarianism and the Theory of Justice’.
18. J. Griffin, *Well-Being: Its Meaning, Measurement, and Moral Importance*, Part 1 (Oxford: Clarendon, 1986). Griffin allows for incomplete individual goodness relations, but such considerations are beyond the scope of this article.
19. L. Sumner, *Welfare, Happiness, and Ethics* (Oxford: Clarendon, 1996), Ch. 6.
20. See C. Blackorby, W. Bossert and D. Donaldson, ‘Information Invariance in Variable-Population Social-Choice Problems’, *International Economic Review* 40 (1999): 403–22.
21. E. Carlson, ‘Mere Addition and the Two Trilemmas of Population Ethics’, *Economics and Philosophy* 14 (1998): 283–306.
22. P. Dasgupta, *An Inquiry into Well-Being and Destitution* (Oxford: Clarendon, 1993), Ch. 13, advocates rules such as this.
23. A version of the example (without Euclid) is due to Parfit, ‘On Doing the Best’ and ‘Future Generations’.
24. Parfit, ‘On Doing the Best’ and ‘Future Generations’.
25. See the Appendix for a proof. In inter-temporal settings, we call a related axiom ‘independence of the utilities of the dead’. See Blackorby, Bossert and Donaldson, ‘Intertemporal Population Ethics’; ‘Intertemporally Consistent Population

Principles'; and 'Birth-Date Dependent Population Ethics: Critical-Level Principles', *Journal of Economic Theory* 77 (1997): 260–84. Similar axioms have been suggested by G. Arrhenius and K. Bykvist, 'Future Generations and Interpersonal Compensations', *Uppsala Prints and Preprints in Philosophy*, No. 21 (Uppsala University, 1995); P. Hammond, 'Consequentialist Demographic Norms and Parenting Rights', *Social Choice and Welfare* 5 (1988): 127–46; J. McMahan, 'Problems of Population Theory', *Ethics* 92 (1981): 96–127; and Parfit, 'On Doing the Best' and 'Future Generations'.

26. Carlson, 'Mere Addition'.
27. See, for example, Blackorby, Bossert and Donaldson, 'Critical-Level Utilitarianism'.
28. Sikora's Pareto-plus axiom requires the *ceteris paribus* addition of a person whose lifetime utility level is above neutrality to be ranked as an improvement. In conjunction with our basic axioms and non-negative critical levels, his axiom implies that all critical levels must be zero. See Sikora, 'Is it Wrong'.
29. Parfit, 'On Doing the Best'; 'Future Generations'; *Reasons and Persons*, Ch. 17.
30. See also T. Cowen, 'What do we Learn from the Repugnant Conclusion?', *Ethics* 106 (1996): 754–95.
31. Heyd, *Genethics*, p. 57.
32. C. Blackorby, W. Bossert, D. Donaldson and M. Fleurbaey, 'Critical Levels and the (Reverse) Repugnant Conclusion', *Journal of Economics* 67 (1998): 1–15.
33. For this result and several variants, see, G. Arrhenius, 'An Impossibility Theorem for Welfarist Axiologies', *Economics and Philosophy* 16 (2000); Blackorby, Bossert, Donaldson and Fleurbaey, 'Critical Levels'; C. Blackorby and D. Donaldson, 'Normative Population Theory: A Comment', *Social Choice and Welfare* 8 (1991): 261–7; Carlson, 'Mere Addition'; McMahan, 'Problems of Population Theory'; and Parfit, 'On Doing the Best', 'Future Generations', *Reasons and Persons*, Ch. 19.
34. T. Sider, 'Might Theory X be a Theory of Diminishing Marginal Value?', *Analysis* 51 (1991): 202–13.
35. See, Arrhenius and Bykvist, 'Future Generations and Interpersonal Compensations'.
36. G. Arrhenius, 'An Impossibility Theorem'; 'Future Generations: A Challenge for Moral Theory', PhD dissertation (Uppsala University, 2000): Ch. 10.
37. Blackorby, Bossert and Donaldson, 'Population Principles'.
38. Arrhenius, 'An Impossibility Theorem'.
39. Arrhenius provides impossibility theorems using, for example, avoidance of the sadistic conclusion as well as a non-sadistic condition which is weaker. See Arrhenius, 'An Impossibility Theorem' and 'Future Generations: A Challenge', Ch. 10.
40. For sufficient conditions on general population principles (not necessarily same-number utilitarian) that guarantee the existence of a value function, see C. Blackorby, W. Bossert and D. Donaldson, 'Population Ethics and the Existence of Value Functions', *Journal of Public Economics* 82 (2001): 301–8; J. Broome, 'Representing an Ordering when the Population Varies', in *Social Choice and Welfare* (forthcoming).
41. Carlson, 'Mere Addition'; A. Carter, 'Moral Theory and Global Population', *Proceedings of the Aristotelian Society* 99 (1999): 289–313; Parfit, *Reasons and*

