CHAPTER 1

Introduction

In discussions of environmental problems, descriptions of histories of the world over the next century or two are often linked to policy alternatives. In these histories, standards of living as well as the size and composition of populations may differ at different times. It might be said, for example, that, if present policies continue, we can expect a large population, low average standards of living, environmental degradation, and wars as the relative scarcity of land increases. On the other hand, if responsible environmental and population policies are adopted, we might be told to expect a smaller population with higher average standards of living and a world with more resources and less environmental damage for subsequent generations.

Other policy decisions have population consequences as well. Examples include the allocation of public funds to prenatal care, the design of aid packages to developing countries, public funding of education, legalization of birth-control devices, resource conservation, expenditures on public health, and the design of social security systems. In addition, individual fertility decisions affect population composition and size both directly and indirectly, as the children have children. In each of these cases, history will depend, to some extent, on actions taken.

It is important to ask what it means to say that one of these histories is better than another. Because population size and individual identities may be different, this is not a straightforward question. To attempt to provide an answer, we investigate principles for social evaluation that can be used, together with factual information, to rank histories. Such an exercise requires examining the idea of the common good, extended so that it can take account of differences in population composition and size. We call principles that make comparisons of alternative histories with respect to their social goodness population principles.

Because changes in policies rarely lead to gains for everyone, principles must be able to balance gains and losses. In addition, one policy might result in a larger population with a lower standard of living than another does. Principles therefore have an ethical dimension. The idea of the common or social good provides the needed trade-offs by means of a social goodness relation. Instead of using two relations – one for betterness and one for equal goodness – we

focus on a single at-least-as-good-as relation, which combines the two. Two histories are equally good if and only if each is at least as good as the other, and one is better than another if and only if it is at least as good as the other and it is not the case that the other is at least as good as it. Goodness relations may or may not be complete. If they are, every pair of histories is ranked.

Our main focus is on principles for social evaluation that are commonly called welfarist (Sen 1979). These principles use information about individual well-being to rank histories, disregarding all other information. Welfarism rests, in the main, on the idea that any two histories with the same population in which everyone is equally well off are equally good, a condition commonly called the Pareto-indifference axiom.

Welfarist principles regard values such as individual liberty and autonomy as instrumental: valuable because of their contribution to well-being. In addition, virtues and fair procedures may have instrumental value. Because of this, it is important to employ a comprehensive notion of well-being such as that of Griffin (1986) or Sumner (1996). We focus on lifetime well-being and include enjoyment, pleasure and the absence of pain, good health, length of life, autonomy, liberty, understanding, accomplishment, and good human relationships as aspects of it. Individuals who are autonomous and fully informed may have self-regarding preferences that accord with their well-being, but we do not assume that they do. In addition, sentient nonhuman animals have experiences, and it is possible to take account of their interests. We reserve our discussion of that possibility for Chapter 11, however, and focus on human beings in the rest of the book.

Most principles are impartial in the sense that individual identities do not matter. Because impartiality is ethically fundamental, all the principles that we investigate satisfy it.

In a large part of the book, we require the social at-least-as-good-as relation to be complete. We are not committed to this, but we see the exercise as important. We want to know whether there are reasonable principles that provide complete orderings of alternative histories. Some principles, however, fail to rank some of them. This may occur for one of two reasons: it may be impossible to obtain enough information, or there may be no fact of the matter. In the first case, it is possible to generate an incomplete ranking by using several principles that generate complete rankings when information is perfect. In the second case, principles must be able to take account of the incommensurabilities directly. See Chapter 7 for a discussion.

Although the idea of the common good plays an important role in discussions of government policies and consequentialist moral theories, we do not consider the normative status of actions. In situations in which actions lead with certainty to particular histories, we use population principles to rank actions according to their goodness but do not provide a link with obligations. As a result, our work is consistent with rule consequentialism, which assesses moral rules by estimating their probable consequences, and with supererogation, the idea that some actions that have very good consequences but require great sacrifices by the agents taking them are beyond the call of duty.

