1 Introduction

1.1 A look at inequality analysis

Thinking about inequality is not always a fashionable topic amongst economists. But thinking about inequality actually goes on all the time. Perceptions of inequality affect economic choices and political decisions. A sensitivity to inequality coupled with compassion for the poor motivates charitable giving by individuals and states. Notions about inequality appear to inform popular views about the appropriateness or otherwise of pay awards. And any parent with two or more children needs no formal analysis to be persuaded of the importance of distributive justice. Fashionable or not, thinking about inequality plays a part in the judgments and actions of politicians, planners and ordinary people.

Of course the study of economic inequality has not just been a matter of fashion. It has been an integral part of the general historical development of political economy and economics, and the approach to the topic has changed with the passage of time. While this is not the place for an extensive treatise on the history of economic thought about inequality, a brief sketch to introduce conventional wisdom on the subject may help to put into context what we want to tackle in this book.

This century has witnessed a shift in emphasis in thinking about inequality. It used to be commonplace to set the analysis of economic injustice within a particular social or institutional framework – such as Ricardo's or Marx's class-based theories of political economy. Alternatively, issues of inequality used to be cast in terms of specific models of income distribution – such as Vilfredo Pareto's famous laws of distribution. However, in more recent times, there has been a move away from these narrowly focused perceptions of the problem to an approach founded upon general principles. What principles?

A cursory review of recent literature suggests that the principles encompass a wide range of theoretical and applied economics. But all the same it is possible

to simplify them down to a relatively few essentials. In practice we may usefully distinguish four major building blocks that are required in the analysis of income distributions:

- *The definition of income.* We need to specify carefully, or to be told clearly, what the thing called 'income' is.
- *The income recipient.* We also need to be clear about the nature of the entities – persons, families, households or whatever – that receive those incomes.
- *The reference group.* We should explicitly define the 'universe': the collection of persons or groups within which inequality comparisons are to be made.
- *The calibration system.* The 'inequality thermometer' – the inequality measurement tool – has to be precisely specified.

These four main components of inequality analysis get unequal treatment in the literature. As a sweeping generalisation we may state that items one, two and four in this list get a lot of attention; item three does not.

For example, theoretical economists focus principally on the fourth item in the list: the specification of a system of calibration. This forms a natural extension to a substantial literature on social choice and welfare economics. The way the analytical problems are formulated has close connections with other related issues such as the assessment of risk, the meaning of individual utility and the construction of index numbers of prices and income. In fact, inequality presents a classic theoretical measurement problem, and is typically treated in a classic fashion by setting out a system of axioms that appear to be reasonable and by formulating key propositions that follow from the axiomatic base.

On the other hand, applied economists and statisticians usually pay close attention to the first two issues: it is widely recognised that practical matters in defining income (or wealth, consumption expenditure, or whatever) or the family unit that is the income receiver are essential to understanding levels and trends of inequality within most economies.

But as far as the third point on the list is concerned – the appropriate reference group – one is immediately struck by the lack of references in the mainstream economics literature. Why this apparent neglect of one of the main components of income distribution analysis? Perhaps the answer is that to many researchers the issue seems obvious or self-defining. For example, in an empirical study, the sample is what it is. The population which the sample represents – so it might be argued – does not really need more than the most cursory discussion. Yet in principle the 'universe set' on which income distributions are to be defined and inequality to be assessed is a matter of theoretical as well as practical debate. For example, in the world of Plato or Aristotle the issue of distributional justice was applied only to free men since, in a social system that tolerated slavery, economic injustice for slaves was not a particularly relevant concept (and, of course, women did not get a look in). Similar difficulties have been raised in connection with modern theories of justice: who is to be counted within the

ambit of such theories, or who is to be party to the social contract? The voting public? All adults? The whole population? If the jurisdiction of nation-states can abruptly change, even this last broad definition may be imprecise. Matters become yet more complicated if we try to take account of all the citizens of the world or persons yet unborn. The question has also been raised as to whether the principles that are applied to people should also be applied to cats, dogs and other animals.[1]