Persons, Ch. 18. See Arrhenius, 'Future Generations: A Challenge', for a critical discussion.

42. Although we do not require it, all welfarist principles can be extended to cover the null alternative (the one in which no one ever lives). This can be accomplished by specifying a critical level for the null alternative.
43. Fixed critical levels were proposed by Parfit, 'On Doing the Best' and 'Future Generations'. The critical-level utilitarian family was characterized by Blackorby, Bossert and Donaldson, 'Intertemporal Population Ethics'; 'Uncertainty and Critical-Level Population Principles'; and Blackorby and Donaldson, 'Social Criteria'.
44. Blackorby, Bossert and Donaldson, 'Uncertainty and Critical-Level Population Principles' and 'Utilitarianism and the Theory of Justice'.
45. Blackorby, Bossert and Donaldson, 'Population Principles'.
46. Ibid.
47. Ng, 'Social Criteria'.
48. The discussion of the NDU family is based on Blackorby, Bossert and Donaldson, 'Population Principles'.
49. T. Hurka, 'Value and Population Size', *Ethics* 93 (1983): 496–507.
50. Arrhenius, 'An Impossibility Theorem' and 'Future Generations: A Challenge', Ch. 4.
51. Hurka, 'Comment on "Population Principles"'.
52. Blackorby, Bossert and Donaldson, 'Population Principles'.
53. Heyd, *Genethics*, Ch. 2; and Parfit, 'On Doing the Best', 'Future Generations', and *Reasons and Persons*, Ch. 17.
54. Non-decreasing critical levels are consistent with Carlson's intuition, which is discussed in subsection 3.1.
55. F. Feldman, 'Justice, Desert, and the Repugnant Conclusion', *Utilitas* 7 (1995): 189–206; also published in F. Feldman, *Utilitarianism, Hedonism, and Desert: Essays in Moral Philosophy* (Cambridge: Cambridge University Press, 1997), pp. 193–214.
56. Feldman, *Utilitarianism, Hedonism, and Desert*, p. 194; see also Arrhenius, 'Future Generations: A Challenge', Ch. 9.
57. Carlson, 'Mere Addition'.
58. See the Appendix for formulas for the generalized principles.
59. Arrhenius, 'An Impossibility Theorem'.
60. Ibid.
61. Ng, 'Social Criteria'.
62. Blackorby, Bossert and Donaldson, 'Population Principles'.
63. Hurka, 'Comment on "Population Principles"'.
64. Blackorby, Bossert and Donaldson, 'Population Principles'.
65. Carlson, 'Mere Addition'.
66. A principle is strictly inequality averse if, and only if, when comparing alternatives with the same population size and the same total utility, its ranking is consistent with the Lorenz criterion.
67. See, for example, Carter, 'Moral Theory and Global Population'.
68. Blackorby, Bossert and Donaldson, 'Population Principles'.