1.1 FRAMEWORK AND METHOD OF INVESTIGATION

We use the term *alternative* as a label for a complete history of the universe. Each alternative is associated with a list of the identities of all those individuals who ever live. Information about the well-being of each person on the list (welfare information) is available for each alternative. In addition, social nonwelfare information together with individual nonwelfare information is known for each person on the list. Nonwelfare information may include individual information such as dates of birth, lengths of life, and character traits. Social nonwelfare information may include information about individual liberties, freedom of the press, and so on.

We distinguish several cases. In the simplest one – the single-profile case – the information associated with each alternative is fixed. The other cases allow some of the information to vary and, in that case, we say that there are multiple information profiles.

A goodness relation for an individual ranks alternatives according to their goodness for that person. Individual lifetime well-being is measured by an index called a utility function, which is a representation of an individual goodness relation. If the value of the index is greater for one alternative than for another, the individual's life is better, for him or her, in that alternative. Similarly, equal utility values correspond to alternatives that are equally good for the person. Uniqueness properties of utility functions are discussed in Chapter 2.

If a person's life, taken as a whole, is worth living, we say that the level of well-being associated with his or her life is above neutrality, and, if it is not, the lifetime well-being of the individual is below neutrality. If lifetime well-being is at neutrality, it is not the case that, as a whole, the life is worth living and it is not the case that it is not worth living. We adopt the standard convention that neutrality is associated with a utility value of zero. We do not believe a person can be made better off by being born into a life above neutrality (see Chapter 2): well-being requires being.

Population principles, welfarist or not, can be captured in a formal model by employing social-evaluation functionals. A social-evaluation functional associates an ordering of the alternatives with every profile of information.[1] Welfarism obtains if and only if there is a single ordering of vectors of individual utilities that can be used, together with the information about well-being in a profile, to order the alternatives. Some welfarist principles, such as classical and average utilitarianism and all members of the critical-level utilitarian class, rank some alternatives with different population sizes differently but coincide when ranking alternatives with the same population size.

Most of the principles we investigate require interpersonal comparisons of well-being. Statements such as "person 10's life is better, as a whole, in

---

[1] Orderings are at-least-as-good-as relations that are reflexive (every alternative is as good as itself), complete, and transitive (if x is at least as good as y and y is at least as good as z, x is at least as good as z). See Chapter 2, Sections 2.2 and 2.6, for discussions.

alternative x than person 25's life is in alternative y" are assumed to be meaningful, provided that 10 is alive in x and 25 is alive in y.

Our focus is, for the most part, on welfarist principles, but we also investigate other principles. An example is a principle that uses information about individual birth dates to discount future utilities. In addition, we consider principles that can rank uncertain alternatives: lists of possible alternatives with probabilities assigned to each component. Such principles can be used to rank actions of individuals or governments according to their goodness when the consequences of actions are uncertain.

Throughout the book, we use axioms to examine a large collection of problems. Axioms for population principles, for example, isolate hypothetical situations and impose simple properties on the behavior of the principles in these situations (Thomson 2001). One such axiom is minimal individual goodness. If two alternatives have the same population and one is ranked as better than the other, the axiom requires it to be better for at least one person. The underlying idea is that our intuitions are reliable in such situations.

Population principles make two kinds of trade-offs: those that determine the relative weight assigned to the interests of different individuals, and those that trade off inequality-adjusted per-capita well-being against population size. Axioms place conditions on those trade-offs and allow us to discover and formulate classes of population principles.

Many of the theorems that we prove are characterizations of classes of principles such as critical-level generalized utilitarianism or number-dampened generalized utilitarianism. Such theorems show that each member of the class satisfies all the axioms used and that no other principle does. Some collections of axioms are inconsistent, however. Formal statements of inconsistencies are called impossibility theorems and they provide important tests of the intuitions that motivate the axioms, calling for the assignment of priorities. Consequently, some of the discussion of axioms necessarily involves assessments of their relative worth.

