

Cahier 2001-06

The Axiomatic Approach to Population Ethics

BLACKORBY, Charles

BOSSERT, Walter

DONALDSON, David

Département de sciences économiques

Université de Montréal

Faculté des arts et des sciences

C.P. 6128, succursale Centre-Ville

Montréal (Québec) H3C 3J7

Canada

<http://www.sceco.umontreal.ca>

SCECO-information@UMontreal.CA

Téléphone : (514) 343-6539

Télécopieur : (514) 343-7221

Ce cahier a également été publié par le Centre interuniversitaire de recherche en économie quantitative (CIREQ) sous le numéro 06-2001.

This working paper was also published by the Center for Interuniversity Research in Quantitative Economics (CIREQ), under number 06-2001.

ISSN 0709-9231

CAHIER 2001-06

**THE AXIOMATIC APPROACH
TO POPULATION ETHICS**

Charles BLACKORBY¹, Walter BOSSERT² and David DONALDSON³

¹ University of British Columbia and GREQAM

² Centre de recherche et développement en économique (C.R.D.E.) and Département de sciences économiques, Université de Montréal

³ University of British Columbia

March 2001

The authors wish to thank Thomas Hurka who suggested the restricted number-dampened family in his comments on a paper presented at the International Society for Utilitarian Studies Conference, *Utilitarianism 2000*, at Wake Forest University, March 2000. His suggestion inspired the discovery of the restricted number-sensitive critical-level family. Financial support through a grant from the Social Sciences and Humanities Research Council of Canada is gratefully acknowledged.

RÉSUMÉ

Cet article étudie diverses familles de principes de population d'un point de vue axiomatique. En plus de l'utilitarisme avec niveau critique, de l'utilitarisme avec niveau critique sensible à (ou tempéré par) la taille de la population, et de leurs généralisations, nous considérons la famille restreinte suggérée par Hurka et nous en proposons deux nouvelles : la famille restreinte avec niveau critique et la famille restreinte avec niveau critique sensible à la taille de la population. Certains sous-ensembles de ces familles restreintes possèdent un niveau critique non négatif et évitent les conclusions répugnantes et sadiques, mais violent une condition d'indépendance importante. Nous défendons les niveaux critiques positifs.

Mots clés : éthique de population, méthodologie axiomatique

ABSTRACT

This paper 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 families and their generalized counterparts, we consider the restricted number-dampened family (suggested by Hurka) and introduce two new families : the restricted critical-level and restricted number-dependent critical-level families. Subsets of the restricted families have nonnegative critical levels and avoid both the repugnant and sadistic conclusions but fail to satisfy an important independence condition. We defend the critical-level principles with positive critical levels.

Key words : population ethics, axiomatic methodology

1. Introduction

The axiomatic method was introduced into population ethics by Parfit [1976, 1982, 1984] when he accused classical utilitarianism of leading to the repugnant conclusion. Because average utilitarianism, the only other population principle that was well-known at the time, does not have this property but has other defects, Parfit’s criticism encouraged the search for population principles that make better trade-offs. Avoidance of the repugnant conclusion had become an axiom that acceptable principles should satisfy.

The axiomatic investigation of principles for ranking alternatives in a fixed-population environment began with Arrow’s [1951, 1963] book *Social Choice and Individual Values* 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.

In this paper, we attempt to test population principles by using axioms that are, in our view, the ones that are the most attractive on ethical grounds. We investigate principles that rank alternatives—complete histories of the universe from remote past to distant future—according to their goodness. 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.

Welfarist population principles employ only information about the lifetime well-being of the individuals alive in various alternatives in order to rank them. 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 can rank actions whose consequences are uncertain (Blackorby, Bossert and Donaldson [1998, 2001a]). Because welfarist principles use welfare information alone, it is important that they be partnered with a comprehensive account of well-being such as the one provided by Griffin [1986] or that of Sumner [1996].

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. A neutral life is one which is as good as one in which the individual has no experiences.¹ 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 nonexistence. It makes perfect sense, of course, to say that an individual gains

¹ See Broome [1993, 1999].

or loses by continuing to live because of surviving a life-threatening illness, say. Such a change affects length of life, not existence itself.²

Although we consider only human populations here, it is possible to extend all the principles to take account of the well-being of sentient non-humans. Sidgwick [1907, 1966, p. 414] argues that we should “extend our concern to all the beings capable of pleasure and pain whose feelings are affected by our conduct”. Our simplification makes our presentation more transparent. Readers who are interested in the extension of welfarist principles to non-humans are referred to Blackorby and Donaldson [1992].

It is possible to apply population principles to single periods of time but the results may be inconsistent with timeless application of the same principle and may recommend killing people whose lives are worth living (see Blackorby, Bossert and Donaldson [1996b, 1997b]). For example, the average-utilitarian principle, applied to a single period, recommends the painless killing of people whose level of well-being is below the average as long as no one else is affected. For that reason, we focus exclusively on the social evaluation of complete histories using lifetime utilities. In that case, killing shortens a life and does not change population size. As a consequence, those counterintuitive recommendations do not arise.

Our investigation uses four basic axioms—strong Pareto, continuity, anonymity and existence of critical levels—and six additional axioms, five of which apply to the ranking of alternatives with different population sizes. Three axioms in the latter group ask principles to ignore the utilities or number 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 both the utilities and existence 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 nonnegative. It ensures that a principle does not sanction the *ceteris-paribus* creation of people whose lives are not worth living. The last two axioms ask principles to avoid the repugnant and the sadistic conclusions. A principle leads to the repugnant conclusion if any 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 is above neutrality but arbitrarily close to it. A principle implies the sadistic conclusion (Arrhenius [2000]) if every alternative in which each person’s utility level is below neutrality is ranked as better

² For discussions, see Heyd [1992, Chapter 1], McMahan [1996] and Parfit [1984, Appendix G].

than some alternative in which each person in a larger population has a utility level above neutrality.

