Introduction

Many decisions of public concern cannot be left to the market because cooperative opportunities will not be efficiently utilized by decentralized actions of the agents. The most prominent examples include the provision of public goods, pricing of a natural monopoly, as well as all decisions taken by vote. To remedy these market failures, welfare economists have come up with a variety of normative solutions and tried to convince the decision makers of their relevance.

The theoretical foundation of these normative arguments is axiomatic. This point was clearly made in Amartya Sen's landmark book (*Collective Choice and Social Welfare,* first published in 1970). Since then the axiomatic literature has considerably expanded its scope and refined its methods: The whole theory of cooperative games has played a central role in the analysis of cost sharing when there are increasing returns to scales; our understanding of voting rules now encompasses the impact of strategic manipulations; several refined measurements of inequality have been constructed and abundantly tested; and so on.

This book describes the recent successes of the axiomatic method in four areas: welfarism (the construction of collective utility functions and inequality measures as well as the axiomatic bargaining approach); cooperative games (the core and the two most popular value operators – Shapley value and nucleolus); public decision making (cost sharing of a public good and pricing of a regulated monopoly, in both the first-best and strategic-second-best perspectives); and voting and social choice (majority voting à la Condorcet and scoring methods à la Borda; the impossibility of aggregating individual preferences into a social preference).

I attempt to provide a comprehensive and unified presentation of these technically heterogeneous subjects. The link between them is the collection of axioms (in total, about 30 of them are discussed). Many of them are so versatile that they apply, with very little modification, to radically different problems. A good example is the core property that profoundly

Table I.1. *Interprofile axioms*

Independence[a]	Monotonicity[a]
Anonymity: all chapters	
Neutrality: 8, 9, 11	
Variable preferences	
Arrow IIA: 11	Pigou–Dalton principle: 2
Zero independence: 2, 5	Monotonicity: 9, 11
Scale independence: 2, 3	Strategyproofness: 8, 10
Independence of common utility pace: 2	Strong monotonicity: 10
Independence of common zero: 2	
Independence of common scale: 2	
Marginalist (value operator): 5	
Decentralizability: 6	
Variable issues	
Nash IIA: 3, 11	Issue monotonicity: 3
Aizerman: 11	Coalition monotonicity: 5
Expansion: 11	Technological monotonicity: 7
Path independence: 11	
Additivity: 3, 5, 6	
Variable population	
Separability: 2, 3	Population monotonicity: 3
Reduced-game properties: 5	Participation: 9
Consistency: 6	Reinforcement: 9

[a] Numbers refer to chapters.

influences the discussion of surplus-sharing methods in Chapters 6 and 7 as well as that of strategic voting in Chapter 10. Or consider the separability axiom, which plays an important role in the analysis of collective utility functions (Chapters 2 and 3), in values of cooperative games (Chapter 5), as well as in cost- and surplus-sharing methods (Chapter 6). Many more such connections exist, as will become clear from the classification outlined in Table I.1.

Throughout the book we utilize the axiomatic method to guide decision making. Of course, each microeconomic problem that we will examine has more than one plausible solution: There are many "good" voting rules, several plausible "values" for cooperative games, many useful inequality indices, and so on. Ideally, the axiomatic method can help our choice by, first, reducing the number of plausible solutions as much as possible and, second, providing us with a specific axiomatic characterization of each of these plausible solutions. For instance, we find essentially two interesting value operators for cooperative games and provide a characterization of each (Part II, Chapter 5). See also the discussion of collective utility functions in Part I (Chapter 2), built on the two polar examples of egalitarianism and classical utilitarianism; or, in Part IV, see the two focal families of voting rules based, respectively, on scoring systems and on majority comparisons.

As a rule, we pay little attention to the innumerable impossibility results that the axiomatic approach generates too easily. The only exceptions are the impossibility theorems of Arrow (Chapter 11) and Gibbard–Satterthwaite (Chapter 10), of which the mathematical content is so rich.

We cover many different economic models where the axiomatic method has scored significant successes. Our choice of topics is, however, not exhaustive. One notable missing subject is the allocation of private goods, whether through exchange or as a pure distribution issue (cake division problems). Despite a substantial normative literature on these questions, the existing results (surveyed in Thomson and Varian [1985]) do not appear sufficiently focused upon particular solutions to justify a textbook presentation. The only (major) exception is the "Edgeworth proposition," characterizing the competitive equilibrium by means of the core property when the agents become vanishingly small. But the allocation of private goods between a small, fixed number of agents has not yet found clear-cut axiomatic answers except in some very special cases such as matching problems.