Even on the empirical level the issue of the reference group can have a dramatic impact on the picture that emerges about the pattern of world inequality. As a simple instance of this consider the study of international income comparisons by Summers and Heston (1988, 1991). Their hundred-plus countries are divided into six broad groups (Africa, North and Central America, South America, Asia, Europe and Oceania) so that it is possible to obtain a broad-brush picture of world income inequality in 1985 and 1988. But at first glance this broad-brush picture looks rather extraordinary: we find that in 1985 per capita income in Oceania was remarkably low – below Asia and South America so that it ranked fifth out of the six world regions (were the New Zealanders and Australians really so hard-pressed?); but in 1988 per capita income in Oceania had seemingly jumped so that it ranked third out of six (after Europe, but above South America). The answer to this conundrum is not hard to find: in the 1985 data compilation the relatively poor Indonesia – with its 160 million inhabitants – was classified as being part of Oceania; in 1988 it was lumped in with Asia. So, by respecifying the groups only very slightly – in effect just relabelling one country – a substantially different story emerges of income inequality among different regions of the world. Clearly too, whether one counts Indonesia as an Asian country or part of Oceania is going to have a dramatic impact on the perceived inequality within Oceania.

This brief mention of theoretical and practical difficulties is not intended to imply that clear comparisons of inequality are usually impossible or meaningless. But it serves to highlight the importance of what might appear to be mere background features of the problem in making sensible inequality comparisons.

1.2 A second look

In our view there are deeper problems associated with the issue of the reference group. In fact it is arguable that the issue lies at the root of some of the more intractable problems in the assessment of income distribution. One of these problems – which we shall be taking up later in the book – is the relationship between the analysis of economic inequality and the analysis of poverty. Over recent years each of these two related topics has been extensively developed in terms of a mathematical approach founded upon a set of formal assumptions or axioms. But they have been developed separately, each using a distinct set of axioms as an intellectual basis. The intellectual divorce between the two branches of the

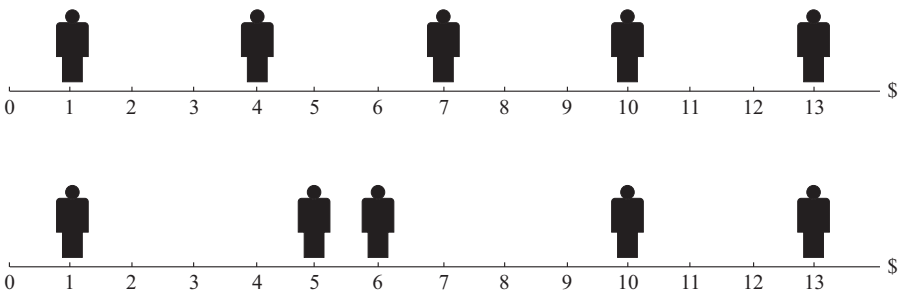

Figure 1.1. A simple distributional experiment.

subject can to some extent be explained in terms of different approaches to the idea of a reference group as we shall see further in chapter 7.

The way that reference groups are perceived also has a bearing upon some basic propositions in inequality analysis. In effect, what people mean by inequality can be crucially dependent on their perception of the relevant reference groups and in the ways that these groups are interlinked. The problem of the reference group and the way in which it relates to people's thinking about inequality is actually a convenient introduction to the case for a second look at the basics of inequality and income distribution analysis.

As an example of what is involved here, try a simple experiment. Figure 1.1 shows two possible income distributions in a very elementary economy. Each distribution contains five persons who have been arranged on an income scale in positions corresponding to their incomes, and the two distributions have the same total income ($35). The units of income are irrelevant in the experiment (the '$' sign has an unspecified value) but let us suppose that income tells us all that we might need to know about the 'well-offness', economic status or whatever of the persons; and we might as well assume that the five anonymous persons are as identical as the caricature suggests them to be. The experiment is simply this: write down which of the two distributions appears to you at first sight to be the more unequal and, if possible, give reasons for your answer.

Now, noting that the difference between the two distributions directly affects only two of the persons in the experiment, consider the slight modification of the diagram that is presented in figure 1.2. Here we have explicitly divided the population of five into two component groups, left and right, as indicated by the shading, but the distributions are in reality just the same as in figure 1.1. Notice that in each of the two subgroups taken separately it is arguable that the situation at the bottom of the diagram represents greater inequality than that at the top. The richest person in the left-hand group has a higher income ($5 rather than $4) and the poorest person in the right-hand group has a lower income ($6 rather than $7); so in both cases the income gaps within each reference group widen as we go from the top of the diagram to the bottom. However, that is not the end of the story.