Although the axioms are desirable on a priori grounds, we state a result in the appendix (Theorem 2) that shows that there is no population principle that satisfies the basic axioms, utility (and therefore existence) independence, has nonnegative critical levels and avoids both the repugnant and sadistic conclusions. It is possible, however, to find population principles that satisfy some but not all of the axioms. We consider the critical-level utilitarian family (Blackorby and Donaldson [1984]), the restricted critical-utilitarian family (introduced in this paper), the number-sensitive critical-level utilitarian family (Blackorby, Bossert and Donaldson [2000]), the number-dampened utilitarian family (Ng [1986]) and the restricted number-dampened utilitarian family (Hurka [2000]). Although all of these principles rank alternatives with the same population size using utilitarianism, there are expanded families which rank those alternatives with generalized utilitarianism. It uses transformed utilities and allows for aversion to utility inequality (see Blackorby, Bossert and Donaldson [2000]).

Using our axioms, we examine all these principles and defend the critical-level principles with positive critical levels. In addition, the performance of all the principles in the tests provided by the axioms offers readers a way to rank the principles for themselves.

Section 2 sets out the axioms and Section 3 presents the principles that are consistent with same-number utilitarianism and investigates their consistency with the axioms. Section 4 provides a defence of the critical-level utilitarian principles with positive critical levels and Section 5 concludes. The main part of the paper offers several examples but contains no mathematics: mathematical statements of the axioms, principles and theorems are in the appendix.

2. The Axioms

Population principles rank alternatives according to their goodness. Although there are some principles that do not rank all possible alternatives, the principles that we investigate in this paper provide *orderings* of alternatives: at-least-as-good-as relations that are complete, reflexive and transitive.³ 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 .⁴

³ See Blackorby, Bossert and Donaldson [1996a] and Broome [1999] for discussions of principles that provide incomplete rankings.

⁴ 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 y and y is at least as good as z , x is better than z .

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, they 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 . Because it is possible for the same person to be born at different times in different alternatives, this axiom has the effect of ruling out discounting of the well-being of future generations.⁵ It can be weakened to accommodate discounting, however (see Blackorby, Bossert and Donaldson [1997a]).

Our second basic axiom is *continuity*. It requires fixed-population trade-offs to be gradual, without sudden jumps from better to worse.

Our third basic axiom is an impartiality condition called *anonymity* and it requires any two alternatives which are the same except for the identities of those alive to be ranked as equally good. Given strong Pareto, this implies that any two alternatives in which the same number of people have the same levels of lifetime utility are equally good. Anonymity provides a solution to Parfit's [1984] 'non-identity problem' and ensures all individual's interests receive equal treatment.

Anonymous population principles provide a single ordering of utility vectors of all dimensions. 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 . Both the classical and average-utilitarian principles 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. The critical level for the first alternative is that level of utility for the added person that makes the two equally good. We assume that critical levels exist for all alternatives and call our fourth basic assumption *existence of critical levels*.

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.

⁵ See Broome [1994], Cowen [1992] and Cowen and Parfit [1992].

2.1. 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 5 for the child and 6 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 6, 0 and 6. 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.⁶

	Parent	First Child	Second Child	Euclid
x	6	5		same in both
y	6	0	6	same in both line

Table 1

The parent wants to know which alternative is better. Parfit [1976, 1982] considers this example and he assumes that utility levels other than those of family members and potential members are irrelevant. That assumption is satisfied if principles such as classical utilitarianism are used to rank the alternatives but not when 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 4, average utility is 5 in x and 4 in y and, if his utility level is -10 , average utility is $1/3$ in x and $1/2$ 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 3.7 in x and 3 in y so x is regarded as better in this 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 are regarded as morally irrelevant. The simplest 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 are alive and have the same utility levels in

⁶ This example, without Euclid, is due to Parfit [1976, 1982].

both. Same-number independence is satisfied by all the principles considered in this paper including those based on generalized utilitarianism.

Utility independence requires the ranking of any two alternatives to be independent of the utility levels of individuals who are alive 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 are alive and have the same utility levels in both. It is the strongest of the independence axioms and implies the other two.⁷ It allows population principles to be applied to affected individuals only.

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 space ship and, after drifting 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 depends on the number of people in the other group even though the number is unaffected by the decisions under consideration. And, if the population principle satisfies same-number independence only, the ranking of the feasible alternatives depends 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. An argument for utility independence could be made, however; it would focus on the total number of people who ever live. For that reason, the numbers of the long dead, of unaffected independent groups or of unaffected people in the far future would count in social rankings.

⁷ In intertemporal settings, we call a related axiom 'independence of the utilities of the dead'. Hammond [1988] and McMahan [1981] have suggested similar axioms.

Existence independence is attractive for practical reasons as well. Information about the number and utility levels of the long dead or of future people whose existence and well-being are unaffected by decisions taken in the present is very difficult to obtain. For that reason alone, we might want to use population principles that satisfy existence independence and do not need such information.

2.2. *Nonnegative 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 *nonnegative critical levels* and it has an obvious justification.⁸ 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 6 and -14 , its critical level is -4 and the addition of a person whose utility level is -2 is ranked as desirable.

2.3. *The Repugnant Conclusion*

A population principle implies the repugnant conclusion (Parfit [1976,1982,1984])⁹ if any alternative in which each member of the population has a positive utility level, no matter how large, is ranked as worse than an alternative in which a sufficiently large population has a utility level that is above neutrality but arbitrarily close to it. Such principles may recommend the creation of a large population in which each person is poverty stricken. As Heyd [1992, p. 57] 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?”

Classical utilitarianism leads to the repugnant conclusion and this is a powerful criticism. It is possible to show that any principle which satisfies our basic axioms, is weakly inequality averse and has critical levels that are all nonpositive implies the repugnant conclusion.¹⁰ It follows that, if the repugnant conclusion is to be avoided, some critical levels must be positive. Average utilitarianism is a principle which has some positive critical levels and does not imply the repugnant conclusion.

