PART I

# Scope and limitations

CHAPTER 1

# Introduction

We present here a synthesis of recent developments in normative welfare economics that have direct application in public finance. Within this rather broad subject, we will emphasize operational prescriptions for public policy in a "mixed" economy consisting of a public sector, a private-market sector, and some limited links and controls between them.

Since the book is mostly about normative economics, we start with a discussion of social objectives. Our approach is that of the "old welfare economics," and since the term "old" may connote "outdated" to some, we argue the point at some length. In the process, we air our views on the methodology of interpersonal comparisons and distributional equity. Having settled on a normative methodology, we use it to justify our conception of the mixed economy. Chapters 2 and 3 begin this development squarely in the tradition of competitive theory. We outline reasons why "direct" decision making is inadequate and summarize the case for market decentralization.

Staying in this tradition, we define collective goods by intrinsic characteristics that make them unmarketable. We develop a conceptual framework for studying such goods that is general enough to encompass many common examples such as natural monopolies, pure public goods [as defined by Samuelson (1954)], and "club" goods [as defined by Buchanan (1965)]. Although all these examples share an important common element, distinctions among them provide foundations for a normative theory of "fiscal federalism"; pure public goods are best allocated by a national entity, whereas club goods are best allocated by a more local entity. The ways in which these entities ought to do business and the relationships between them are explored in Chapters 5–7.

Initially, we search for planning mechanisms that simultaneously elicit necessary information and implement first-best decisions for collective-goods allocation much as the market system does for private goods. Although we can identify many well-known attempts along these lines, almost all of them suffer from incentive problems: Agents find that it is not in their interests to behave according to the specified rules. We end up rejecting planning mechanisms as first-best procedures in most cases but find some of the examples useful as building blocks for less ambitious

approaches to come later and in identifying biases that are likely to emerge when first-best planning is not possible.

Among the more controversial issues discussed here are the proposition [put forward by Tiebout (1956)] that less interference is required in the allocation of club goods (by local constituencies) than in the allocation of pure public goods and the view [argued by Bradford and Oates (1971)] that general revenue sharing (by the national government to the local governments) is irrelevant in the presence of a national tax system that correctly deals with equity considerations.

Part II also includes some general discussion of optimal public fund-raising. The question of tax versus debt finance is raised in Chapter 6, whereas issues involving optimal public pricing and indirect taxation are dealt with in Chapter 8. We explore and evaluate recent incarnations of the "Ricardian equivalence" doctrine, which dictates that the choice of tax versus debt finance is irrelevant. Although we do not accept strong versions of this doctrine, we do generally accept the view that issues involving the choice of debt can be reduced to issues involving the timing of taxation.

When discussing optimal taxation, we focus on the parallels between the rules of Ramsey public pricing and the rules of optimal-commodity taxation. We argue that these rules are useful and operational only if certain simple structural assumptions can be invoked; these assumptions seem reasonable in some applications of the Ramsey analysis but much less so for the optimal-commodity tax problem. Indeed, we suggest that most common interpretations of the optimal-commodity-tax rules are misleading at best.

Many people have argued that the "best" tax system we could have would be one that is broadly based in that most goods would be taxed at the same rate. We ask whether such claims could be justified on the basis of optimal-tax theory and discuss some of the practical considerations involved in constructing such a tax.

Beginning in Part III our general focus changes considerably. To that point, we have aimed at a comprehensive treatment of public issues; that is, we have sought institutional arrangements and planning procedures that can deal with the full range of public decision making simultaneously. However, from a pragmatic point of view, such aims turn out to be too ambitious, requiring, among other things, too much centralized information processing. At the very least, we require some political decentralization so that separate issues can be dealt with by semiautonomous agencies. *Benefit–cost analysis* provides a theoretical framework for implementing this decentralization. Projects (or programs) are evaluated one at a time, and various measures are developed to account for general equilibrium interactions.

Part III examines benefit–cost analysis as applied to projects or programs that are sufficiently small. The meaning of "small" here is roughly that implementation of the project will not change consumer marginal rates of substitution (a precise definition is given in the text). This restriction is quite useful in reducing information-gathering requirements, as we shall see. Although we couch the discussion in terms of evaluating small projects, the analysis can also be used to determine and evaluate first-order conditions for optimal allocations or to point the direction of large movements toward full optimality. However, some caution is required in adopting these interpretations. First, when projects are characterized by important elements of scale economies, the analysis of small versions may tell us very little about the desirability of large ones. For this reason, it is important to have a different methodology for evaluating large projects.