Our aim in using the axiomatic method is to contribute to readers' understanding of the various classes of principles. An additional contribution is provided, we believe, by applying the principles to simple choice situations. For that reason, we examine several examples. Some of these focus on choices in a certain universe and others deal with choices when consequences are uncertain.

Anonymity axioms capture the idea of impartiality in social evaluation. One version applies to pairs of alternatives in which the same number of people live and requires two alternatives to be ranked as equally good if the social nonwelfare information is the same in both, and individual welfare *and* nonwelfare information is the same but attached to different individuals. As a special case, anonymity also applies to comparisons of alternatives with the same population.

A principle implies the repugnant conclusion (Parfit 1976, 1982, 1984, Chapter 19) if every alternative in which each person experiences a utility level above neutrality is ranked as worse than some alternative in which each member of

a larger population has a utility level that is above neutrality but 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. Avoidance of the repugnant conclusion is an important axiom that ethically acceptable principles should satisfy. Although classical utilitarianism leads to the repugnant conclusion, many principles avoid it but some of these have other undesirable features. To make further progress in finding the class of acceptable principles, therefore, additional axioms are needed.

It is possible to rank alternatives by using information in a single period of time only. If that is done and information in other periods is the same, it is important that the resulting rankings coincide with those obtained from a timeless point of view, using lifetime utilities. In Chapter 9, we find that such a consistency requirement, when combined with a few standard axioms, results in principles that lead to the repugnant conclusion. The reason is that the shortening of a life may appear to be a population change if viewed from within a single period. We conclude that history must matter to some extent. The critical-level generalized utilitarian principles allow the existence of every person whose life is over to be ignored but require that we take account of the well-being of any person whose life may extend into the present or future. In addition, these principles can be applied consistently to any group of people, such as a nation or a particular generation, if its members are the only ones affected by a change.

If each member of a set of feasible actions is associated with an alternative (without uncertainty), a possibility for rational action is to choose the action that leads to the best alternative. (If more than one alternative is best, an action that leads to any one of them will do.) Such choices are said to be rationalized by an ordering of the alternatives. To investigate choices directly, we use axioms that specify various conditions on the choices themselves. Two issues arise. The first is whether observed choices can be rationalized by an ordering of the alternatives. The second is conditional on the existence of such an ordering and is concerned with its properties.

For any alternative, consider another with one additional person alive and suppose, hypothetically, that each member of the common population has the same level of well-being in both. A critical level for the utility vector that corresponds to the first alternative is a level of utility for the added person that makes the two alternatives equally good. In general, critical levels may depend on utility levels of those who are alive and their number. Critical levels exist for all the principles considered here.

Most of the principles that we discuss rank alternatives with the same population size by using generalized utilitarianism. It ranks these alternatives with the sum of transformed utilities, where the transformation is increasing and continuous. Same-number utilitarian principles are also same-number generalized utilitarian, but the converse is not true. Same-number generalized utilitarian principles can be averse to inequality of well-being or, equivalently, give priority to the interests of people with low levels of well-being.

The best-known welfarist principles are classical utilitarianism, which uses the sum of utilities to rank alternatives, and average utilitarianism, which uses average utility. Because classical utilitarianism leads to the repugnant conclusion and average utilitarianism sometimes declares the ceteris paribus addition of a person whose life is not worth living to be good, these principles typically are rejected. There are, however, many principles that avoid both of these properties. Our inquiry groups principles into the following classes: critical-level generalized utilitarianism, restricted critical-level generalized utilitarianism, number-sensitive critical-level generalized utilitarianism, restricted number-sensitive critical-level generalized utilitarianism, number-dampened generalized utilitarianism (Ng 1986), and restricted number-dampened generalized utilitarianism (Hurka 2000). The classical-utilitarian and average-utilitarian principles belong to some of these classes. In addition to the above classes of principles, we consider variable-population extensions of maximin and leximin as well as principles that have been suggested by Carlson (1998) and Sider (1991). Members of the critical-level generalized utilitarian class have critical levels that are independent of utilities and population size. In the number-sensitive class, critical levels may depend on the number of people alive but not on their utilities.

