

14.1 Introduction

A social welfare function permits the evaluation of economic policies that cause redistribution among consumers—a task that Pareto-efficiency can never accomplish. Although the concept of a social welfare function is a simple one, previous chapters have identified numerous difficulties on the path between individual utility and aggregate social welfare. The essence of these difficulties is that if the individual utility function corresponds with what is theoretically acceptable, then its information content is too limited for social decision-making.

The motivation for employing a social welfare function was to be able to address issues of equity as well as issues of efficiency. Fortunately a social welfare function is not the only way to do this, and as this chapter will show, we can construct measures of the economic situation that relate to equity and that are based on observable and measurable information. This provides a set of tools that can be, and frequently are, applied in economic policy analysis. They may not meet some of the requirements of the ideal social welfare function, but they have the distinct advantage of being practically implementable.

Inequality and poverty provide two alternative perspectives on the equity of the income distribution. Inequality of income means that some households have higher incomes than others—which is a basic source for an inequity in welfare. Poverty exists when some households are too poor to achieve an acceptable standard of living. An inequality measure is a means of assigning a single number to the observed income distribution that reflects its degree of inequality. A poverty measure achieves the same for poverty. Although measures of inequality and poverty are not directly social welfare functions, the chapter will reveal the closeness of the link between the measures and welfare.

The starting point of the chapter is a discussion of income. There are two aspects to this: the definition of income and the comparison of income across families with different compositions. In a setting of certainty, income is a clearly defined concept. When there is uncertainty, differences can arise between *ex ante* and *ex post* definitions. Given this, we look at alternative definitions and relate these to the treatment of income for tax purposes. If two households differ in their composition (e.g., one household is a single person and the other is a family of four), a direct comparison of their income

levels will reveal little about the standard of living they achieve. Instead, the incomes must be adjusted to take account of composition and then compared. The tool used to make the adjustment is an equivalence scale. We review the use of equivalence scales and some of the issues that they raise.

Having arrived at a set of correctly defined income levels that have been adjusted for family composition using an equivalence scale, it becomes possible to evaluate inequality and poverty. A number of the commonly used measures of each of these concepts are discussed and their properties investigated. Importantly, the link is drawn between measures of inequality and the welfare assumptions that are implicit within them. This leads into the idea of making the welfare assumptions explicit and building the measure up from these assumptions. To measure poverty, it is necessary to determine who is “poor,” which is achieved by choosing a level of income as the poverty line and labeling as poor all those who fall below it. As well as discussing measures of poverty, we also review issues concerning the definition of the poverty line and the very concept of poverty.

Although the aim of this chapter is to move away from utility concepts toward practical tools, it is significant that we keep returning to utility in the assessment and improvement of the tools. In attempting to refine, for example, an equivalence scale or a measure of inequality, it is found that it is necessary to comprehend the utility basis of the measure. Despite intentionally starting in a direction away from utility, the theory returns us back to utility on every occasion.

14.2 Measuring Income

What is income? The obvious answer is that it is the additional resources a consumer receives over a given period of time. The reference to a time period is important here, since income is a flow, so the period over which measurement takes place must be specified. Certainly evaluating the receipt of resources is the basis of the definition used in the assessment of income for tax purposes. This definition works in a practical setting but only in a backward-looking sense. What an economist needs in order to understand behavior, especially when choices are made in advance of income being received, is a forward-looking measure of income. If the flow of income is certain, then there is no distinction between backward- and forward-looking measures. It is when income is uncertain that differences emerge.

The relevance of this issue is that both inequality and poverty measures use income data as their basic input. The resulting measures will only be as accurate as the data that are employed to evaluate them. The data will be accurate when information is carefully

collected and a consistent definition is used of what is to be measured. To evaluate the level of inequality or poverty, a necessary first step is to resolve the issues surrounding the definition of income.

The classic backward-looking definition of income was provided by Henry C. Simons in 1938. This definition is “Personal income may be defined as the algebraic sum of (1) the market value of rights exercised in consumption and (2) the change in the value of the store of property rights between the beginning and end of the period in question.” The essential feature of this definition is that it makes an attempt to be inclusive so as to incorporate all income regardless of the source.