We now propose a categorization of the many axioms populating this book. The one and only axiom that pervades through the entire book is Pareto optimality (also called unanimity or efficiency). This should come as no surprise since 40 years ago the new welfare economists agreed that it was the only indisputable principle on which legitimate welfare analysis could ever develop. Note that Pareto optimality is an *intraprofile* property, in the sense that it can be defined for one single preference profile for one particular problem. The only other intraprofile axioms are those placing lower or upper bounds (mostly lower) on the welfare of individual agents and/or coalitions of agents. The main examples are the individual rationality and core properties (Chapters 4–8 and 10).

The *interprofile* axioms consider a specific change in the parameters of the model and state some condition on the induced change of the solution. If the condition says that the solution must not change as the parameter changes, we have an *independence* axiom; if it says that the solution must shift in a certain direction related to the shift of the parameter, we have a *monotonicity* axiom.

The most popular independence axiom is anonymity. It says that the solution should treat agents equally (one person, one vote): Specifically, when the problem is altered by exchanging the characteristics of two agents (including their preferences), its solution should not change. This basic equity principle appears explicitly in most chapters and is implicit in all. A related equity property for voting rules is neutrality, requiring that no candidate should be a priori discriminated against. Neutrality plays an important role in Part IV.

The bulk of the interprofile axioms falls into six classes corresponding to the variations of different parameters of the model. The parameters can be (a) the agents' preferences, (b) the set of feasible outcomes (the issue), or (c) the set of concerned agents (the population). In each case the axiom can be either an independence or a monotonicity property, so that, as stated, we have six classes of axioms.

Variable preferences

This is by far the most frequent parameter variation since Arrow's pioneer discussion of the independence of irrelevant alternatives (IIA; an independence axiom) as well as of positive responsiveness (a monotonicity axiom). In the welfarist theory (Part I), Sen [1977] pointed to the paramount importance of independence axioms that economize on information gathering. Indeed, these axioms are all over Part I (see the independence of individual utility scale or common zero of utility). In the analysis of public decision making, monotonicity of the solution with respect to (w.r.t.) preferences is the key to incentive compatibility (see the equivalence between strategyproofness and strong monotonicity in Chapter 10).

Variable issues

The issue monotonicity property in axiomatic bargaining (Chapter 3) says that when the set of feasible outcomes expands, the welfare of no agent should decrease. A similar axiom applies to production economies (Chapter 7) and values of cooperative games (Chapter 5).

On the other hand, abstract choice functions (Chapter 11) are classified by means of independence axioms where the variable parameter is the issue. See Nash's version of the IIA (in Chapters 3 and 11).

Variable population

The prominent independence axiom is separability. A collective utility function is separable if it can compare two utility distributions restricted to a subgroup of agents independently of the rest of the utility distribution. Separability in various forms is also defined for inequality indices (Chapter 2), axiomatic bargaining solutions (Chapter 3), cooperative games (see the reduced-game properties in Chapter 5), and surplus-sharing problems (see the consistency axiom in Chapter 6). In each of these contexts, it drives a very powerful functional equation.

Several population monotonicity properties play an important role in voting (Chapter 9) and in axiomatic bargaining (Chapter 3).

Table I.1 gives the detail of the foregoing classification.

Overview

Two broad categories of models are considered: sharing divisible surplus (Parts I, II, and III) and collective choice of an indivisible public decision (Part IV).

Part I presents the *welfarist* theory of utilitarianism. The central postulate is that only the individual utility levels matter when comparing any two outcomes, and those levels are intercomparable across agents. The main concept is that of a social welfare ordering (or the almost equivalent concept of a collective utility function) aggregating individual utilities into a social preference. The two principal examples are the classical utilitarian collective utility (the sum of individual utilities) and the egalitarian collective utility (the minimal individual utility), which are contrasted in Chapter 1. The next step is to construct collective utility functions that compromise between these two, thereby generating a family of reasonable indices to measure inequality (Chapter 2). Finally, the model of axiomatic bargaining generalizes the welfarist approach by making the feasibility constraint an ingredient of the choice method itself (Chapter 3).