⁸ Sikora’s [1978] 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 the basic axioms and nonnegative critical levels, his axiom implies that all critical levels must be zero.

⁹ See also Cowen [1996].

¹⁰ See Arrhenius [2000], Blackorby, Bossert, Donaldson and Fleurbaey [1998], Blackorby and Donaldson [1991], Carlson [1998], McMahan [1981] and Parfit [1976, 1982, 1984].

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

2.4. *The Sadistic Conclusion*

Arrhenius [2000] has suggested that population principles should avoid the *sadistic conclusion* which obtains if and only if every alternative in which each person has a negative utility level is ranked as better than an alternative in which each person in a sufficiently large population has a level above neutrality. Such principles rank alternatives in which each person is below neutrality as better than some alternatives in which each person is above it. This suggests that population principles should satisfy an axiom which we call *avoidance of the sadistic conclusion*.¹¹

A stronger axiom would require a population principle to rank all alternatives in which each person is above neutrality as better than all those in which each person is below it. As we shall see, all of the restricted principles satisfy this but, because our impossibility result does not need the stronger version, we do not use it.

3. Welfarist Population Principles

A welfarist population principle provides a single ordering of utility vectors which is used, along with information about well-being associated with alternatives, to order them. Many principles have value functions that are able to do the same job. 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. If a principle has a value function, one alternative is better than another if and only if the value of the first, as measured by the function, is greater than the value of the second. In addition, the two alternatives are equally good if and only if their values are the same.

Not all population principles that satisfy our basic axioms have value functions. In order to guarantee the existence of such a function, an extension of continuity to different numbers can be used (Blackorby, Bossert and Donaldson [2001b]). All of the principles discussed in this paper have value functions and all but a few satisfy the extended continuity axiom. The mathematical results that we summarize below do not, however, require extended continuity.

If a principle satisfies same-number independence, then its same-number subprinciples must be generalized utilitarian: one utility vector is at least as good as another of the same

¹¹ A related axiom requires principles to avoid the ‘reverse repugnant conclusion’, which is also implied by all principles that imply the sadistic conclusion. It is difficult to see why it is repugnant and, for that reason, we have omitted it. See Blackorby, Bossert, Donaldson and Fleurbaey [1998] and Carlson [1998].

dimension if and only if the sum of transformed utilities is at least as great in the first.¹² A continuous and increasing function is used to transform utilities and, if it is (strictly) concave, the principle expresses (strict) aversion to utility inequality. Such a transform might give a greater weight to the utilities of those who are below neutrality, for example. Without loss of generality, the transform can be chosen so that the transformed value of zero is zero. Same-number independence does not imply that the transforms used are the same for all population sizes, but both utility and existence independence do imply that (Blackorby, Bossert and Donaldson [2000]). Consequently, if same-number independence is the only independence axiom used, it is necessary to assume directly that the transforms are the same for different population sizes.

All principles whose same-number subprinciples are utilitarian are also principles whose same-number subprinciples are generalized utilitarian but the utilitarian same-number orderings are not implied by any of our independence axioms. The axiom *incremental equity* is sufficient for same-number utilitarianism. It demands impartiality with respect to utility increases or decreases. If a single individual's utility level changes by a given amount, the goodness or badness of the change does not depend on the identity who receives the increment. Incremental equity and weak Pareto (an increase in each person's well-being is ranked as an improvement) together characterize same-number utilitarianism (Blackorby, Bossert and Donaldson [2001a, Theorem 12]).

The representative utility for a utility vector is that level of utility which, if assigned to each person, produces a vector which is as good as the original. If a principle's same-number subprinciples are utilitarian, representative utility is average utility. If a principle has a value function, it can be written in terms of representative utility and population size.¹³ The value function must be increasing in representative utility but its response to population size may be different (increasing, decreasing or unresponsive) for different levels of representative utility.

Carter [1999] (and, implicitly, Parfit [1984]) has suggested that value functions for principles whose same-number subprinciples are utilitarian should be expressible in terms of average utility and total utility. Because population size is equal to total utility divided by average utility, the value function for any principle that has one 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 1) that any principle with this property has

¹² See, for example, Blackorby, Bossert and Donaldson [1995, 2001a].

¹³ See Blackorby, Bossert, and Donaldson [2001b] for sufficient conditions on general population principles—not necessarily same-number utilitarian—that guarantee the existence of a value function.

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, however, we work with the representative-utility – population-size representation.

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 (see Blackorby, Bossert and Donaldson [2000]). All the principles discussed below are members of larger families whose same-number subprinciples are generalized utilitarian. In this paper, we discuss only principles whose same-number subprinciples are utilitarian and refer the reader to Blackorby, Bossert and Donaldson [2000] for their generalized counterparts.

3.1. Classical Utilitarianism

The value function for classical utilitarianism (CU) is the sum of utilities. Consequently, one alternative is better than another if and only if the sum of the utilities of those who ever live is greater in it. The value function for CU can also be written as the product of population size and average utility. 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 isovalue curves. Points on isovalue 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 sixty, thirty, zero and minus thirty 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 nonnegative critical levels and avoids the sadistic conclusion. As is well known, however, it leads to the repugnant conclusion. The repugnant conclusion is implied because the isovalue curve for any average-utility – population-size pair with positive average utility approaches the population-size axis as numbers increase. This is true of isovalue curves I and II in Figure 1. As a consequence, it is possible to find a population size such that an arbitrarily small average-utility level paired with that population size is ranked as better.

3.2. 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 (Blackorby and Donaldson [1984], Blackorby, Bossert and Donaldson [1995, 1998]). If the critical level is zero, classical utilitarianism results. The 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.

Critical-level utilitarianism with a critical level of thirty is illustrated in Figure 2. The four isovalue curves are constructed in the same way as the isovalue 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.

Isovalue curves for average-utility – population-size pairs with average utility above the critical level do not drop below isovalue curve II. If the critical level is positive, therefore, CLU avoids the repugnant conclusion and has nonnegative critical levels.