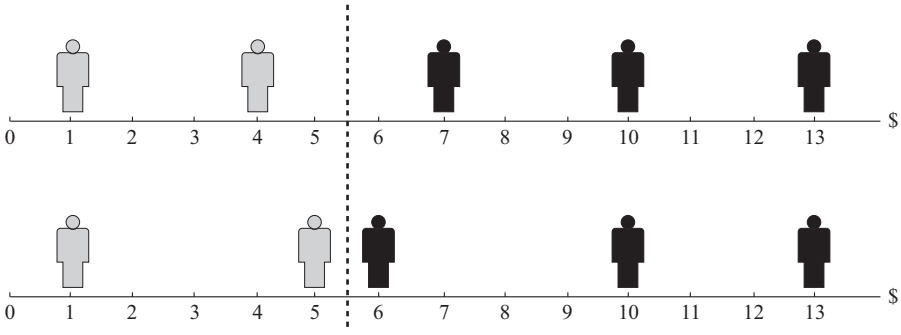

Figure 1.2. A simple distributional experiment: second view.

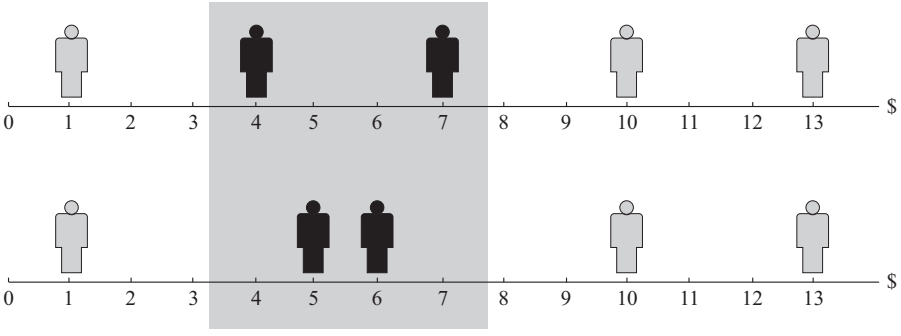

Figure 1.3. A simple distributional experiment: third view.

As a final step in the experiment have a look at figure 1.3, which again merely retouches the picture that was originally displayed in figure 1.1. In this case we have highlighted just the two persons whose incomes are directly affected in going from one distribution to the other. Put in this fashion there appears to be no argument whatsoever as to which distribution is the more unequal. Disregarding for the moment the persons whose income positions remain unchanged in going from one distribution to the other (the people with incomes \$1, \$10, \$13) it is clear that there has been an unambiguous reduction in the gap between the two remaining persons: the gap closes from (\$4, \$7) to (\$5, \$6). Put another way, if we consider the top income distribution as the 'before' picture of inequality, and the bottom distribution as the 'after' picture, then there has been a redistribution of \$1 from a richer to a poorer person: according to this view inequality *must* have fallen.

So we seem to have at least two stories about what is going on in this trivial problem of inequality comparison. How does one go about resolving the apparently contradictory pictures of inequality that emerge from even a very simple experiment such as the one we have been considering? Indeed, is there any point in trying to resolve such contradictions? Evidently the way that one tries to

answer this sort of question will strongly influence one's entire conception of the meaning of inequality comparisons.

The conventional approach to the subject has been twofold. On the one hand there is the horny-handed practical approach to evaluating empirical income distributions: having argued about the right way to measure income, and whether we should view income receivers as individuals, families, households or other groups, you pick a measure of dispersion off the shelf and you estimate this along with mean income and other statistics (we shall be looking at some of these off-the-shelf measures later in the book). Alternatively there is a theoretical approach to the problem that appears to be full of the intellectual promise that accompanies analytical rigour: this method is to introduce a particular set of axioms which collectively define what is meant by inequality comparisons and hence what is meant by economic inequality itself. It is essentially the picture of inequality characterised by figure 1.3 that is taken to be the standard paradigm for the majority of theoretical and empirical work in the economics literature.

The ambiguity of answers from the simple experiment raises issues that are considerably wider than the particular principle or principles which may be brought to bear on the particular distributional problem highlighted in figure 1.1. It prompts the question as to whether the way in which inequality is conventionally presented in the literature on economics and in other related disciplines is in some sense 'appropriate'.

1.3 A guide to the book

Those who know the economics literature on inequality will say that we have presented the pictures in our little experiment the wrong way round. That is actually quite true. We deliberately put the unorthodox view of the inequality comparison experiment first, and followed up with the standard story. The reason for this has little to do with the grand sweep of the history of thought on the subject, and much to do with a small domestic incident.

