

WELFARE ECONOMICS

Robin W. Boadway &
Neil Bruce

Welfare Economics

ROBIN BOADWAY

NEIL BRUCE

Basil Blackwell

© Robin Boadway and Neil Bruce 1984

First published 1984

Basil Blackwell Publisher Limited

108 Cowley Road, Oxford OX4 1JF, England

Basil Blackwell Inc.

432 Park Avenue South, Suite 1505

New York, NY 10016, USA

All rights reserved. Except for the quotation of short passages for the purposes of criticism and review, no part of this publication may be reproduced, stored in a retrieval system, or transmitted, in any form or by any means, electronic, mechanical, photocopying, recording or otherwise, without the prior permission of the publisher.

Except in the United States of America, this book is sold subject to the condition that it shall not, by way of trade or otherwise, be lent, re-sold, hired out, or otherwise circulated without the publisher's prior consent in any form of binding or cover other than that in which it is published and without a similar condition including this condition being imposed on the subsequent purchaser.

British Library Cataloguing in Publication Data

Boadway, Robin W.

Welfare economics.

1. Welfare economics

I. Title II. Bruce, Neil

330.15'5 HB99.3

ISBN 0-631-13326-7

ISBN 0-631-13327-5 Pbk

Typeset by Unicus Graphics Ltd, Horsham, West Sussex
Printed in Great Britain by Bell & Bain Ltd, Glasgow

To Bernie and Janis

Preface

Most economists are prepared, even eager, to make prescriptions and recommendations about economic matters based on their specialized knowledge. This practice is so widespread that a catalogue of examples would easily fill this volume. It is a source of concern to many people that economists who are so eager to prescribe policies often disagree among themselves as to which policies should be followed, and even advocate conflicting policies. This is a legitimate concern, especially since the basis for their differing views is often obscure, not only to the interested listener, but even to the economists themselves.

In all cases where economists make policy prescriptions they are transcending the scientific territory of positive economics, which is predictive and descriptive, and are practising normative or welfare economics. Positive economics is an important element in welfare economics because it describes what is feasible and predicts the observable implications of particular policies. However, positive economics is not the only element in welfare economics. In order to conclude that a situation is 'good' or 'bad', or that policy x is more desirable than policy y , it is necessary to accept certain ethical premises. These premises may be strong or weak. They also may represent the peculiar ethics of the economist making policy statements or, more usefully, may represent the economist's view of an ethical consensus.

Inevitably, many economists practise welfare economics without distinguishing the ethical premises and the positive economics on which their conclusions rest. When they differ about positive economics there is hope that such differences can be resolved by careful empirical research. When they differ on ethical premises it is necessary to make one's premises explicit and champion them with an appeal to persuasion and consensus.

In this book we systematically develop the foundations of modern welfare economics. We particularly stress the twin pillars of positive economics and ethical premises, and examine both the weaker ethical premises underlying economic efficiency and the stronger premises underlying social welfare maximization. We believe a good understanding of welfare economics is an important ingredient in the constructive use of economics in policy-making. It also adds interest and zest to the study of economics itself.

In Part I of this book we present the pure theory of welfare economics and construct a coherent, if abstract, framework for the practice of welfare economics. In Part II we focus on the technical problems encountered in putting the pure theory to work. In particular, we address the problem of finding a criterion for measuring welfare change which synthesizes the positive and ethical components of welfare economics. While particular applications of welfare economics lie beyond the scope of this book, we do conclude with a chapter on cost-benefit analysis where many of the principles developed are put into practice.

We have tried to make our discussion as comprehensive as possible and have included both mainstream material and material not normally covered in treatises on the subject. Thus, we discuss the notion of efficiency as well as the measurement of welfare change under uncertainty, and in an intertemporal framework. A detailed account is made of the sources of market failure and of the measurement of welfare change in distorted economies. We have tried to present a comprehensible account of the vast, and often daunting, literature on social welfare functions. And, we have discussed in detail the notion of a compensation test and its implementation as a measure for ordering Pareto non-comparable states. The overriding theme of the entire book is the notion of welfare economics as the study of ordering social states, and the relevance of the concepts of economic efficiency, the compensation principle, and the conventional measures of welfare change to that.