Because the above principles provide complete orderings of alternatives, every pair of alternatives is ranked. In Chapter 7, we study a class of principles that rank some, but not all, pairs. It is closely related to the critical-level generalized utilitarian class and is called critical-band utilitarianism. The band is an interval of critical levels and each principle in the class ranks one alternative as at least as good as another if and only it is at least as good according to critical-level generalized utilitarianism for all critical levels in the band.

The book can be read as a defense of various classes of principles. The strongest case is made for the number-sensitive generalized utilitarian principles with nonnegative critical levels that are positive above some population size. An important part of the case for those principles is that their generalizations to principles that rank uncertain alternatives satisfy the most important axioms for that environment (see Chapter 7). A consistency axiom has the effect of ruling out all restricted principles. Additional axioms, which apply to environments with certainty (Chapter 5) or uncertainty (Chapter 7), provide support for various subclasses: the critical-level generalized utilitarian principles with positive critical levels, the number-sensitive critical-level utilitarian class, and the critical-level utilitarian principles with positive critical levels.

A good case, without considerations of uncertainty, can be made for the critical-level generalized utilitarian principles. If critical levels depend on population size, ethical significance must be attached to absolute numbers. If the axiom existence independence is employed, this problem vanishes because critical levels, if they exist, must be constant. Alternatively, the axiom extended replication invariance requires only relative population sizes to matter for social evaluation, and the only number-sensitive principles that satisfy it have constant critical levels.

In our view, the best class of principles consists of the critical-level utilitarian principles with positive critical levels. Our best defense of same-number utilitarianism (rather than same-number generalized utilitarianism) rests on acceptance of a property formulated in a model that accommodates uncertain consequences (Chapter 7). The crucial feature of the argument we use is the hypothesis that actual utility is given by one of the functions that can be used to compare uncertain outcomes. If this hypothesis is replaced by a condition that requires individual attitudes toward uncertainty about well-being to coincide across individuals and with the social attitude toward inequality of well-being, it is possible to make a case for critical-level generalized utilitarian principles that have positive critical levels and are weakly inequality averse.

To use a critical-level principle, the critical level must be chosen. The critical level must be positive to avoid the repugnant conclusion, but it must not be so large as to exclude lives that are reasonably good. For that reason, the critical-band principles may be an acceptable alternative. However, their application requires abandoning the requirement that all alternatives can be compared.

Although we offer defenses of these principles, we attempt to do more. The axiomatic investigation, if done well, should provide an understanding of the properties that are necessary for other classes, such as the restricted critical-level utilitarian and restricted number-dampened utilitarian classes. In addition, we expect our analysis will interest readers who are not welfarists, because they typically believe that considerations of well-being have some relevance for social evaluation.

1.2 ORGANIZATION

Each chapter after this one is divided into two parts, with mathematically advanced material relegated to Part B. Chapter 6 offers characterizations of the principles that are presented and discussed in Chapter 5. As a consequence, some results discussed in Chapter 5 are proved in Chapter 6.

Because the measurability and interpersonal comparability of well-being is a fundamental issue in our framework, we begin in Chapter 2 with an examination of the most important informational bases of individual well-being. If the number of alternatives is finite or countably infinite, there is a utility function for each person that ranks alternatives in the same way that the individual's goodness relation does. If there is an uncountable infinity of alternatives, goodness relations that satisfy a continuity axiom also have representations. These representations are ordinally measurable: unique up to increasing transformations.