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 leads to the sadistic conclusion, however. This can be seen in Figure 2 by looking at isovalue 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 minus thirty is ranked as better than any alternative in which each of four people has a utility level of ten. 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 one. 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 two is applied to the family alone, values are 7 for x and 6 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 (Blackorby, Bossert and Donaldson [2001a, Theorem 16]; see also Blackorby and Donaldson [1984] and Blackorby, Bossert and Donaldson [1995]). Existence independence implies that critical levels are the same for all alternatives. If a single individual is added to a utility-unaffected population at the critical level, the two alternatives are equally good by definition of the critical level. 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. Avoidance of the repugnant conclusion requires a positive critical level in both the generalized-utilitarian and utilitarian cases.

3.3. *Restricted Critical-Level Utilitarianism*

Critical-level-utilitarian principles with positive critical levels can be modified so that all members of the resulting family avoid both the repugnant and sadistic conclusions. 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 average utility divided by the parameter less one if average utility is positive and no greater than the parameter, and by total utility less one if average utility is nonpositive. It is illustrated for a parameter value of thirty in Figure 3 (isovalue curves I–IV are defined as before and isovalue curve V is added). We call the resulting family restricted critical-level utilitarianism (RCLU). It ranks alternatives whose average utilities are greater than thirty using CLU (isovalue curve I), alternatives whose average utilities are positive and no greater than thirty using average utilitarianism (isovalue curves II and V), and alternatives whose average utilities are nonpositive with classical utilitarianism (isovalue 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 nonnegative 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 nonpositive. Consequently, all critical levels are nonnegative.

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

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 the critical-level parameter equal to two is used to rank x and y . If Euclid’s utility level is 10, average utility is 7 in x and 6.5 in y which are both greater than 2. Consequently, values are 15 for x and 18 for y , and the alternative with one child is better. Now suppose

that Euclid’s utility level is -14 . Then average utilities are -1 for x and $-1/2$ for y and values are -4 for x and -3 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.

3.4. 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 covered, its critical level is c_0 and if it is 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, \dots, c_{n-1} from average utility and multiplying by 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 thirty for population sizes greater than one. If average utility is constant, population increases are bad for average-utility – population-size pairs that lie below isovalue curve III (and have negative average utility) but, for pairs that are above curve III, the goodness or badness of population-size increases is ambiguous and 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 thirty.

Satisfaction of nonnegative critical levels requires all critical levels to be nonnegative. In this case, 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 (Blackorby, Bossert and Donaldson [2000, Theorem 7]). If, in addition, critical levels are nondecreasing, the repugnant conclusion is avoided if and only if there is at least one positive critical level. Nondecreasingness of critical levels ensures, in that case, that all the critical levels for higher population sizes are positive.

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 nine for population sizes above three. Without loss of generality, c_0 may be chosen to be zero. Writing Euclid’s utility level as ξ , values are $11 + \xi$ for x and $12 + \xi$ for y and y is better than x for all values of ξ . Suppose, now, that we discover that Euclid had an identical twin brother whose lifetime utility level was also equal to ξ . In that case, values are $11 + 2\xi$ for x and $3 + 2\xi$ 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 (Blackorby, Bossert and Donaldson [2000, Theorem 3]). 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 imply the sadistic conclusion (Blackorby, Bossert and Donaldson [2000, Theorem 19]). This can be seen in Figure 4 by noting that isovalue curve IV crosses the population-size axis. We show, in the following subsection, that it is possible to modify these principles so that they imply neither but, in that case, neither existence nor utility independence is satisfied.

3.5. 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 nonnegative, nondecreasing 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, \dots, c_{n-1} where c_0 is nonnegative. 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 nonpositive. 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 isovalue 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 average utility greater than \bar{c}_n are better than all pairs with 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 nonpositive. RNCLU ranks alternatives with average utilities above \bar{c}_n with the corresponding NCLU principle and alternatives with nonpositive 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 positive and no greater than \bar{c}_n , and zero if average utility is nonpositive. Consequently, all of the RNCLU principles satisfy nonnegative 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 nondecreasing and at least one is positive.

All of the restricted number-sensitive critical-level utilitarian principles avoid the sadistic conclusion. The reason is that all alternatives with negative average utility are

ranked as worse than all those with positive average utility. An example is provided by isovalue curve IV in Figure 5: it does not cross the population-size axis.

All 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.

3.6. 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 isovalue curves indicate 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 one person enjoys a trivially higher utility level.

Because the addition of a single 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 nonnegative critical levels is not satisfied.

Isovalue 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, the sadistic conclusion is avoided. The discussion of average utilitarianism following the disabled-child example summarized in Table 1 demonstrates that AU satisfies neither utility nor existence independence.

3.7. Number-Dampened Utilitarianism

The number-dampened utilitarian (NDU) family (Ng [1986]) 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.

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 [1983]. Such a case is illustrated in Figure 7. For that principle, the function takes on the values 1, 2, 2.6 and 3 for population sizes one, two, three, and four or more respectively. For population sizes one and two, the principle coincides with CU and, for population sizes of four or more, the principle coincides with AU.

It can be shown (see 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, .31, .47 and 1 for population sizes one, two, three, and four or more.

An important case 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.

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 (Blackorby, Bossert and Donaldson [2000, Theorem 12]). 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. We have been able to prove that all of the 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 (Blackorby, Bossert and Donaldson [2000, Theorem 17]). In the case that the ratio is nonconstant, the requirement that it be nondecreasing is consistent with Hurka's suggestion (Hurka [1983]). In order to avoid the repugnant conclusion, any NDU principle for which the ratio of critical levels to average utilities is a nondecreasing, positive real number between zero and one must approximate average utilitarianism as population size becomes large (Blackorby, Bossert and Donaldson [2000, Theorem 18]).

Because every NDU principle ranks all alternatives with positive average utilities as better than all alternatives with negative average utilities, all of the NDU principles avoid the sadistic conclusion.