We have attempted to present the material in a reasonably non-technical manner. We do presume, however, that the reader has a good understanding of undergraduate microeconomics, but we review the more advanced theoretical concepts utilized in the presentation. We also expect that the reader is familiar with basic multivariate calculus. We hope the level of presentation will make the material accessible to advanced undergraduate students in addition to graduate students and professional economists.

In preparing this book we have benefited from comments and discussion on various chapters by several people including Chuck Blackorby, Richard Harris, Jack Mintz, Tony Shorrocks, Dan Usher, John Weymark and David Wildasin. We also were assisted in the preparation of the manuscript by Simon Anderson, Steve Clark and Helen Gaglia, all graduate students at Queen's. We are grateful to all of these people. We also owe a great intellectual debt to the many authors who have contributed to the body of knowledge we present and synthesize in this book. We hope we have done them justice. Finally, we would like to express our appreciation to Dorothy MacKenzie who spent many weeks on the word processor preparing the manuscript.

Contents

CHAPTER 1	The Study of Welfare Economics	1
1	The Scope and Purpose of Welfare Economics	1
2	The Welfare of the Household	8
3	Pareto Optimality and Market Equilibrium	11
4	Social Welfare Functions	16
5	Welfare Change for a Single Household	19
6	Welfare Change in a Many-Household Economy	23
7	Cost-Benefit Analysis	26

PART I THE PURE THEORY OF WELFARE ECONOMICS

CHAPTER 2	The Welfare of the Household	31
1	Introduction	31
2	Household Choice in a Static, Certain Environment	32
3	Household Welfare in a Static, Certain Environment	39
4	Household Welfare with Quantity Constraints	43
5	Household Welfare in an Intertemporal Setting	46
6	Household Welfare under Uncertainty	50
7	The Measurement of Household Utility	56
8	Summary	59

CHAPTER 3	Welfare in a Many-Household Economy: The Pareto Criterion	61
1	Introduction	61
2	Economic Efficiency in an Exchange Economy	63
3	Economic Efficiency in a Production Economy	71
4	The Relevance of Pareto Optimality in Welfare Economics	82
5	Extending the Fundamental Theorems	84
6	The Compensation Principle and the Efficiency Criterion	96
7	Summary	101

CHAPTER 4	Market Failure and the Theory of the Second Best	103
1	Introduction	103
2	Market Failure: Causes and Cases	103
3	Market Failure and the Role of the Public Sector	129
4	The Theory of the Second Best	131
5	Conclusions	135
CHAPTER 5	Social Welfare Orderings: Requirements and Possibilities	137
1	Introduction	137
2	The Framework of Normative Social Choice Theory	139
3	Informational Restrictions and the Social Welfare Orderings	143
4	Non-comparability and Dictatorship Possibilities	147
5	SWO Possibilities with Full Comparability	152
6	SWO Possibilities with Partial Comparability	161
7	Summary and Interpretation	166
CHAPTER 6	Fairness, Equity and Distributive Justice	170
1	Introduction	170
2	Fairness as Absence of Envy	171
3	The Theory of Justice	175
4	Interhousehold Utility Comparability and Moral Preferences	184
5	Conclusions	189
PART II	APPLIED WELFARE ECONOMICS: THE THEORY OF WELFARE CHANGE MEASUREMENT	
CHAPTER 7	The Measurement of Welfare Change for an Individual	195
1	Introduction	195
2	The Problem Defined	196
3	The Compensating Variation (CV)	201
4	The Equivalent Variation (EV)	205
5	An Example – The Deadweight Loss of Excise Taxation	206
6	Approximate Measures of Welfare Change	211
7	Producer's Surplus or Rent	220
8	Intertemporal Welfare Change Measures	223
9	Measurement of Welfare Change under Uncertainty	226