Although income definitions for tax purposes also adopt the backward-looking viewpoint, they do not precisely satisfy the Simons’s definition. The divergence arises through the practical difficulties of assessing some sources of income especially those arising from capital gains. According to the Simons’s definition, the increase in the value of capital assets should be classed as income. However, if the assets are not liquidated, the capital gain will not be realized during the period in question and will not be received as an income flow. For this reason capital gains are taxed only on realization. In the converse situation when capital losses are made, most tax codes place limits on the extent to which they can be offset against income.

We have so far worked with the natural definition of income as the flow of additional resources. To proceed further, it becomes more helpful to adopt a different perspective and to view the level of income by the benefits it can deliver. Since income is the means to achieve consumption, the flow of income during a fixed time period can be measured as the value of consumption that can be undertaken, while leaving the household with the same stock of wealth at the end of the period as it had at the start of the period. The benefit of this perspective is that it extends naturally to situations where the income flow is uncertain. Building on it, in 1939 John R. Hicks provided what is generally taken as the standard definition of income with uncertainty. This definition states that “income is the maximum value which a man can consume during a week and still expect to be as well-off at the end of the week as he was at the beginning.”

This definition can clearly cope with uncertainty, since it operates in expectational terms. But this advantage is also its major shortcoming when a move is made toward applications. Expectations may be ill-defined or even irrational, so evaluation of the expected income flow may be unreasonably high or low. A literal application of the definition would not count windfall gains, such as unexpected gifts or lottery wins, as income because they are not expected despite such gains clearly raising the potential level of consumption. For these reasons the Hicks definition of income is informative but not perfect.

These alternative definitions of income have highlighted the distinctions between *ex ante* and *ex post* measures. Assessments of income for tax purposes use the backward-looking viewpoint and measure income as all relevant payments received over the measurement period. Practical issues limit the extent to which some sources of income can be included, so the definition of income in tax codes does not precisely satisfy any of the formal definitions. This observation just reflects the fact that there is no unambiguously perfect definition of income.

14.3 Equivalence Scales

The fact that households differ in size and age distribution means that welfare levels cannot be judged just by looking at their income levels. A household of one adult with no children needs less income to achieve a given level of welfare than a household with two adults and one child. In the words of the economist William M. “Terence” Gorman, “When you have a wife and a baby, a penny bun costs threepence.” A larger household obviously needs more income to achieve a given level of utility, but the question is how much more income? Equivalence scales are the economist’s way of answering this question and provide the means of adjusting measured incomes into comparable quantities.

Differences among households arise in the number of adults and the number and ages of dependants. These are called *demographic variables*. The general problem in designing equivalence scales is to achieve the adjustment of observed income to take account of demographic differences in household composition. Several ways exist to do this, and these are now discussed.

The first approach to equivalence scales is based on the concept of *minimum needs*. A bundle of goods and services that is seen as representing the minimum needs for the household is identified. The exact bundle will differ among households of varying size, but it typically involves only very basic commodities. The cost of this bundle for families with different compositions is then calculated and the ratio of these costs for different families provides the equivalence scale. The first application of this approach was by Seebohm Rowntree in 1901 in his pioneering study of poverty. The bundle of goods employed was just a minimum acceptable quantity of food, rent, and a small allowance for “household sundries.” The equivalence scale was constructed by assigning the expenditure for a two-adult household with no children the index of 100 and measuring costs for all other household compositions relative to this. The scale obtained from expenditures calculated by Rowntree is given in the first column of table 14.1. The

Table 14.1
Minimum needs equivalence scales

	Rowntree (1901)	Beveridge (1942)	US poverty scale (2003)
Single person	60	59	78
Couple	100	100	100
+1 Child	124	122	120
+2 Children	161	144	151
+3 Children	186	166	178
+4 Children	223	188	199

Sources: B. S. Rowntree (1901, *Poverty: A Study of Town Life*, Macmillan), W. H. Beveridge, (1942, *Social Insurance and Allied Services*, HMSO), US Bureau of the Census (2003, www.census.gov/hhes/poverty/threshld/thresh03.html).

interpretation of these figures is that the minimum needs of a couple with one child cost 24 percent more than for a couple with no children.

A similar approach was taken by William Beveridge in his 1942 construction of the expenditure requirements that provided the foundation for the introduction of social assistance in the United Kingdom. In addition to the goods in the bundle of Rowntree, Beveridge added fuel, light, and a margin for “inefficiency” in purchasing. Also the cost assigned to children increased with their age. The values of the Beveridge scale in the second column of table 14.1 are for children in the 5 to 10 age group.