Every member of the number-dampened utilitarian and the number-dampened generalized-utilitarian families satisfies same-number independence. None of them other than classical utilitarianism and classical generalized utilitarianism satisfy either utility or existence independence. This follows from the fact that utility independence, together with our basic axioms, characterizes the number-sensitive critical-level generalized utilitarian family (Blackorby, Bossert and Donaldson [2000, Theorem 2]). The only principles in that family that are also members of the number-dampened generalized-utilitarian family are the classical generalized-utilitarian principles. Because the only classical generalized-utilitarian principle in the number-dampened utilitarian family is CU, the result follows.

3.8. Restricted Number-Dampened Utilitarianism

Suggested by Hurka [2000], the restricted number-dampened utilitarian (RNDU) family of principles provides a partial solution to one of the most important difficulties 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 nonpositive. The restricted version of the example of Figure 7 is illustrated in Figure 8. Above the population-size axis, the isovalue curves are the same for both principles. But, below the population-size axis, isovalue curves 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 nonpositive average utilities are all zero and, hence, nonnegative. Critical levels for alternatives with positive average utilities are not necessarily nonnegative, 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, because the average of four and minus two multiplied by four is equal to four, the critical level is minus two. For any NDU principle, Theorem 14 in Blackorby, Bossert and Donaldson [2000] demonstrates that the ratios of critical levels to average utilities are nonnegative if and only if the ratio of the function that multiplies average utility to population size does not increase as population size increases. This condition, applied to RNDU principles, is clearly necessary and sufficient for nonnegative critical levels.

Two special cases of RNDU have, however, nonnegative 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 nonpositive. Its critical levels are equal to average utility for alternatives with positive average utility and zero for alternatives whose average utility is nonpositive. The restricted version of the NDU family for which the ratio of critical levels to average utilities is a positive constant between zero and one also has nonnegative critical levels: they are equal to the positive constant when average utility is positive and to zero when average utility is nonpositive.

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 avoid the sadistic conclusion.

Of all the RNDU principles, only classical utilitarianism satisfies utility independence. This observation follows from the discussion of independence and the NDU principles.

4. Which Principles Are Best?

It is interesting to ask whether the axioms that we have presented tell us which principles are best. To investigate that question, we have summarized the results of the previous section in Table 2. The table is divided into two sections: the first lists families of principles that can be regarded as generalizations of CU and the second list 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 (RAU) 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 positive constant between zero and one and the third selects NDU principles for which the ratio of critical levels to average utilities is positive, nondecreasing 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 a general result. The answer is ‘yes’ and we state it in the appendix (Theorem 2; see also Blackorby, Bossert and Donaldson [2000, Theorem 10]). There is no population principle that satisfies our basic axioms and utility independence, has non-negative critical levels, and avoids both the repugnant and sadistic conclusions. 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 and avoidance of the repugnant and sadistic conclusions. In searching for the best 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 the table. 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 [1992] and Parfit [1976, 1982, 1984], 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 lead to the sadistic conclusion and those that satisfy neither utility nor existence independence but do not imply the sadistic conclusion. 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 avoidance of the sadistic conclusion. 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. 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 utility 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.

Although existence independence is attractive on both ethical and practical grounds, the principles that satisfy it and avoid the repugnant conclusion necessarily imply the sadistic conclusion. It is important to ask, therefore, whether avoidance of the sadistic conclusion is important enough to make us abandon existence independence. Our view is that it is not. To understand our first reason, consider an alternative in which each person has a negative utility level and choose any utility level that is positive but less than the critical level. An alternative in which each person has this utility level and the same number of people as the original is ranked as better by CLU with a positive critical level. But, because the higher utility level is below the critical level, population increases with added people at the same level are bad and it is possible to find a larger population size such that the resulting alternative is regarded as worse than the original one. But the principle clearly states that such expansions are undesirable. Our second reason relies on the fact that, in a set of feasible alternatives, an alternative that maximizes the CLU value function is the most desirable. If there is a feasible alternative with average utility above the critical level, it is best according to both CLU and its restricted counterpart. If the average utility of those who ever live is above the critical level, therefore, we need not be concerned with the sadistic conclusion.

It is important to note that, in our framework, axioms do not refer exclusively to those who are presently alive. Theorem 2 shows that principles that satisfy nonnegative critical levels, avoidance of the repugnant conclusion and avoidance of the sadistic conclusion necessarily satisfy neither utility nor existence independence. Consequently, avoiding the sadistic conclusion imposes informational demands that are very great. In the example of the disabled child summarized in Table 1, we must have information about the long dead in order to rank changes that affect only present and future generations.

Suppose, however, that the view that both the repugnant and sadistic conclusions should be avoided is accepted. If we restrict attention to principles with nonnegative critical levels, this requires us to give up both utility and existence independence. Subfamilies in Table 2 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, nondecreasing and approaches one as numbers increase.

Restricted average utilitarianism may be attractive to some, but it retains the stark trade-offs of AU for alternatives with positive average utilities and this leads us to reject it. 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 solar system and universe. If that occurs, then value becomes confounded with constraints. The same consideration applies to the number-sensitive critical-level utilitarian principles. If this argument is accepted, we are left with restricted CLU as the only satisfactory family whose members avoid both the repugnant and sadistic conclusions.

If we were forced to choose between the CLU subfamily with positive critical levels and restricted CLU, we would choose the former because we think that satisfaction of existence independence is more important than avoidance of the sadistic conclusion. But the choice may not be necessary. If we can somehow know that the average utility of those who will ever live is above the critical-level parameter in all feasible alternatives, then the RCLU and CLU rank them in the same way. We suspect, moreover, that, in practical situations, decision makers will be forced to behave as if they had adopted a principle that satisfies existence independence because of the unavailability of information. If our conjecture is correct, the ethical decision that matters is the choice of a positive critical level, not the choice between a restricted and an unrestricted CLU principle.