A few years ago one of the authors, Yoram Amiel, was asked by his wife Hayka (who is a school-teacher and not an economist) to explain the topic of his research. To put the main ideas over concisely he gave her a little numerical example as an illustration – something close to the experiment that we have just been considering, in fact. Faced with the choice between the two distributions, Hayka gave the 'wrong' answer. Yoram concisely pointed this out: the answer should have been clear, unambiguous and the exact opposite to hers – according to the standard theory of inequality measurement. Hayka's reply was similarly concise: 'So change the theory.'

We make no claim to be changing the theory of inequality measurement in this book. But this issue did prompt an extensive research project which, amongst other things, resulted in the book. Along the way it also raised a number of provocative questions which we make no claim to have resolved. Is the standard

theory 'right' and, furthermore, what does it mean for a theory to be 'right' in this context? These questions have in turn prompted the theoretical and empirical analysis which is reported in the following chapters.

Chapters 2 to 4 are principally about the problem of inequality in its purest form, the problem of dividing a cake of fixed size amongst a fixed number of people. Chapter 2 gives a summary guide to the standard approach in the literature on inequality measurement, chapter 3 explains the method we used to investigate the assumptions underlying the approach – a series of specially structured questionnaires – and chapter 4 reports the results of these investigations. These chapters also deal with elementary issues of how one can compare situations that have different sizes of cake or different numbers of people sharing the cake.

Chapters 5 to 7 extend the approach to three areas closely related to the pure inequality problem: social welfare (chapter 5), the relationship between income growth and perceptions of inequality (chapter 6), and poverty (chapter 7). Each of these additional topics requires additional assumptions on top of the structure used for the pure inequality problem, and we subject these to the same sort of investigative strategy; they also provide us with an opportunity to check our results on the pure inequality issue.

Chapter 8 makes a comparison of responses to our various questionnaire studies across countries and across academic disciplines, while chapter 9 sums up and suggests directions in which thinking about inequality may yet go. Finally, those readers who like to have assumptions and propositions tidied up in a concise mathematical format may want to use appendix A which sets out the main results in the conventional approach to inequality measurement; all this material has been parked in this unglamorous location because, although it has its uses, it is no substitute for thinking about inequality.

Notes

1 See, for example, Sen's discussion of Rawls's concept of the 'original position' (Sen 1970, p. 124).

2 What is inequality? The economists' view

2.1 The axiomatic approach

As we mentioned in chapter 1 the standard approach to the problem of inequality comparisons is based upon a formal structure that is usually expressed in terms of precise assumptions – or axioms – and mathematical propositions. So, in order to grasp what is going on and what it is that we are investigating in the rest of this book, it may be helpful to have a brief introduction to the axiomatic methodology.

The axiomatic methodology is a grand name for an essentially simple approach to our subject. It consists of a rule-based system of thought which enables us to state precisely what we mean by inequality comparisons, and thereby what we mean by inequality. The axioms are formal assumptions which are taken as fundamental: they are not derived from even more basic assumptions; they do not have to be based on any real-world experience or observation; they require no appeal to any external value system. Whether the axioms are 'true' or not, and what is meant by the 'truth' of a set of axioms are moot points. What are the circumstances under which the axiomatic approach is valuable, and what is one trying to achieve by adopting this approach?

Perhaps a rough-and-ready description of the principal advantage of the axiomatic approach is that it acts as a systematic antidote to the trial-and-error approach of picking apparently suitable ready-made statistics. This description both overstates and understates the case. It is a slight overstatement because picking an axiom system can in fact amount to little more than the trial-and-error approach in a rather more sophisticated guise: if you happen to want to use a particular measurement tool anyway, it is not too difficult to write down a set of 'basic' assumptions which will imply that your pet measure is in fact the only one that is available for use. Our simple description also understates the rôle of the axiomatic approach: actually it is potentially quite powerful because, by adopting this method of analysis, it is possible to set out the exact relationship between

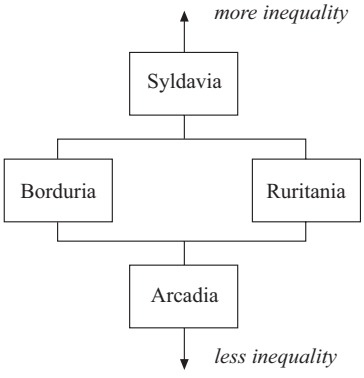

Figure 2.1. An inequality ranking.

particular principles or rules of comparison, and the types of specific mathematical formulae available for summarising income distributions. In fact it aids in formulating basic ideas about the meaning and structure of inequality comparisons.