CHAPTER 8	Welfare Change Measurement in Single-Consumer Economies	235
1	The Interpretation of Single-Consumer Economies	235
2	The Welfare Cost of Taxation in a Single-Consumer Economy	237
3	Welfare Change Measures in a Distorted Economy	241
4	Project Evaluation Criteria	243
5	The Social Cost of Public Funds	246
6	An Alternative Way to Express Welfare Change Measures	250
7	An Algebraic Demonstration	255
8	The Theory of Optimal Economic Policy	257
CHAPTER 9	Measuring Welfare Changes in a Many-Consumer Economy	262
1	Introduction	262
2	What Does the Aggregate CV Actually Measure?	263
3	The Aggregate CV and the Strong Test	265
4	The Aggregate CV and the Weak Test	267
5	Other Tests and Other Measures	269
6	Measuring Welfare Change Using a Social Welfare Function	271
7	The Use of Aggregate Income Distribution Statistics	281
CHAPTER 10	Cost-Benefit Analysis	292
1	Introduction	292
2	The Appropriate Decision Rule	294
3	Evaluating the Net Benefits of a Project	297
4	Choice of a Social Discount Rate	314
5	The Treatment of Risk	322
	References	329
	Author Index	337
	Subject Index	339

CHAPTER 1

The Study of Welfare Economics

1 The Scope and Purpose of Welfare Economics

Welfare economics is the framework within which the normative significance of economic events is evaluated. In order to make statements about the consequences for economic welfare of an event we must go beyond the study of *positive* economics, which is concerned with the effects of an event on objectively measurable economic variables such as price and quantity. That is, the welfare economist wishes to determine the desirability of a particular policy – not in terms of his or her own values, but in terms of some explicitly stated ethical criteria.

This book presents the modern theory of welfare economics, and the techniques of its application, as a unified whole. At the outset it is useful to outline precisely our view of the nature of welfare economics (or, more generally, normative economics) and of its scope and limitations. In this introductory chapter we shall attempt to summarize the contents of the subsequent chapters and to indicate how each is related to the underlying theme of welfare economics.

Welfare economics can be viewed as an investigation of methods of obtaining a *social ordering* over alternative possible *states of the world*. A social ordering permits one to compare all states of the world and rank each one as ‘better than’, ‘worse than’, or ‘equally as good as’ every other. Ideally, we would like the ordering to be complete (so that all states could be ranked or ordered) and transitive (so that if state A is at least as good as state B, and state B is at least as good as state C, then state A must be at least as good as state C in the ranking). The term ‘state of the world’ can be interpreted as a complete description of a possible state of an economy including economic characteristics, political conditions such as freedom of speech and non-discrimination, physical characteristics such as the weather, and so on. We are primarily interested in ranking states which differ in economic characteristics, such as outputs of commodities, supplies of factors to different uses, and distributions of commodities over households. That is, we are interested in ranking different *allocations of resources*, where this is used in its broadest sense to refer to the combinations of commodities produced and consumed by each decision-maker in

the economy and the combinations of factors used in the production of each commodity. When we refer to a state of the world, or a *social state*, we are usually referring to a particular allocation of resources. We abstract from all other possible characteristics of the state of the world by assuming them to remain the same over the states under comparison. This convention ought to be construed as meaning not that non-economic characteristics are irrelevant in ranking social states, but only that most (though not all) of the tools we are going to develop and use are designed specifically to rank states differing only in economic characteristics.