It is true, of course, that we do not have to choose principles whose same-number sub-principles are utilitarian. Both the critical-level utilitarian and restricted critical-level utilitarian families are themselves members of two larger families: the critical-level generalized-utilitarian and the restricted critical-level generalized-utilitarian families. The axioms we have employed do not rule out the inequality aversion that generalized utilitarian families can represent. In addition, the critical-level generalized-utilitarian subfamily with positive critical levels exhausts the possibilities for principles that satisfy the basic axioms, existence independence and avoidance of the repugnant conclusion.

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 (negative) orthant of \mathcal{R}^n is denoted by \mathcal{R}_{++}^n (\mathcal{R}_{--}^n). For $n \in \mathcal{Z}_{++}$, $\mathbf{1}_n$ is the vector consisting of n ones. Individuals are indexed by positive integers. 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 $i, j \in \{1, \dots, n\}$. We use the notation $\Omega = \cup_{n \in \mathcal{Z}_{++}} \mathcal{R}^n$. Letting u_\emptyset denote the utility vector associated with the alternative where no one is alive, we define $\Omega_\emptyset = \Omega \cup \{u_\emptyset\}$.

For each alternative, the associated population size is denoted by $n \in \mathcal{Z}_{++}$ and the associated vector of individual identities by $\pi = (\pi_1, \dots, \pi_n) \in \mathbf{Z}^n$. Utilities for those alive are $u = (u_1, \dots, u_n) \in \mathcal{R}^n$ where, for all $i \in \{1, \dots, n\}$, u_i is individual π_i 's lifetime utility in the alternative in question.

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 any $n \in \mathcal{Z}_{++}$, $(\pi, u), (\rho, v) \in \mathbf{Z}^n \times \mathcal{R}^n$, if ρ is a permutation of π and v is the same permutation of u , the $(\pi, u) \overset{*}{I} (\rho, v)$.

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

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

Existence of Critical Levels: For all $n \in \mathcal{Z}_{++}$, for all $(\pi, u) \in \mathbf{Z}^n \times \mathcal{R}^n$, there exist $j \in \mathcal{Z}_{++} \setminus \{\pi_1, \dots, \pi_n\}$ 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 on Ω 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 $u R v$. Furthermore, the same-number restrictions of R must be anonymous: if v is a permutation of u , $u I v$. 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. This can be done without loss of generality because the four basic axioms are maintained throughout.

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 $u, v \in \mathcal{R}^n$, for all $w, s \in \mathcal{R}^m$,

$$(u, w)R(v, w) \iff (u, s)R(v, s). \quad (1)$$

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 $r \in \mathcal{Z}_{++}$, for all $w, s \in \mathcal{R}^r$,

$$(u, w)R(v, w) \iff (u, s)R(v, s). \quad (2)$$

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 $u, v, w \in \Omega$,

$$(u, w)R(v, w) \iff uRv. \quad (3)$$

The requirement that critical levels be nonnegative is defined as follows.

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

A population principle implies the repugnant conclusion (Parfit [1976, 1982, 1984]) if and only if, for any population size $n \in \mathcal{Z}_{++}$, any positive utility level ξ , and any utility level $\varepsilon \in (0, \xi)$, 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 $\varepsilon \in (0, \xi)$ such that, for all $m > n$, $\xi \mathbf{1}_n R \varepsilon \mathbf{1}_m$.

A population principle implies the sadistic conclusion (Arrhenius [2000]) if and only if, for any population size $n \in \mathcal{Z}_{++}$ and any negative utility level ξ , there exist a population

size $m > n$ and a positive utility level ε such that an alternative in which n people experience utility level ξ is better than an alternative in which m people experience ε .

Avoidance of the Sadistic Conclusion: There exist $n \in \mathcal{Z}_{++}$ and $\xi \in \mathcal{R}_{--}$ such that, for all $m > n$ and for all $\varepsilon \in \mathcal{R}_{++}$, $\varepsilon \mathbf{1}_m R \xi \mathbf{1}_n$.

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

$$\mu = \mu^n(u) = \frac{1}{n} \sum_{i=1}^n u_i. \quad (4)$$

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 subprinciples), the value function W can be written in terms 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 $u \in \mathcal{R}^n$, for all $v \in \mathcal{R}^m$,

$$u R v \iff W(n, \mu^n(u)) \geq W(m, \mu^m(v)). \quad (5)$$

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

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

$$W_{CU}(n, \mu) = n\mu = \sum_{i=1}^n u_i. \quad (6)$$

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

Critical-level utilitarianism (CLU) is a family of principles, one for each value of a fixed utility level which is the critical level for every alternative.¹⁴ The CLU value functions are given by

$$W_{CLU}(n, \mu) = n(\mu - \alpha) = \sum_{i=1}^n (u_i - \alpha) \quad (7)$$

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.

¹⁴ The CLU principles are introduced in Blackorby and Donaldson [1984]. See Blackorby, Bossert and Donaldson [1995, 1998] for further discussions.

The restricted critical-level utilitarian (RCLU) family of principles is introduced in this paper. Its value functions are given by

$$W_{RCLU}(n, \mu) = \begin{cases} n(\mu - \alpha) = \sum_{i=1}^n (u_i - \alpha) & \text{if } \mu > \alpha, \\ \mu/\alpha - 1 = \sum_{i=1}^n u_i / (n\alpha) - 1 & \text{if } 0 < \mu \leq \alpha, \\ n\mu - 1 = \sum_{i=1}^n u_i - 1 & \text{if } \mu \leq 0, \end{cases} \quad (8)$$

where α is positive. W_{RCLU} 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 nonpositive. Consequently, all alternatives whose average utility is above the critical-level parameter are better than all whose average utility is positive and no greater than it and these alternatives are, in turn, better than all whose average utilities are nonpositive. 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 both the repugnant and sadistic conclusions.