2.2 Inequality rankings and orderings

Inequality comparisons are the basic idea with which we need to work. From several of these comparisons we may then build up an inequality ranking or, perhaps, an inequality ordering. An inequality comparison is simply a coherent rule for comparing distributions of income (or wealth, consumption, and so on) in two situations. Has the income distribution in our country become more equal over the last five years? Is Syldavia a more unequal society than Borduria? For an inequality comparison to be meaningful, then the answer 'yes, Syldavia is more unequal than Borduria' must also imply, in the other direction, 'Borduria is less unequal than Syldavia.' Of course this simple pairwise rule does not say anything about whether such pairwise comparisons can always be made, or whether connecting up different pairs of comparisons (where they can be made) is also possible. We can do more if the comparison rule is transitive, which means that inequality comparisons can be chained together: the statements 'Syldavia is more unequal than Borduria' and 'Borduria is more unequal than Arcadia' together imply the statement 'Syldavia is more unequal than Arcadia.' Given transitivity we can produce an inequality ranking such as that illustrated in figure 2.1.

A transitive ranking of distributions may nevertheless leave certain gaps in the set of possible pairwise comparisons. This is also illustrated in figure 2.1 where the inequality comparison rule implies that Syldavia is more unequal than Ruritania which in turn is more unequal than Arcadia, but that Borduria and Ruritania cannot be compared in terms of the rule. Notice that we are not saying

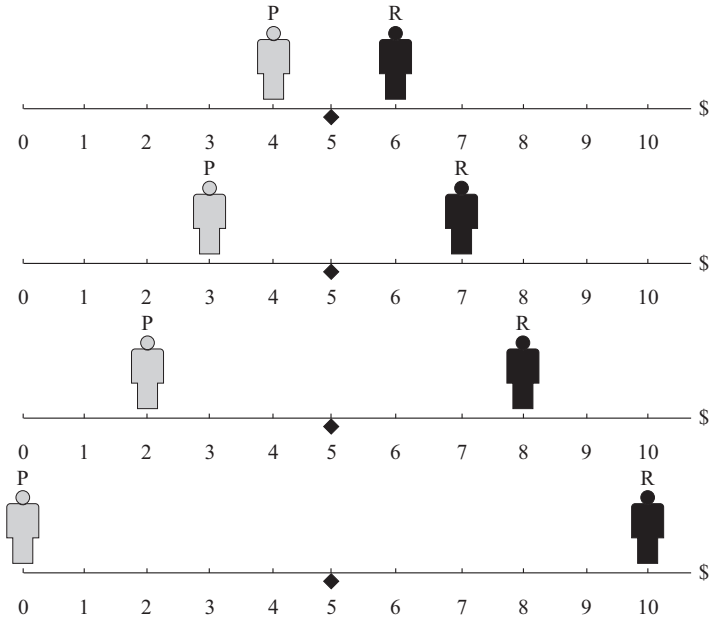

Figure 2.2. Inequality in a two-person world.

that the inequality rule indicates that Borduria and Ruritania have the same inequality, but rather that the rule is indecisive in this case. According to the rule that generated this ranking we just do not know which of the two is more unequal.

However, if the inequality comparison rule is always decisive (or 'complete' to use the standard jargon), then, combined with the property of transitivity we could obtain an inequality ordering. At one level this is just what we mean by 'measuring' inequality. As we shall see, the meaning of inequality comparisons depends critically upon the axiomatic basis that is specified for the inequality comparison rule.

In one very special case, virtually nothing is required in terms of axiomatisation. Given a two-person world with a fixed total income there is very little to say in terms of inequality comparisons. A brief look at the four distributions in figure 2.2 is sufficient to make the point. Because there is by assumption a fixed total income (in this case \$10) the two persons (P and R) must be equally spaced around the position of mean income (\$5); it is clear that as we move from the top of the diagram to the bottom, inequality steadily increases as the rich person R and the poor person P move farther apart. So what is inequality in this case? We could measure it as the income gap between the two figures (\$2 in the first line, \$4 in the second, and so on), or we could measure it as the proportionate gap between the two (\$2/\$10 = 0.20 in the first line, and so on); or, if we wanted to, we could use the square or the cube-root of the distance between persons P and R: it matters little except in terms of the scaling of the 'thermometer' by which we propose to measure inequality.