The ranking of social states is inevitably a normative procedure; that is, it involves making *value judgments*. For each set of value judgments adopted, a different social ordering results. Therefore there is no objective or unique way to order social states. This does not imply that nothing useful can be said. On the contrary, some value judgments might, in fact, command widespread support, and rankings based on them might therefore legitimately form the basis for actual policy prescriptions. The use of welfare economics for policy purposes is, we would argue, based on this premise. Much of the welfare economic analysis underlying policy prescriptions is based on a certain set of value judgments which are widely accepted among economists, including ourselves.

A good deal of welfare economics is based on the concept of *economic efficiency*. This concept is used to order social states on the basis of some minimal value judgments. Two main value judgments are involved. The first is that the social ordering ought to be based on individual orderings of alternative social states, that is, on individual preferences, where it is implicitly assumed that each individual is the best judge of his or her own preferences. This assumption is called *individualism*. There are obviously a number of problems with it, the main one being that some individuals may be judged incompetent to formulate their own preferences, for example because of insanity or youth. This is partly overcome by our practice of referring interchangeably to the basic decision-maker as an individual and a household, thereby allowing for the fact that parents may determine preferences on behalf of their children. A second problem concerns the treatment of preferences of persons not alive at the time (e.g. the deceased and the unborn). The households upon whose preferences the social ordering is based are assumed to have preferences that are complete and transitive and, more generally, satisfy the axioms of rational consumer behaviour usually posited in microeconomics and discussed in chapter 2. It is important to note that we do not preclude the possibility that an individual's ordering of social states depends upon the commodities consumed by others as well as the commodities consumed by himself or herself; that is, individuals can have sympathetic or antipathetic preferences. In this way individuals can rank entire resource allocations or states of the world.

The second widely accepted value judgment is the *Pareto principle*, which says, in its strongest form, that if state A is ranked higher than state B for one person, and all other persons rank A at least as high as B, then A

should be ranked higher than B in the social ordering. Although one can imagine certain circumstances in which the Pareto principle may conflict with other fundamental value judgments (such as freedom of speech), much of both theoretical and applied welfare economics is based upon it. In our analysis we assume it to be a desired property of the social ordering.

In order to make use of these two ethical principles, it is necessary to have a theory of how an individual's level of welfare or utility is determined by the state of the world. This theoretical basis is provided by the utility maximizing theory of consumer choice. This theory posits that consumers will rank alternative states according to a set of *preference orderings*. We can then attempt to 'aggregate' the preferences of different households using the 'dominance' notions contained in the Pareto principle. Unfortunately, this procedure permits only a *partial* ordering of social states. In particular, we cannot compare two states where one is preferred by one or more households and the other is preferred by one or more other households. Such states are said to be *Pareto non-comparable* and can only be ranked if additional value judgments are made. Since many states are Pareto non-comparable, this is an important limitation. For example, we are precluded from making judgments about policies which have income redistributive effects (unless individuals are sufficiently altruistic that everybody feels better off after the redistribution).

Despite this discouraging fact, the normative analysis of economic efficiency proceeds with only individualism and the Pareto principle as value judgments. Since individualism and the Pareto principle represent rather weak value judgments, to be able to rank two states according to these criteria has strong appeal. There is an important relationship between resource allocations generated by the market economy and those ranked highly by the Pareto criterion. This relationship is summarized in what are referred to as the *two basic theorems of welfare economics*. First, under certain assumptions concerning the convexity of individual preferences and technology, any allocation of resources generated as a general equilibrium of a perfectly competitive economy is *Pareto optimal*; that is, there is no other feasible allocation which ranks above it according to the Pareto principle. Secondly, the converse is also true: any Pareto optimal allocation of resources can be achieved by the solution to a general equilibrium in a competitive economy.

Certain assumptions must be satisfied to ensure that these two basic theorems hold. If any of these assumptions are not satisfied, then the market allocation of resources may not necessarily be Pareto optimal – that is, there may be some other allocation of resources which dominates those actually attained – and there is said to be *market failure*. The existence of market failure implies that there are some opportunities for mutual gains that are not being exploited (since there are other feasible allocations not attained in which everyone can be better off). Such situations are described as 'inefficient'. The study of welfare economics is useful in identifying such inefficiencies and in recommending and evaluating 'corrective' policies.