Second, even if all technologies are characterized by full convexity, it is still possible for first-order analysis to be a misleading guide to the optimal direction of movement. A germaine example has been provided by Diamond and Mirrlees (1973), who show that the optimal level of an external diseconomy can be higher than it would be in the absence of regulation (first-order analysis applied to an unregulated status quo would naturally dictate lower levels of the externality). This example and others like it depend heavily on the presence of strong income effects (as the inefficiency is removed, consumers are "better off," causing them to care much more about the good responsible for generating the externality). Generally, we think such examples are pathological since it is hard to believe that income effects are going to be so potent. It is partly for this reason that we place such a strong emphasis on first-order analysis in this book.

We identify two distinct, though related, methodologies associated with first-order benefit–cost analysis. The first of these values directs project inputs and outputs at market prices (when possible) and accounts for interactions between one project and the outside environment by adding (or subtracting) a sequence of correction terms. The other seeks to administer the corresponding corrections directly on valuations of the inputs and outputs.

The first method relates closely to the general theory of second best, since each correction term is associated with a particular second-best distortion present in the underlying mixed economy. For example, there will be a term associated with the presence of indirect taxation that can be evaluated using information about tax rates and tax base elasticities. Aside from the practical advantages of such decompositions, the approach provides a useful way to establish the optimality of market decentralization when there are no second-best distortions. We develop the relevant

decompositions in Chapter 9 and use them to generate important operational rules for project analysis.

The second approach is identified with methods of *shadow pricing*. Namely, we seek to measure the marginal social benefit for particular project outputs and the marginal opportunity cost of project inputs and aggregate these to get a bottom line. Although the two approaches naturally must lead to the same bottom line, they involve quite different intermediate calculations and illustrate distinct theoretical principles. We develop the theory of shadow pricing in Chapter 10 and show how it can be used to illustrate and amplify on a number of theoretical propositions in the project analysis–development planning literature.

Chapters 11 and 12 apply the methodology of the previous two chapters to the study of problems involving spatial and temporal allocations. Applying the decomposition principles to problems in local community choice enables us to identify and measure the interactions between one community and the rest of the economy. These measures enable us to quantify the concept of *fiscal externality* and to identify several potential biases in local choice with respect to size and composition of local public-goods provision. We present a unified treatment that illustrates the relationships between various types of potential distortion.

Most discussions of project analysis in an intertemporal context focus attention on the social rate of discount. We apply the shadow-pricing theory of Chapter 10 to evaluate this rate in the mixed-economy environments of this book. Our analysis explains why there has been such a divergence of opinion concerning the correct size of the social discount rate, since the answer turns out to be quite sensitive to crucial elasticity assumptions.

When evaluating the impact of risk in social decision making, the decomposition approach again turns out to be most useful. We show how discount rates can be decomposed into a term reflecting pure impatience, a term associated with uninsurable risk, and a term associated with failure to allocate risk in an optimal way. We find that the second of these terms is closely related to risk measures found in the capital asset-pricing model, even though the latter is based on a much simpler structure. Moreover, we see that the third term can explain a number of examples of "pecuniary externality" that have appeared in the recent literature on incomplete markets.

Chapter 13 concludes Part III with a discussion of measurement techniques for identifying shadow prices. Since shadow pricing becomes problematic precisely when market prices misrepresent opportunity value, we must draw on methods outside of standard market econometrics. We focus on indirect methods such as using individual choice between life-

threatening versus safe jobs to infer something about the shadow price of life or using land rent differentials to infer the shadow price of air pollution. We argue that procedures of this type have a common underlying theoretical structure (we label them *spanning methods*) and that the limitations of this structure are rather severe. The ensuing analysis is used to discuss and evaluate some commonly used hedonic methods.

First-order project analysis represents a relatively well developed methodology. And although it does present some measurement difficulties (as we discover in Chapter 13), this methodology constitutes a useful practical framework for analyzing and evaluating public programs. Unfortunately, however, a significant class of public projects cannot be considered small in the sense required by first-order analysis. These are projects that must be instituted (if at all) at a sufficiently large scale so as to change the relative valuation of commodities they use or produce. Part IV of this book looks at the state of the art in evaluating such projects.