The final column of the table is generated from the income levels that are judged to represent poverty in the United States for families with different compositions. The original construction of these poverty levels was undertaken by Mollie Orshansky in 1963. The method she used was to evaluate the cost of food for each family composition using the 1961 Economy Food Plan. Next it was observed that if expenditure on food, F , constituted a proportion θ of the family’s budget, then total needs would be $\left[\frac{1}{\theta}\right]F$. For a family of two, $\frac{1}{\theta}$ was taken as 3.7, and for a family of three or more, $\frac{1}{\theta}$ was 3. The exception to this process was to evaluate the cost for a single person as 80 percent of that of a couple. The minimum expenditures obtained have been continually updated, and the third column of the table gives the equivalence scale implied by the poverty line used in 2003.

Table 14.1 shows that these equivalence scales all assume that there are returns to scale in household size so that, for example, a family of two adults does not require twice the income of a single person. Observe also that the US poverty scale is relatively generous for a single person compared to the other two scales. The fact that the

single-person value was constructed in a different way from the other values for the poverty scale (as a fixed percentage of that for a couple rather than as a multiple of food costs) has long been regarded as a contentious issue. Furthermore only for the Beveridge scale is the cost of additional children constant. The fact that the cost of children is nonmonotonic for the poverty scale is a further point of contention.

There are three major shortcomings of this method of computing equivalence scales. First, by focusing on the cost of meeting a minimum set of needs, they are inappropriate for applying to incomes above the minimum level. Second, they are dependent on an assessment of what constitutes minimum needs—and this can be contentious. Most important, the scales do not take into account the process of optimization by the households. The consequence of optimization is that as income rises, substitution between goods can take place, and the same relativities need no longer apply. Alternative methods of constructing equivalence scales that aim to overcome these difficulties are now considered.

In a similar way to the Orshansky construction of the US poverty scale, the Ernst Engel approach to equivalence scales is based on the hypothesis that the welfare of a household can be measured by the proportion of its income that is spent on food. This is a consequence of Engel's law, which asserts that the share of food in expenditure falls as income rises. If this is accepted, equivalence scales can be constructed for households of different compositions by calculating the income levels at which their expenditure share on food is equal. This is illustrated in figure 14.1 in which the expenditure share on food, as a function of income, is shown for two households with family compositions d^1 and d^2 . For example, d^1 may refer to a couple and d^2 to a couple with one child. Incomes M^1 and M^2 lead to the same expenditure share, s , and so are equivalent for the Engel method. The equivalence scale is then formed from the ratio $\frac{M^2}{M^1}$.

Although Engel's law may be empirically true, it does not necessarily provide a basis for making welfare comparisons, since it leaves unexplored the link between household composition and food expenditure. In fact there is ground for believing that the Engel method overestimates the cost of additional children because a child is largely a food-consuming addition to a household. If this is correct, a household compensated sufficiently to restore the share of food in its expenditure to its original level after the addition of a child would have been overcompensated with respect to other commodities. The approach of Engel has been extended to the more general isoprop method in which the expenditure shares of a basket of goods, rather than simply food, becomes the basis for the construction of scales. However, considering a basket of goods does not overcome the basic shortcomings of the Engel method.




Figure 14.1
Construction of Engel scale

A further alternative is to select for attention a set of goods that are consumed only by adults, termed “adult goods,” and such that the expenditure on them can be treated as a measure of welfare. Typical examples of such goods that have been used in practice are tobacco and alcohol. If these goods have the property that changes in household composition only affect their demand via an income effect (so changes in household composition do not cause substitution between commodities), then the extra income required to keep their consumption constant when household composition changes can be used to construct an equivalence scale. The use of adult goods to construct an equivalence scale is illustrated in figure 14.2. On the basis that they generate the same level of demand, \bar{x} , as family composition changes, the income levels M^1 and M^2 can be classed as equivalent, and the equivalence scale can be constructed from their ratio.

There are also a number of difficulties with this approach. It rests on the hypotheses that consumption of adult goods accurately reflects welfare and that household composition affects the demand for these goods only via an income effect. Furthermore the ratio of M^1 to M^2 will depend on the level of demand chosen for the comparison except in the special case where the demand curves are straight lines through the origin. The ratios may also vary for different goods. This leads into a further problem of forming some average ratio out of the ratios for the individual goods.