The main limitation of welfarism is that it deliberately ignores ethical considerations suggested by the very nature of the economic decisions at stake. One such consideration is the protection of cooperative opportunities open to individual and to coalitions of agents. The theory of *cooperative games,* to which Part II is devoted, extends the welfarist model by taking those opportunities into account.

The stand-alone test is a general equity principle for cost sharing, translated in Chapter 4 as the *core* property. It prevents a coalition from getting a surplus share below its own cooperative opportunities. Although in most of Part II we make the simplifying assumption that money is available to transfer utility across agents, the search for a value solution (a deterministic surplus sharing defined for all cooperative games) is anything but simple. In Chapter 5 we contrast the two main solutions, one of them picking a central point of the core (nucleolus) and the other assigning to an agent the average of his marginal contribution to various coalitions (Shapley value).

The economic applications of cooperative games are very diverse. In Part III we discuss two especially important ones in the more general perspective of public decision mechanisms. These are the pricing of a regulated monopoly and the provision and cost sharing of a public good.

We start with two deceptively simple binary decision problems in Chapter 6: the cost sharing of an indivisible public good and the surplus sharing of an indivisible cooperative venture. Guided by the core property, we find no less than five reasonable solutions for the cost-sharing problem only. In Chapter 7 we consider in full generality a one-input, one-output regulated monopoly, producing either a public or a private good. We find the traditional marginal pricing methods challenged by two simple welfare egalitarian solutions.

In Chapter 8 we discuss the incentive properties of public good mechanisms. In general, a mechanism is called incentive compatible if its outcome is robust against strategic manipulations of selfish agents. The easiest manipulation consists of misreporting one's preference. A mechanism is strategyproof if such misreport does not occur. The pivotal mechanism was invented in the early seventies as an example of strategyproof mechanism. We discuss its incentive and axiomatic properties and compare it with other strategyproof mechanisms. We also look at the incentive properties of some familiar first-best mechanisms.

Part IV studies the *voting* problem: A pure public decision must be selected from the conflicting opinions of a given set of agents. The key difference with all earlier models is the absence of some *numéraire,* a transferable good allowing compensation for those agents who dislike the chosen decision at the expense of those who like it. In voting, we typically have finitely many outcomes and purely ordinal preferences only. In spite of this, the axioms for voting rules are fairly similar to those used in earlier parts.

The two-centuries-old debate of Condorcet versus Borda inspires the discussion of Chapter 9. We contrast voting rules based on majority comparisons with methods assigning scores to candidates from each voter's ballot. In Chapter 10 we discuss the two main incentive properties, strategyproofness (it turns out that no reasonable rule can be strategyproof if at least three candidates compete) and the core property. The latter leads to the minority principle, attentive to preserve some decision power for all coalitions however small. Finally, in Chapter 11 we survey the main results of Kenneth Arrow's aggregation-of-preferences approach. As in Part I in the cardinal context, one seeks to derive from the profile of individual preferences a collective ordering of all candidates. Despite many impossibility results limiting the scope for practical applications, this line of research provides useful insights into the logics of collective action.

We provide an average of nine exercises at the end of each chapter, many of them establishing some fine points alluded to in the course of the chapter; the exercises are also used to draw more links across chapters with the help of "continued" examples.

The complete bibliography is gathered at the end of the volume. The literature provides some good surveys to complement many of our chapters. We list them here for the benefit of the teaching-oriented reader.

Chapters 1 and 2: The classical books of Sen [1970] and Kolm [1972] can be complemented by the survey of d'Aspremont [1985]. On inequality indices, see Shorrocks [1985] and Foster [1985].

Chapter 3: The monograph by Roth [1979] is slightly outdated and should be supplemented by the recent survey of Kalaï [1985] and the forthcoming book by Thomson. The latter covers many quite recent results and all classical ones.

Chapters 4 and 5: The game theory textbooks typically do not cover the cost allocation applications; the best references seem to be the article by Young [1985b] for the transferable utility (TU) games and the monograph by Sharkey [1982] for the applications to the pricing problem; see also Ichiishi [1983] for a thorough exposition of nontransferable utility (NTU) games.

Chapters 7 and 8: A good elementary introduction is in the books of Feldman [1980] and Mueller [1979]; the monograph by Green and Laffont [1979] gives a more advanced account of the demand-revealing mechanisms (Chapter 8).

Chapter 9: The small book of Straffin [1980] is a remarkable introduction to voting rules (it also discusses the pivotal mechanism). See also the survey by Moulin [1985a]. The book by Mueller [1979] gives a detailed account of the applications of the voting model to public choice.