Chapter 14 traces the development of exact methods for measuring net welfare of large projects. Although it is possible to construct such measures from sufficiently rich information about choices among sets of alternatives, we find that practical difficulties with implementation are far more formidable than was the case in Part II. Therefore, we devote the following chapter to a discussion of possible approximation methods.

Chapter 16 looks at ways of estimating the theoretical measures developed in the previous two chapters and asks how such estimates can be used to solve practical problems in benefit–cost analysis. Although we do not have much to say about actual econometric technique, we do comment briefly on the strengths and weaknesses of various methods. Finally, we analyze in Chapter 17 a class of problems that incorporates certain "large-project" features as well as a number of other characteristics that have appeared prominently in the book. Although hardly a neat summary of what went before, this chapter does serve to bring together a number of our major themes.

CHAPTER 2

# Social objectives and direct decision making

Any book on normative economics ought to begin with some discussion of social objectives. Since the meat of this book involves practical methods for evaluating such objectives, and since the extant literature on social choice theory is voluminous, we choose not to spend a great deal of time on the subject here.[1] However, a short discussion of our basic philosophy is in order.

Fundamental to our approach is the implicit acceptance of "consumer sovereignty." Operationally, this principle means that only measures of individual preference affect social choice. We reject the paternalist view that government knows better than its citizens what is good for them. Having taken this position, it remains to decide how to measure individual preference and how to reconcile differences among these individual preferences.

We also accept the classical view of economic person as utility maximizer. If we think of preference as revealed by choice, the underlying choice functions must be consistent. Although violations of this principle of rationality can be found in the experimental literature, they seem to be exceptions rather than the rule, at least in a certainty context. Assumptions of rationality in an *uncertainty* context generate more serious objections, which we will comment on later. Unfortunately, we would not make much progress in this book if we tried to drop all principles of rationality unless we also dropped consumer sovereignty; the planner will have a hard time keeping hands off consumer preferences when they are inconsistent.

Operationally, we are going to represent individual preference on social alternatives by a utility function on these alternatives. We assume that these functions are at least twice continuously differentiable. Although such "smoothness" can be given an intuitively appealing justification, we adopt it here more for convenience than anything else; it enables us to conduct "first-order analysis" and to develop measures of welfare by integrating first-order measures. Without smoothness, first-order measures would always be approximations, and our analysis would be messy at best.

[1] See Sen (1970) or Mueller (1976) for an introduction to this literature.

So, we take smooth utility functions as the inputs into social decision making. Up to now, nothing we have done would be very controversial among economists, although we would certainly get some argument from other social scientists. However, when we specify a method of aggregation, the subject quickly becomes quite controversial. Disagreements arise as soon as we try to specify which properties of the utility functions can be used by the planner in the course of the aggregation. For example, can we use *cardinal* properties of these functions or must the derived social ordering be invariant to monotone transformations of individual utility functions? It has been argued that since only ordinal preference can be observed (at least in the certainty context), only ordinal information should be used. Unfortunately, it is difficult to formalize this restriction in a useful way. However, stronger restrictions in the same spirit can be formalized. For example, if we stipulate that social choice between any *pair* of alternatives depends only on individual ordinal rankings of these same two alternatives, we have Arrow's famous "independence from irrelevant alternatives" (IIA).[2] The added strength may be desirable per se since it guarantees that we will not be making any interpersonal comparisons of utility units in our social choice procedures.

But if we accept IIA, we will quickly find ourselves up against the bane of this literature: impossibility theorems. Indeed, Arrow showed that if we want a social welfare ordering that is universal (i.e., it can handle all potential individual preference configurations) and nondictatorial (i.e., it does not agree exactly with any particular individual preference) and satisfies the Pareto principle (to be discussed), it cannot satisfy IIA.[3] This is pretty unfortunate since we can think of quite a few other "reasonable" assumptions we would like to add before we would be completely satisfied with any resulting aggregation function. A voluminous literature now exists demonstrating that even if we are willing to drop IIA, we run into impossibilities as soon as a few other reasonable assumptions are added.[4]

The sad fact here is that we cannot have everything we want. Realizing this, there are several ways to go. We can forgo those desirable axioms on comparability, we can limit the set of issues to be resolved through a general social choice procedure, or we can impose some restrictions on the types of individual preference allowable as input to social choice. The first two approaches are associated with the labels "old welfare economics" and "new welfare economics," respectively. We compare and contrast these two first and take up the third later in the chapter.