Figure 14.2
Adult good equivalence scale

All of the methods described so far have attempted to derive the equivalence scale from an observable proxy for welfare. A general approach that can, in principal, overcome the problems identified in the previous methods is illustrated in figure 14.3. To understand this figure, assume that there are just two goods available. The outer indifference curve represents the consumption levels of these two goods necessary for a family of composition d^2 to obtain welfare level U^* , and the inner indifference curve the consumption requirements for a family with composition d^1 to obtain the same utility. The extent to which the budget line has to be shifted outward to reach the higher curve determines the extra income required to compensate for the change in family structure. This construction incorporates both the potential change in preferences as family composition changes and the process of optimization subject to budget constraint by the households.

To formalize this process, let the household have preferences described by the utility function $U(x_1, x_2; d)$, where x_i is the level of consumption of good i and d denotes information on family composition. For example, d will describe the number of adults, the number and ages of children, and any other relevant information. The consumption plan needed to attain a given utility level, U , at least cost is the solution to

$$\min_{\{x_1, x_2\}} p_1 x_1 + p_2 x_2 \quad \text{subject to} \quad U(x_1, x_2; d) \geq U^*. \quad (14.1)$$




Figure 14.3
General equivalence scale

Denoting the (compensated) demand for good i by $x_i(U^*, d)$, the minimum cost of attaining utility U with characteristics d is then given by

$$M(U^*, d) = p_1 x_1(U^*, d) + p_2 x_2(U^*, d). \quad (14.2)$$

The equivalent incomes at utility U^* for two households with compositions d^1 and d^2 are then given by $M(U^*, d^1)$ and $M(U^*, d^2)$. The equivalence scale is derived by computing their ratio. The important point obtained by presenting the construction in this way is the observation that the equivalence scale will generally depend on the level of utility at which the comparison is made. If it does, there can be no single equivalence scale that works at all levels of utility.

The construction of an equivalence scale from preferences makes two further issues apparent. First, the minimum needs and budget share approaches do not take account of how changes in family structure may shift the indifference map. For instance, the pleasure of having children may raise the utility obtained from any given consumption plan. With the utility approach it then becomes cheaper to attain each indifference curve, so the value of the equivalence scale falls as family size increases. This conclusion, of course, conflicts with the basic sense that it is more expensive to support a larger family.

The second problem centers around the use of a household utility function. Many economists would argue that a household utility function cannot exist; instead, they would observe that households are composed of individuals with individual preferences. Under the latter interpretation, the construction of a household utility function suffers from the difficulties of preference aggregation identified by Arrow's Impossibility Theorem. Among the solutions to this problem now being investigated is to look within the functioning of the household and to model its decisions as the outcome of an efficient resource allocation process.

14.4 Inequality Measurement

Inequality is a concept that has immediate intuitive implications. The existence of inequality is easily perceived: differences in living standards between the rich and poor are only too obvious both across countries and, sometimes to a surprising extent, within countries. The obsession of the media with wealth and celebrity provides a constant reminder of just how rich the rich can be. An increase in inequality can also be understood at a basic level. If the rich become richer, and the poor become poorer, then inequality must have increased.

The substantive economic questions about inequality arise when we try to move beyond these generalizations to construct a quantitative measure of inequality. Without a quantitative measure it is not possible to provide a precise answer to questions about inequality. For example, a measure is required to determine which of a range of countries has the greatest level of inequality and to determine whether inequality has risen or fallen over time.

What an inequality measure must do is to take data on the distribution of income and generate a single number that captures the inequality in that distribution. A first approach to constructing such a measure is to adopt a standard statistical index. We describe the most significant of these indexes. Looking at the statistical measures reveals that there are properties, particularly how the measure is affected by transfers of income between households, that we may wish an inequality measure to possess. These properties can also be used to assess the acceptability of alternative measures. It is also shown that implicit within a statistical measure are a set of welfare implications. Rather than just accept these implications, the alternative approach is explored of making the welfare assumptions explicit and building the inequality measure on them.

14.4.1 The Setting

The intention of an inequality measure is to assign a single number to an income distribution that represents the degree of inequality. This section sets out the notation employed for the basic information that is input into the measure and defines precisely what is meant by a measure.