Even in the absence of market failure, the fact that competitive economies are Pareto optimal is of limited use. There are, in general, a large number of Pareto optimal allocations of resources in the economy, each associated with different combinations of levels of well-being of the various households in the economy. For example, in a competitive market economy, the equilibrium actually achieved depends upon the pattern of ownership of society's primary factors of production (e.g. land, labour, capital) among the households of the economy. Unfortunately, the many Pareto optimal outcomes themselves are Pareto non-comparable. Furthermore, even a particular Pareto optimal allocation does not dominate all non-Pareto optimal allocations on Pareto grounds; that is, certain non-Pareto allocations will be Pareto non-comparable with respect to Pareto optimal ones. This reflects the fact that rankings based solely upon the Pareto principle are incomplete.

One attempt to deal with this incompleteness is to extend the Pareto principle by means of a *hypothetical compensation test*. By this is meant the following. Even if two resource allocations are Pareto non-comparable, it may still be the case that if it is possible for those who are better off in state A than in state B to compensate those who are worse off such that, if the compensation actually were paid, all persons would be better off than in state B. In this case the hypothetical compensation test would be satisfied and state A would be ranked higher than state B according to this criterion. This criterion is used extensively in applied welfare economics in order to separate the efficiency issues, discussed above, from the equity issues to be discussed below.

There are some difficulties with relying upon the Pareto criterion, augmented by the compensation test, as a method of ordering social states. First, the acceptance of the criterion requires that one adhere to the strong value judgment inherent in the hypothetical compensation test. If some persons are worse off in state A than in state B, the fact that they could *hypothetically* be made better off by a transfer of some sort may be small consolation. Secondly, it transpires that there are different notions of compensation which depend upon the nature of transfers that are assumed to be hypothetically possible. For example, the compensation may be viewed as taking the form of lump-sum distributions of either a bundle of goods or purchasing power, or the compensation may be carried out by distorting taxes and transfers. Finally, even if one did accept the compensation test, a complete ordering of social states is still not possible. Some states, such as the Pareto optimal ones, cannot be ranked by the compensation test.

For these reasons, a complete study of welfare economics attempts to go beyond the concepts of economic efficiency based on the Pareto principle. This involves devising some means of weighting the utilities of different households, and this, in turn, requires that stronger value judgments be made. The value judgments can be codified into a *social welfare function* (SWF) or ordering. The SWF is an important conceptual tool in welfare economics, since it is the means by which a complete social ordering is obtained.

The choice of an SWF is constrained by the ethical assumptions to be incorporated and the information about individual utilities assumed to be available. The latter depends on the extent to which individual utilities are measurable and the extent to which the utility measures can be compared among households. The least restrictive informational assumptions are that the individual utilities are only *ordinally* measurable and *non-comparable* among households. The types of social welfare functions available under such assumptions are extremely limited and violate some reasonable equity assumptions. A necessary condition to obtain a wider choice set of social welfare functions is that household utilities be comparable with each other in some way. It is also useful if household utility is measurable on a scale that is more informative than an ordinal scale. In general, as the comparability and measurability assumptions are strengthened, the choice set of social welfare functions expands.

As mentioned, an SWF will incorporate both informational assumptions and ethical ones. A fairly weak set of ethical judgments will lead to an SWF which is said to be *welfaristic*, where the term 'welfaristic' means that the SWF depends only on the utilities achieved by the individuals in the economy and not on any other information. The ethical principles underlying a welfaristic SWF include the Pareto principle as well as some slightly more technical ones known as the independence of irrelevant alternatives and unrestricted domain. These will be introduced in more detail in section 4.