The number-sensitive critical-level utilitarian (NCLU) family is a generalization of the CLU family.¹⁵ 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).¹⁶ The value functions for the NCLU principles can be written as

$$W_{NCLU}(n, \mu) = n(\mu - \bar{c}_n) = \sum_{i=1}^n (u_i - c_{i-1}) \quad (9)$$

where

$$\bar{c}_n = \frac{1}{n} \sum_{i=1}^n c_{i-1}. \quad (10)$$

\bar{c}_n is the average of c_0 and the critical levels for population sizes 1 to $n - 1$. 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 (RNCLU), the second family introduced in this paper, are represented by value functions that are generalizations of those for RCLU and they are derived from those for NCLU. \bar{c}_n is defined

¹⁵ See Blackorby, Bossert and Donaldson [2000].

¹⁶ If the null alternative were included, c_0 would be its critical level.

by (10) and we assume that $c_0 \geq 0$, the c_n are nondecreasing, and at least one c_n is positive. The value functions can be written as

$$W_{RNCLU}(n, \mu) = \begin{cases} n(\mu - \bar{c}_n) = \sum_{i=1}^n (u_i - c_{i-1}) & \text{if } \mu > \bar{c}_n, \\ \mu/\bar{c}_n - 1 = \sum_{i=1}^n u_i / \sum_{i=1}^n c_{i-1} - 1 & \text{if } 0 < \mu \leq \bar{c}_n, \\ n\mu - 1 = \sum_{i=1}^n u_i - 1 & \text{if } \mu \leq 0. \end{cases} \quad (11)$$

It is possible to have $\bar{c}_n = 0$ for some n and, in that case, the middle branch of (11) 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 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 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 nonpositive average utility, the critical level is zero. All RNCLU principles avoid both the repugnant and sadistic conclusions.

The value function for average utilitarianism (AU) is average utility, that is,

$$W_{AU}(n, \mu) = \mu = \frac{1}{n} \sum_{i=1}^n u_i. \quad (12)$$

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

The number-dampened utilitarian (NDU) family (Ng [1986]) includes both average and classical utilitarianism as members. Its value functions can be written as

$$W_{NDU}(n, \mu) = f(n)\mu = \frac{f(n)}{n} \sum_{i=1}^n u_i, \quad (13)$$

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

$$f(n)\mu = f(n+1) \frac{n\mu + c}{n+1}. \quad (14)$$

Consequently,

$$c = h(n)\mu = \left[\frac{f(n)(n+1) - f(n+1)n}{f(n+1)} \right] \mu. \quad (15)$$

The function h must satisfy $h(n) > -n$ for all n (Blackorby, Bossert, and Donaldson [2000, Theorem 11]).

Restricted number-dampened utilitarianism (Hurka [2000]) uses the value function

$$W_{RNDU}(n, \mu) = \begin{cases} f(n)\mu = f(n) \sum_{i=1}^n u_i/n & \text{if } \mu > 0, \\ n\mu = \sum_{i=1}^n u_i & \text{if } \mu \leq 0. \end{cases} \quad (16)$$

These principles coincide with number-dampened utilitarianism for positive average utilities and with classical utilitarianism for nonpositive average-utility levels. Critical levels are given by (15) for alternatives with positive average utilities and are equal to zero for all others. All RNDU principles avoid the sadistic conclusion but not all avoid the repugnant conclusion.

Instead of the value function W , a value function V which depends on average utility μ and total utility $\tau = n\mu$ can be employed to represent population principles. Provided that average utility is nonzero, $V(\tau, \mu) = W(\tau/\mu, \mu)$. It has been suggested that the value function should be increasing in both total utility and average utility (see, for example, Carter [1999]). If this monotonicity requirement is combined with our basic axioms strong Pareto, anonymity and existence of critical levels, then some critical levels must be negative. Therefore, we obtain the following impossibility result.

Theorem 1: *There is no population principle that satisfies strong Pareto, anonymity, existence of critical levels, and nonnegative critical levels if the associated value function V exists and is increasing in total utility and 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 $\overset{*}{R}$ because of anonymity). Consider the utility vector $\xi \mathbf{1}_n$ with $n \in \mathcal{Z}_{++}$ and $\xi < 0$. Because the value function V is increasing in average utility, we have $(\xi \mathbf{1}_n, 0) P \xi \mathbf{1}_n$ because, in moving from $\xi \mathbf{1}_n$ to $(\xi \mathbf{1}_n, 0)$, average utility increases and total utility is unchanged. Furthermore, $\xi \mathbf{1}_n P \xi \mathbf{1}_{n+1}$ because average utility is unchanged and total utility decreases when moving from $\xi \mathbf{1}_n$ to $\xi \mathbf{1}_{n+1}$. By strong Pareto, the critical level c for $\xi \mathbf{1}_n$ must satisfy $\xi < c < 0$, contradicting nonnegative critical levels. ■

We conclude with a second impossibility result due to Blackorby, Bossert and Donaldson [2000, Theorem 10]. It establishes an incompatibility between avoidance of the repugnant and sadistic conclusions and utility independence, given our basic axioms.

Theorem 2: *There is no population principle that satisfies strong Pareto, continuity, anonymity, existence of critical levels, utility independence, nonnegative critical levels, avoidance of the repugnant conclusion and avoidance of the sadistic conclusion.*

The theorem of Blackorby, Bossert and Donaldson [2000] is slightly stronger because existence of critical levels can be weakened to an axiom that requires only the existence, for each population size, of one utility vector with a critical level.