We assume that there are H households and label these $h = 1, \dots, H$. The labeling of the households is chosen so that the lower is the label, the lower is the household's income. The incomes, M^h , then form an increasing sequence with

$$M^1 \leq M^2 \leq M^3 \leq \dots \leq M^H. \quad (14.3)$$

The list $\{M^1, \dots, M^H\}$ is the income distribution whose inequality we wish to measure. Given the income distribution, the mean level of income, μ , is defined by

$$\mu = \frac{1}{H} \sum_{h=1}^H M^h. \quad (14.4)$$

The purpose of an inequality measure is to assign a single number to the distribution $\{M^1, \dots, M^H\}$. Let $I(M^1, \dots, M^H)$ be an inequality measure. Then income distributions $\{\tilde{M}^1, \dots, \tilde{M}^H\}$ has greater inequality than distribution $\{\hat{M}^1, \dots, \hat{M}^H\}$ if $I(\tilde{M}^1, \dots, \tilde{M}^H) > I(\hat{M}^1, \dots, \hat{M}^H)$. Typically the inequality measure is constructed so that a value of 0 represents complete equality (the position where all incomes are equal) and a value of 1 represents maximum inequality (all income is received by just one household).

The issues that arise in inequality measurement are encapsulated in determining the form that the function $I(M^1, \dots, M^H)$ should take. We now investigate some alternative forms and explore their implications.

14.4.2 Statistical Measures

Under the heading of “statistical” fall inequality measures that are derived from the general statistical literature. That is, the measures have been constructed to characterize the distribution of a set of numbers without thought of any explicit economic application or motivation. Even so, the discussion will later show that these statistical measures make implicit economic value judgments. Accepting any one of these measures as the “correct” way to measure inequality means the acceptance of these implicit

assumptions. The measures that follow are presented in approximate order of sophistication. Each is constructed to take a value between 0 and 1, with a value of 0 occurring when all households have identical income levels.

Probably the simplest conceivable measure, the *range* calculates inequality as being the difference between the highest and lowest incomes expressed as a proportion of total income. As such, it is a very simple measure to compute. The definition of the range, R , is

$$R = \frac{M^H - M^1}{H\mu}. \quad (14.5)$$

The division by $H\mu$ in (14.5) is a normalization that ensures the index is independent of the scale of incomes (or the units of measurement of income). Any index that has this property of independence is called a *relative index*.

As an example of the use of the range, consider the income distribution {1, 3, 6, 9, 11}. For this distribution $\mu = 6$ and

$$R = \frac{11 - 1}{5 \times 6} = 0.3333. \quad (14.6)$$

The failure of the range to take account of the intermediate part of the distribution can be illustrated by taking income from the second household in the example and giving it to the fourth to generate new income distribution {1, 1, 6, 11, 11}. This new distribution appears to be more unequal than the first, yet the value of the range remains at $R = 0.3333$.

Given the simplicity of its definition, it is not surprising that the range has deficiencies. Most important, the range takes no account of the dispersion of the income distribution between the highest and the lowest incomes. Consequently it is not sensitive to any features of the income distribution between these extremes. For instance, an income distribution with most of the households receiving close to the maximum income would be judged just as unequal as one in which most received the lowest income. An ideal measure should possess more sensitivity to the value of intermediate incomes than the range.

The *relative mean deviation*, D , takes account of the deviation of each income level from the mean so that it is dependent on intermediate incomes. It does this by calculating the absolute value of the deviation of each income level from the mean and then summing. This summation process gives equal weight to deviations both above and below the mean and implies that D is linear in the size of deviations. Formally, D is defined by

$$D = \frac{\sum_{h=1}^H |\mu - M^h|}{2[H-1]\mu}. \quad (14.7)$$

The division by $2[H-1]\mu$ again ensures that D takes values between 0 and 1.

The advantage of the relative mean deviation over the range is that it takes account of the entire income distribution and not just the end points. Taking the example used for the range, the inequality in the distribution $\{1, 3, 6, 9, 11\}$ as measured by D is

$$D = \frac{|-5| + |-3| + |0| + |3| + |5|}{2 \times 4 \times 6} = 0.3333, \quad (14.8)$$

and the inequality of $\{1, 1, 6, 11, 11\}$ is

$$D = \frac{|-5| + |-5| + |0| + |5| + |5|}{2 \times 4 \times 6} = 0.4167. \quad (14.9)$$

Unlike the range, the relative mean deviation measures the second distribution as having more inequality. Due to the division by $2[H-1]\mu$ it is easily seen that $D = 1$ with the maximum inequality distribution $\{0, 0, 0, 0, 30\}$ where all income is received by just one household.