An individual might be able to make comparisons such as "the gain to me in moving from y to x is greater than the gain in moving from z to w." If individual utility functions are cardinally measurable – unique up to increasing affine transformations – such comparisons are meaningful. In some environments, the ability to compare utility differences is sufficient to ensure cardinal measurability.

The measurement of well-being by means of utility functions is not, by itself, sufficient for interpersonal comparisons of well-being. Without such comparisons, however, there are no social-evaluation functionals based on ordinally or cardinally measurable individual utility functions that satisfy a few basic axioms (Arrow 1951, 1963; Blackorby, Donaldson, and Weymark 1984). If, however, individual utility functions are cardinally measurable, interpersonal comparisons of well-being at two (or more) levels, which we call norms, are sufficient to provide enough information for any social-evaluation functional. Temperature measurement, which employs a cardinal scale, is analogous. If the freezing and boiling temperatures for water are set at $0°$ and $100°$, the Celsius scale is defined. Once the norms have been chosen, temperature is numerically measurable, and comparisons of temperature levels and temperature differences at different times and places are meaningful.

In addition to the measurement of lifetime well-being, we also investigate the measurement of well-being in periods of time. Because people have integrated lives, such measures may depend on experiences in other periods. Lifetime well-being may be thought of as functionally dependent on levels of well-being in all the periods of life. We take account, explicitly, of the possibility of different lengths of life in different alternatives.

In Chapter 3, we turn to welfarist social evaluation. In both the single-profile and multiprofile environments, welfarism depends critically on the Pareto-indifference axiom. This axiom can be motivated by the requirement that the principle satisfy minimal individual goodness (Goodin 1991), which implies Pareto indifference. An important aspect of the population principles we investigate is that they are impartial. The anonymity axioms that we employ guarantee that, and they take account of differences in nonwelfare characteristics.

Each population principle has subprinciples that can be applied to fixed populations and they are presented and discussed in Chapter 4. Such subprinciples are a necessary part of population principles but, by themselves, are not capable of ranking alternatives with different populations. Although the fixed-population utilitarian principle is indifferent to inequality of well-being, other principles do exhibit aversion to utility inequality. Among these are the generalized utilitarian principles that use strictly concave transformations. They give priority to the interests of people whose well-being is low (Parfit 1997). There are many other inequality-averse principles, such as the social-evaluation ordering corresponding to the Gini index of inequality.

Fixed-population principles can be extended to comparisons of alternatives with the same population size (but possibly different identities of those alive) by using an anonymity axiom. Inequality-averse fixed-population principles can be used to generate ethical indexes of inequality of well-being for each population size. In Chapter 4, the best-known ways of doing this are discussed. An important axiom in this environment is called replication invariance. It requires inequality to be unchanged if a population is replicated. Thus, inequality in a population of two with utility levels of 20 and 40 is the same as inequality in a population of four with utility levels of 20, 20, 40, and 40. An important consequence of this axiom is that, if the same-number principle for each population

size is generalized utilitarian, a single transformation may be chosen for all population sizes.

In Chapters 5 and 6, we turn to population principles themselves. Although the average-utilitarian and classical-utilitarian principles have unfortunate properties, that is not true of all principles. All the classes of principles are examined in light of a set of axioms in an attempt to discover which are ethically attractive and which are not. Because some of the principles lead to the repugnant conclusion or declare the ceteris paribus addition of an individual who is below neutrality to be good, some of the classes can be eliminated easily. Of those that remain, some, such as the critical-level generalized utilitarian principles with positive critical levels, satisfy an axiom called existence independence. That axiom requires the ranking of any two alternatives to be independent of the utilities and number of people whose utility levels are the same in both. A weaker axiom called utility independence, which allows rankings to depend on the number but not the utility levels of those whose utilities are the same in two alternatives, is satisfied by the number-sensitive generalized utilitarian principles.