Given sufficiently strong measurability and comparability assumptions, the welfare economist can choose among many welfaristic social welfare functions. There are several considerations relevant to the choice of an SWF, of which some involve a further ethical judgment and others are purely for analytical convenience. The economist may argue, for example, that it is reasonable to include an 'egalitarian ethic' in the social welfare function because such an ethic would be widely agreed upon. But how egalitarian should the social welfare function be? One can choose an extremely egalitarian SWF or a weakly egalitarian SWF. For this reason it is often useful to choose social welfare functions which embody important variable properties, such as the degree of egalitarianism, into a parameter which can itself be varied. More generally, social welfare functions may be chosen for their tractability and flexibility as well as for the ethical norms they represent.

Despite this rather arbitrary procedure, the use of an SWF is an important element in modern welfare economics. With such a device, the framework of welfare economics is complete and well structured in principle. The remaining problems concern the application of this framework, or the study of *applied welfare economics*, to which we now turn.

In applying welfare economics, we immediately face a conundrum. Much of the theory of welfare economics was designed to produce an ordering or ranking of social states, either at the individual level or for society as a whole. This ordering can be represented by an ordinal function, either an ordinal utility function for the individual or an ordinal

social welfare function. In applied welfare economics we are often asking for something much more; we would like to have a measure of welfare which enables us to measure the *welfare change* in going from one allocation to another. It is important to realize that the yardsticks of measurement used in practice are to some extent arbitrary. Since all we know is a ranking of social states, there are an indefinitely large number of measuring or numbering schemes which can be applied to the ordering and which preserve the ordering intact. In technical terms, if a particular welfare function faithfully ranks alternatives according to an ordering of social states, then any 'composite' welfare function derived from an increasing function of the first welfare function will do just as well. The particular representation or transform that we choose is purely for convenience. Thus we often find it useful to obtain a monetary representation of household utility functions, defined as the amount of money required to attain various utility levels at a set of reference prices. This monetary representation is referred to as a *money metric*, and its value is obtained from a very important technical tool, the *expenditure function*. To repeat, this money metric is simply one of many ways of representing the household's ranking of social states. No significance ought to be attached to the numbers other than as reflections of an ordering. Applied welfare economics makes use of this money metric of household utility to investigate the measurement of welfare change in a variety of circumstances.

The use of this money metric to measure welfare change for an individual is generally restricted to the case in which the individual cares only about his own consumption bundle and not about those of others. Indeed, much of welfare economics is devoted to measuring welfare change in an economy of 'selfish' individuals. This is done partly for simplicity and partly to reflect the prevailing view that most economic decisions are made by a narrowly defined 'caring' unit. It also has the important advantage that our welfare change measures can, in principle, be measured on the basis of observed market prices and quantities. That is because the price of a good measures its marginal benefit to a consumer in terms of the money metric. The tools of applied welfare economics developed under these assumptions would not generally be applicable in an economy consisting of non-selfish individuals.

Welfare change measures for an individual involve obtaining an estimate of the change in value of the money metric resulting from either a change in the bundle of commodities consumed by the individual or, because the commodity bundle chosen by the individual depends upon prices and income, from a change in the prices and income facing the individual. There are a number of alternative money metrics we could use. The two most commonly used are the *compensation variation* (CV) and the *equivalent variation* (EV). For the most part our exposition uses CV.

From the measure of CV for an individual facing exogenous changes in prices and income, it is a simple step to measure the CV for changes in resource allocation in a single-person economy. Now the prices and

income facing the individual are no longer exogenous, but are determined by the interaction of the demand and supply sides of the economy. To analyse the supply side requires that we introduce the production possibilities of the economy. Changes in resource allocation can come about through changes in the primary factor supplies or the technology of production on the supply side. Alternatively, they can come about through changes in the actions of government, in the form of either tax changes or the introduction of public projects.