Although it does take account of the entire distribution of income, the linearity of D has the implication that it is insensitive to transfers from richer to poorer households when the households involved in the transfer remain on the same side of the mean income level. To see an example of this, assume that the mean income level is $\mu = \$20,000$. Now take two households with incomes \$25,000 and \$100,000. Transferring \$4,000 from the poorer of these two households to the richer, so that the income levels become \$21,000 and \$104,000, does not change the value of D —one term in the summation rises by \$4,000 and the other falls by \$4,000. (Notice that if the two households were on different sides of the mean, then a similar transfer would raise two terms in the summation by \$4,000 and increase inequality.) The fact that D can be insensitive to transfers seems unsatisfactory, since it is natural to expect that a transfer from a poorer household to a richer one should raise inequality.

This line of reasoning is enshrined in the *Pigou–Dalton Principle of Transfers*, which is a central concept in the theory of inequality measurement. The basis of this principle is precisely the requirement that any transfer from a poor household to a rich one must increase inequality regardless of where the two households are located in the income distribution.

Definition 14.1 (Pigou–Dalton Principle of Transfers) The inequality index must decrease if there is a transfer of income from a richer household to a poorer household that preserves the ranking of the two households in the income distribution and leaves total income unchanged.

Any inequality measure that satisfies this principle is said to be *sensitive to transfers*. The Pigou–Dalton Principle is generally viewed as a feature that any acceptable measure of inequality should possess and is therefore expected in an inequality measure. Neither the range nor the relative mean deviation satisfy this principle.

The reason why D is not sensitive to transfers is its linearity in deviations from the mean. The removal of the linearity provides the motivation for considering the *coefficient of variation*, which is defined using the sum of squared deviations. The procedure of forming the square places more weight on incomes that are further away from the mean and so introduces a sensitivity to transfers. The coefficient of variation, C , is defined by

$$C = \frac{\sigma}{\mu [H - 1]^{1/2}}, \quad (14.10)$$

where $\sigma^2 = \frac{\sum_{h=1}^H [M^h - \mu]^2}{H}$ is the variance of the income distribution, so σ is its standard deviation. The division by $\mu [H - 1]^{1/2}$ ensures the C lies between 0 and 1. For the income distribution $\{1, 3, 6, 9, 11\}$, $\sigma^2 = \frac{[-5]^2 + [-3]^2 + [0]^2 + [3]^2 + [5]^2}{5} = 13.6$, so

$$C = \frac{[13.6]^{1/2}}{6[4]^{1/2}} = 0.3073, \quad (14.11)$$

and for $\{1, 1, 6, 11, 11\}$, $\sigma^2 = 20.0$, giving

$$C = \frac{[20]^{1/2}}{6[4]^{1/2}} = 0.3727. \quad (14.12)$$

To see that the coefficient of variation satisfies the Pigou–Dalton Principle, consider a transfer of an amount of income $d\varepsilon$ from household i to household j , with the households chosen so that $M^i < M^j$. Then

$$\frac{dC}{d\varepsilon} = \frac{1}{\mu [H - 1]^{1/2}} \frac{d\sigma}{d\varepsilon} = \frac{M^j - M^i}{\sigma H \mu [H - 1]^{1/2}} > 0, \quad (14.13)$$

so the transfer from the poorer household to the richer household decreases measured inequality as required by the Pigou–Dalton Principle. It should be noted that the value of

the change in C depends on the difference between the incomes of the two households. This has the consequence that a transfer of \$100 of income from a household with an income of \$1,000,100 to one with an income of \$999,900 produces the same change in C as a transfer of \$100 between households with incomes \$1,100 and \$900. Most interpretations of equity would suggest that the latter transfer should be of greater consequence for the index because it involves two households of relatively low incomes. This reasoning suggests that satisfaction of the Pigou–Dalton Principle may not be a sufficient requirement for an inequality measure; the manner in which the measure satisfies it may also matter.