Critical-level generalized utilitarian principles with positive critical levels fail to satisfy an axiom that we call priority for lives worth living, which requires all alternatives in which each person is above neutrality to be ranked as better than all those in which each person is below neutrality. That axiom is satisfied by all restricted principles, but neither existence independence nor utility independence is satisfied. In Chapter 5, we find a set of axioms that the members of each class or subclass of principles satisfies. In addition, we explore the informational requirements of the various classes. We examine several examples to illustrate various principles at work.

In Chapter 6, we provide some characterization results and an important impossibility theorem, which demonstrates that there are no principles that satisfy all the axioms that we consider. In our view, the two chapters together provide a defense of the critical-level generalized utilitarian principles with positive critical levels.

Welfarist principles can be used to rank actions if they lead with certainty to particular alternatives. Because the consequences of actions are often uncertain, however, we consider uncertain alternatives in Chapter 7. Probabilities are assumed to be best-information probabilities and are used for individuals and society alike. Because principles that are capable of ranking such alternatives can also deal with the case of certainty, such principles are extensions of ordinary population principles.

Uncertain alternatives can be ranked from a social or individual point of view. We follow Harsanyi (1955, 1977) and require both social and individual rankings to be consistent with the same standard of rationality. In that case, when some additional axioms are used, we show that same-number subprinciples must be utilitarian. If a consistency requirement for critical levels is added, all restricted principles are ruled out and only the number-sensitive critical-level utilitarian principles remain. If, instead, individual preferences are required to reflect a single attitude toward utility uncertainty, social rankings must be made with a number-sensitive critical-level generalized utilitarian principle in which

social inequality aversion matches each individual attitude. This suggests that all the restricted principles should be rejected.

Chapter 7 also contains a characterization of the critical-band generalized utilitarian class of principles. It employs a formalization of a suggestion of Parfit (1976, 1982, 1984) concerning additions of individuals to a utility-unaffected population.

Because people have different birth dates and life spans, an investigation of population principles would be incomplete without an explicit consideration of those facts. Indeed, natural justifications for independence axioms involve populations that are long dead, an observation that strengthens the case for an intertemporal model. In Chapters 8 and 9, we employ an intertemporal framework in which nonwelfare information includes dates of birth and lengths of life. In it, we first use the Pareto-indifference axiom and an anonymity axiom that is, as in Chapter 3, conditional on nonwelfare information. The axiom independence of the existence of the dead requires the ranking of any two alternatives to be independent of the existence of those whose lives are over. The effect of Pareto indifference is to rule out discounting, and the social-evaluation problem becomes equivalent to the timeless problem of Chapters 5 and 6, with existence independence replaced by independence of the existence of the dead. Consequently, a characterization of the critical-level generalized utilitarian class is proved.

If Pareto indifference is replaced with the weaker axiom conditional Pareto indifference, whose application is conditional on birth dates and lengths of life being the same in two alternatives, discounting or upcounting based on people's birth dates is permitted. In addition, lengths of life may matter in social evaluation. When this weaker Pareto-indifference axiom is combined with some standard requirements and birth-date-conditional strong Pareto, which rules out the influence of lengths of life, a very general kind of critical-level principle with discounting results. The stationarity axiom requires only relative dates to matter, and its effect is to make critical levels constant and independent of the birth date of the added person. In addition, with the same axioms as in the characterization of the intertemporal critical-level generalized utilitarian class in the earlier part of the chapter, a class of principles called geometric birth-date-dependent critical-level generalized utilitarianism is characterized. A few arguments against discounting are provided in the chapter, but we find the fact that principles that discount necessarily violate Pareto indifference the most convincing.

Our model allows birth dates to be different in different alternatives. It might be argued, however, that birth dates are linked to identity and, therefore, fixed. Consequently, we address that issue directly and sketch a model that yields the same results.

Chapter 9 continues the exploration of the dynamic environment introduced in Chapter 8, and we ask whether it is possible to apply population principles to a single period in a consistent way. If the information about two alternatives differs in one period only, we require their ranking to be independent of information