Obviously the single-person economy is an analytical myth. However, because it is much more difficult to observe individual commodity bundles than aggregate commodity bundles, it is worth while to determine the conditions under which we can treat the many-person economy as if it were a single-person economy, and treat the aggregate money metric (CV) as if it were a single person's CV. Unfortunately, the conditions are quite stringent. For this reason, applied welfare economics must confront the problems of measuring welfare change in a many-consumer economy in a manner which is consistent with some underlying social ordering or (ordinal) social welfare function. Although this task is difficult, it is not impossible, and simplifying assumptions must inevitably be made. The culmination of this effort lies in the field of cost-benefit analysis or project evaluation, in which the tools of theoretical and applied welfare economics, and in particular the measuring rod of a money metric, are brought to bear on the social evaluation of public projects. Although much cost-benefit analysis is concerned with specific problems encountered in evaluating investment, the general methodological issues are those addressed in welfare economics.

In this book we expound the methodology and techniques of welfare economics as outlined above. We approach the various issues in roughly the same order in which they were presented above; that is, chapters 2-6 deal with the theory of deriving a social ordering over states of the world, and chapters 7-10 deal with the uses to which the social ordering is put. We focus on the theoretical underpinnings of welfare economics and on the general methodological problems encountered in its application. We also survey its application in the form of cost-benefit analysis. Space precludes us from considering all the applications of welfare economics, since they are found in virtually all areas of economics including international trade, industrial organization, public finance and macro-economic policy-making.

The analysis found in the subsequent chapters is meant to be reasonably thorough and carefully argued. In places it may be fairly subtle as well. It might therefore be helpful at this point to provide a precis of the main currents of arguments and results found in the book. This should prove helpful as an overview of each chapter which the reader can use to navigate through the detailed analysis. Also, it will enable readers (or teachers in welfare economics courses) to skip sections or entire chapters without losing the chain of argument or the knowledge of how the different sections fit together as a whole. Without suggesting that he or

she may wish to do so, the reader could even, we think, read either chapters 2-6 or 7-10 independently. What follows, then, is a survey of the contents of the subsequent chapters.

2 The Welfare of the Household

In this book we take the view that society's welfare ultimately depends upon the welfare of its constituent households. Furthermore, we follow much of the literature on welfare economics in assuming that social welfare depends *only* on the welfare of households. This has been referred to as the assumption of *welfarism* by Sen (1979) and can be deduced on the basis of some rather weak value judgments as discussed in chapter 5. The welfare of the household therefore provides the fundamental information upon which social orderings are based, and we begin by establishing exactly what is meant by the welfare of the household.

One further important value judgment required is that the welfare of the household must either correspond with the household's own view of its welfare, or at least be consistent with the household's preferences. The assumption that social welfare must respect household preferences is called *non-paternalism*. This assumption is rather important for many of the results of theoretical and applied welfare economics, since household preference orderings are reflected in the choices the household makes when faced with alternative situations. We can therefore deduce information concerning the household's preference orderings on the basis of choices actually taken. In principle, we could obtain the entire preference ordering if enough choice situations were presented to the household. The bulk of chapter 2 is devoted to reviewing the standard theory of consumer choice. Most of this is done in the context of a selfish household, that is, one that cares only about its own consumption. There is, however, no reason why it cannot be extended to cases in which household preference orderings over social states depend also upon the consumption of other households.

According to the basic theory of household choice, the household has a preference ordering over all alternative bundles of commodities which satisfies certain assumed properties. These properties ensure that the preference ordering can be represented by an ordinal utility function. The consumer is confronted with a social state which is represented by a budget set from which the consumer can choose a bundle of commodities. The budget set will be determined by the set of prices and income facing the household. The household then chooses the most preferred bundle in the budget set, or the one which maximizes its utility function. Since the budget set depends upon prices and income, the bundle chosen does also. As prices and income change, so does the bundle demanded. We characterize the choice by the household as a set of *demand functions*, one for each commodity, and each depending on prices and income. These demand functions reflect what we know about household preferences. Several properties of demand functions are discussed at length in chapter 2.