REFERENCES

- Arrhenius, G., 2000, An impossibility theorem for welfarist axiologies, *Economics and Philosophy* **16**, 247–266.
- Arrow, K., 1951 (second ed. 1963), *Social Choice and Individual Values*, Wiley, New York.
- Blackorby, C., W. Bossert, and D. Donaldson, 1995, Intertemporal population ethics: critical-level utilitarian principles, *Econometrica* **63**, 1303–1320.
- Blackorby, C., W. Bossert, and D. Donaldson, 1996a, Quasi-orderings and population ethics, *Social Choice and Welfare* **13**, 129–150.
- Blackorby, C., W. Bossert, and D. Donaldson, 1996b, Intertemporally consistent population principles: classical utilitarian principles, in *Social Choice Reexamined*, volume 2, K. Arrow, A. Sen, and K. Suzumura, eds., MacMillan, London, 137–162.
- Blackorby, C., W. Bossert, and D. Donaldson, 1997a, Birth-date dependent population ethics: critical-level principles, *Journal of Economic Theory* **77**, 260–284.
- Blackorby, C., W. Bossert, and D. Donaldson, 1997b, Critical-level utilitarianism and the population ethics dilemma, *Economics and Philosophy* **13**, 197–230.
- Blackorby, C., W. Bossert, and D. Donaldson, 1998, Uncertainty and critical-level population principles, *Journal of Population Economics* **11**, 1–20.
- Blackorby, C., W. Bossert, and D. Donaldson, 2000, Population principles with number-dependent critical levels, Discussion Paper No. 00–04, University of British Columbia Department of Economics.
- Blackorby, C., W. Bossert, and D. Donaldson, 2001a, Utilitarianism and the theory of justice, in *Handbook of Social Choice and Welfare*, K. Arrow, A. Sen, and K. Suzumura, eds., Elsevier, Amsterdam, forthcoming.
- Blackorby, C., W. Bossert, and D. Donaldson, 2001b, Population ethics and the existence of value functions, *Journal of Public Economics*, forthcoming.
- Blackorby, C., W. Bossert, D. Donaldson, and M. Fleurbaey, 1998, Critical levels and the (reverse) repugnant conclusion, *Journal of Economics* **67**, 1–15.
- Blackorby, C. and D. Donaldson, 1984, Social criteria for evaluating population change, *Journal of Public Economics* **25**, 13–33.
- Blackorby, C. and D. Donaldson, 1991, Normative population theory: a comment, *Social Choice and Welfare* **8**, 261–267.
- Blackorby, C. and D. Donaldson, 1992, Pigs and guinea pigs; a note on the ethics of animal exploitation, *Economic Journal* **102**, 1345–1369.
- Broome, J., 1993, 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*, P. Koslowski, ed., Springer-Verlag, Berlin, 69–83.

- Broome, J., 1994, Discounting the future, *Philosophy and Public Affairs* **23**, 128–156.
- Broome, J., 1999, *Weighing Lives*, unpublished draft manuscript.
- Carlson, E., 1998, Mere addition and the two trilemmas of population ethics, *Economics and Philosophy* **14**, 283–306. d’Aspremont, C. and L. Gevers, 1977, Equity and the informational basis of collective choice, *Review of Economic Studies* **44**, 199–209.
- Carter, A., 1999, Moral theory and global population, *Proceedings of the Aristotelian Society* **99**, 289–313.
- Cowen, T., 1992, Consequentialism implies a zero rate of intergenerational discount, in *Justice Between Age Groups and Generations*, P. Laslett and J. Fishkin, eds., Yale University Press, New Haven, 162–168.
- Cowen, T., 1996, What do we learn from the repugnant conclusion?, *Ethics* **106**, 754–795.
- Cowen, T. and D. Parfit, 1992, Against the social discount rate, in *Justice Between Age Groups and Generations*, P. Laslett and J. Fishkin, eds., Yale University Press, New Haven, 144–161.
- Griffin, J., 1986, *Well-Being: Its Meaning, Measurement, and Moral Importance*, Clarendon, Oxford.
- Hammond, P., 1988, Consequentialist demographic norms and parenting rights, *Social Choice and Welfare* **5**, 127–146.
- Heyd, D., 1992, *Genethics: Moral Issues in the Creation of People*, University of California Press, Berkeley.
- Hurka, T., 1983, Value and population size, *Ethics* **93**, 496–507.
- Hurka, T., 2000, Comment on ‘Population principles with number-dependent critical levels,’ unpublished manuscript.
- McMahan, J., 1981, Problems of population theory, *Ethics* **92**, 96–127.
- McMahan, J., 1996, Wrongful life: paradoxes in the morality of causing people to exist, University of Illinois, unpublished manuscript.
- Ng, Y.-K., 1986, Social criteria for evaluating population change: an alternative to the Blackorby-Donaldson criterion, *Journal of Public Economics* **29**, 375–381.
- Parfit, D., 1976, On doing the best for our children, in *Ethics and Population*, M. Bayles, ed., Schenkman, Cambridge, 100–102.
- Parfit, D., 1982, Future generations, further problems, *Philosophy and Public Affairs* **11**, 113–172.
- Parfit, D., 1984, *Reasons and Persons*, Oxford University Press, Oxford/New York.
- Roberts, K., 1980, Interpersonal comparability and social choice theory, *Review of Economic Studies* **47**, 421–439.
- Sidgwick, H., 1966 (originally published in 1907), *The Methods of Ethics*, Dover, New York.
- Sikora, R., 1978, Is it wrong to prevent the existence of future generations?, in *Obligations to Future Generations*, R. Sikora and B. Barry, eds., Temple, Philadelphia, 112–166.
- Sumner, L., 1996, *Welfare, Happiness, and Ethics*, Clarendon, Oxford.

◇ ◇ ◇

Table 2

	Same-Number Independence	Utility Independence	Existence Independence	Nonnegative Critical Levels	Avoidance of the Repugnant Conclusion	Avoidance of the Sadistic Conclusion
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: nonnegative, nondecreasing critical levels and some positive critical levels
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 between zero and one
11. Restricted Number-Dampened Utilitarianism: restricted version of 10
12. Number-Dampened Utilitarianism: ratio of critical level to average utility is positive, nondecreasing and approaches one as numbers increase
13. Restricted Number-Dampened Utilitarianism: restricted version of 12




Figure 1: Classical Utilitarianism




Figure 2: Critical-Level Utilitarianism




Figure 3: Restricted Critical-Level Utilitarianism




Figure 4: Number-Sensitive Critical-Level Utilitarianism




Figure 5: Restricted Number-Sensitive Critical-Level Utilitarianism




Figure 6: Average Utilitarianism




Figure 7: Number-Dampened Utilitarianism




Figure 8: Restricted Number-Dampened Utilitarianism