Chapter 10: A good introduction is in Feldman [1980]; for a systematical presentation of strategic voting, see the books of Peleg [1984a] and Moulin [1983].

Chapter 11: Expositions of the aggregation of preferences approach are easy to find in the literature; for instance, among the preceding texts cited see Sen [1970], Feldman [1980], Mueller [1979], Moulin [1983], and Peleg [1984a]. See also Suzumura [1983].

PART I

Welfarism

CHAPTER 1

Egalitarianism versus utilitarianism

Overview

Utilitarianism is a philosophical thesis two centuries old. It judges collective action on the basis of the utility levels enjoyed by the individual agents and of those levels only. This is literally justice by the ends rather than by the means. Welfarism is the name, coined by Amartya Sen, of the theoretical formulation of utilitarianism, especially useful in economic theory and other social sciences. Its axiomatic presentation, developed in the last three decades, is the subject of Chapters 2 and 3.

For the utilitarianist, social cooperation is good only inasmuch as it improves upon the welfare of individual members of society. The means of cooperation (social and legal institutions, such as private contracts and public firms) do not carry any ethical value; they are merely technical devices – some admittedly more efficient than others – to promote individual welfares. For instance, protecting certain rights – say, freedom of speech – is not a moral imperative; it should be enforced only if the agents derive enough utility from it.

This very dry social model rests entirely upon the concept of individual preferences determined by the agents *"libre-arbitre"* while deliberately ignoring all its mitigating factors (among them education and the shaping of individual opinion by the social environment, as well as kinship, friendship, or any specific interpersonal relation). Ever since Bentham, utilitarianists have been aware of these limitations. Yet, oversimplified as it is, the utilitarian model is easily applicable, and its "liberal" ideology has the force of simplicity. The philosophical debate on utilitarianism is still quite active (see, e.g., Sen and Williams [1982]).

Axiomatic welfarism idealizes a collective decision problem by attaching to each feasible alternative (to each possible decision) the vector $(u_1, \dots, u_n)$ of individual utility levels, where u_i is agent i's utility. All relevant information is contained in the set of those feasible utility vectors.

Any information about specific decisions implementing various utility vectors is systematically erased.

Given the set of feasible utility vectors, the collective decision results from a mathematical and deterministic rule that pinpoints one vector as society's choice. This rule conveys the whole ethical policy of the society under consideration. No wonder that its definition raises fierce arguments: It is nothing less than an exhaustive solution of the social problem, a formula computing from the "arithmetics of pleasures and pains" the final outcome that society should enforce. In this chapter we discuss the two main rules advocated by utilitarianists, namely, egalitarianism (seeking to equalize individual utilities) and classical utilitarianism (maximizing the sum of individual utilities).

Egalitarianism is discussed first (Section 1.1). It derives from the oldest and most popular principle of justice: Equal agents must be treated equally. When utility is the only variable of control, application of the principle amounts to equalizing individual utilities. Yet, simple as this may appear, it may conflict with another basic postulate of collective decision making, known as the unanimity principle.

Whenever every concerned agent prefers decision a to decision b, the unanimity principle rejects decision b (we say that b is Pareto inferior to a). Unanimity is the single most important concept of welfare economics and the only axiom that will play a role in each and every chapter of this book. In the welfarist model, unanimity says that the chosen utility vector must be Pareto optimal (that is to say, it should not be Pareto inferior to any feasible utility vector).

The unanimity principle may conflict with the plain equalization of individual utilities. This somehow counterintuitive fact is known as the equality–efficiency dilemma. Its solution takes a more careful definition of the egalitarian program as the maximization of the leximin preordering (see Section 1.1 for details).

Egalitarianism is a very sturdy social glue. When all agents share equally the benefits of cooperation, there is no room for envy or frustration (except, of course, when some agents feel that their contribution to those benefits has been above average, but that is beyond the scope of the welfarist model; see the models of Chapters 6 and 7). The distributive consequences of egalitarianism, however, are sometimes hard to accept: For the sake of bringing one more unit of utility to a single agent, it may forfeit vast amounts of aggregate utility. Think of the distribution of medical care among patients who have subscribed to the same insurance. Say that the announced objective is to provide them with a level of health as equal as possible. This may mean that we must deny forever aspirin and antibiotics to all agents but one in order to pay for expensive equipment that