[2] This condition is one of the axioms introduced in Arrow (1951).
[3] See Arrow (1950) for a proof.
[4] A number of these negative results are discussed in Sen (1970).

According to the new welfare economics, we are to restrict social decisions to those that can be made on the basis of the Pareto principle; that is, we should institute only those policies (or groups of policies) that make some people better off without making anyone else worse off. Stated in terms of utilities, the principle says we should never be satisfied with a vector of utilities that is interior to the set of attainable utility vectors; rather, we should always move toward the boundary in a direction of increasing (or at least nondecreasing) utilities for all. Clearly, as long as nothing more is said about the direction of movement, the rule is *ordinal* but does not fully resolve any existing conflict of interest.

Although the Pareto principle cannot be used to decide all issues, it does generate a well-defined decision procedure in the social contexts we shall consider. These are situations in which there is some status quo ante, and the government (or planner) is considering some change in it (call this a *project*). In the absence of specific transfers, such a project is likely to make some people (the *gainers*) better off but others (the *losers*) worse off. The Pareto principle becomes the "compensation principle with compensation"; a project is initiated if the gainers can (and *do*) compensate the losers in such a way that all of them are better off (or at least none of them are worse off).

Unfortunately, this principle has gained the reputation of being value free, meaning that it does not rely on any judgments about interpersonal comparisons. We argue that the reputation is richly *un*deserved. As soon as the principle is placed in a context where it has operational bite, value judgments are involved in its implementation. To see this, consider the status quo. If this is thought of as representing a position that is "infinitely bad" for everyone, the Pareto principle has no bite as applied to project analysis: Every project will look desirable. But as soon as a real status quo is envisioned, use of the compensation principle with compensation makes the implicit judgment that there is something good about the vector of utilities inherent in the status quo: Projects that improve on all components of *this* vector will be accepted, whereas alternative projects that lead to *other* Pareto-efficient utility vectors will be rejected.

We are not here to argue about whether there is justice in the status quo; there is a substantial "entitlements" literature extant that supports this view but also plenty of dissent.[5] Our point is simply that the compensation principle is not value free, at least not as usually applied. Given this, we prefer methods that bring the element of interpersonal comparison into the open rather than one that prejudges the distributional issue. However, we will find occasion in the sequel to compare and contrast

[5] The entitlements position is expounded by Nozick (1974); see Harsanyi (1955) and Rawls (1971) for some contrasting points of view.

our proposed methods to those involving the compensation principle both with and without compensation.

## 2.1 Interpersonal comparisons

How shall we make interpersonal comparisons given that we must? A natural answer (and the earliest one in the literature) is to postulate a welfare function that aggregates utility levels in the same way individual utilities aggregate individual consumption levels. Using such a function implies that in deciding between two different social states, we look only at individual utilities for those two states. Thus, this formulation is not simply a consequence of consumer sovereignty but requires in addition a "utility independence" assumption.

More generally, we could form a welfare measure that took utility *functionals* as arguments. Unfortunately, these mathematical objects are not easy to work with, and further, it is not clear what violations of utility independence we ought to introduce. So we will confine ourselves here to simple welfare functions.[6] What form should they take? Again, we could easily digress for several chapters discussing this question but instead will confine ourselves to a few remarks. An interesting framework for thinking about the question is the so-called original position. According to this view, the decision maker is someone with the universal traits of humanity but without personal identity; some like to think of this actor as "me before I know who I am."

Assuming that the decision maker is omniscient in the original position, he knows all the possible people he might be and their personal characteristics. The preferences he would exhibit in these circumstances become the social preferences. Thus, if we could agree on what "universal person" would do in these circumstances, we could agree on a social ordering. Unfortunately, the literature of the past 30 years has revealed that there is no agreement within the class of academics, let alone universal agreement.

Harsanyi (1955) and others looked at this problem as one of choice under uncertainty, with universal person having an equal chance of being each real person. Invoking standard axioms of choice under uncertainty, Harsanyi concluded that it was legitimate to treat individual utility functions as cardinal and further that "expected utility" should serve as the social welfare indicator. Assuming that the number of people is fixed, the derived formulation is equivalent to classical utilitarianism (summation of individual utilities).

[6] Utility independence is closely related to the axiom of "neutrality" in the social choice literature. See Sen (1970) for a discussion of this axiom and its shortcomings.