Before moving on to further inequality measures, it is worth describing the Lorenz curve. The Lorenz curve is a helpful graphical device for presenting a summary representation of an income distribution, and it has played an important role in the measurement of inequality. Although not strictly an inequality measure as defined above, Lorenz curves are considered because of their use in illustrating inequality and the central role they play in the motivation of other inequality indexes.

The Lorenz curve is constructed by arranging the population in order of increasing income and then graphing the proportion of income going to each proportion of the population. The graph of the Lorenz curve therefore has the proportion of population on the horizontal axis and the proportion of income on the vertical axis. If all households in the population had identical incomes the Lorenz curve would then be the diagonal line connecting the points (0, 0) and (1, 1). If there is any degree of inequality, the ordering in which the households are taken ensures that the Lorenz curve lies below the diagonal since, for example, the poorest half of the population must have less than half the total income.

To see how the Lorenz curve is plotted, consider a population of 10 with income distribution {1, 2, 3, 4, 5, 6, 7, 8, 9, 10}. The total quantity of income is 55, so the first household (which represents 10 percent of the population) receives $\frac{1}{55} \times 100$ percent of the total income. This is the first point plotted in the lower left corner of figure 13.4. Taking the two lowest income households (which are 20 percent of the population), we have their combined income as $\frac{3}{55} \times 100$ percent of the total. Adding the third household awards 30 percent of the population $\frac{6}{55} \times 100$ percent of total income. Proceeding in this way, we plot the ten points in the figure. Joining them gives the Lorenz curve. In summary, the larger the population, the smoother is this curve.

The Lorenz curve can be employed to unambiguously rank some income distributions with respect to income inequality. This claim is based on the fact that a transfer of income from a poor household to a richer household moves the Lorenz curve farther away from the diagonal. (This can be verified by re-plotting the Lorenz curve in figure 14.4 for





Figure 14.4
Construction of a Lorenz curve

the income distribution $\{1, 1, 3, 4, 5, 6, 7, 8, 10, 10\}$, which is the same as the original except for the transfer of one unit from household 2 to household 9.) Because of this property the Lorenz curve satisfies the Pigou–Dalton Principle, with the curve farther from the diagonal indicating greater inequality.

Income distributions that can, and cannot, be ranked are displayed in figure 14.5. In the left-hand panel, the Lorenz curve for income distribution B lies entirely outside that for income distribution A . In such a case distribution B unambiguously has more inequality than A . One way to see this is to observe that distribution B can be obtained from distribution A by transferring income from poor households to rich households. Applying the Pigou–Dalton Principle, we see that this raises inequality. If the Lorenz curves representing the distributions A and B cross, it is not possible to obtain an unambiguous conclusion by the Lorenz curve alone. The Lorenz curve therefore provides only a partial ranking of income distributions. Despite this limitation the Lorenz curve is still a popular tool in applied economics, since it presents very convenient and easily interpreted visual summary of an income distribution.

The next measure, the *Gini*, has been the subject of extensive attention in discussions of inequality measurement and has been much used in applied economics. The Gini, G , can be expressed by considering all possible pairs of incomes and out of each pair selecting the minimum income level. Summing the minimum income levels and

**Figure 14.5**

Lorenz curves as an incomplete ranking

dividing by $H^2\mu$ to ensure a value between 0 and 1 provides the formula for the Gini:

$$G = 1 - \frac{1}{H^2\mu} \sum_{i=1}^H \sum_{j=1}^H \min\{M^i, M^j\}. \quad (14.14)$$

It should be noted that in the construction of this measure, each level of income is compared to itself as well as all other income levels. For example, if there are three income levels $\{3, 5, 10\}$, the value of the Gini is

$$\begin{aligned} G &= 1 - \frac{1}{3^2 \times 6} \left[\begin{array}{l} \min\{3, 3\} + \min\{3, 5\} + \min\{3, 10\} \\ + \min\{5, 3\} + \min\{5, 5\} + \min\{5, 10\} \\ \min\{10, 3\} + \min\{10, 5\} + \min\{10, 10\} \end{array} \right] \\ &= 1 - \frac{1}{54} \left[3 + 3 + 3 + 3 + 5 + 5 + 3 + 5 + 10 \right] \\ &= 0.259. \end{aligned} \quad (14.15)$$

By counting the number of times each income level appears, we can also write the Gini as

$$\begin{aligned} G &= 1 - \frac{1}{H^2\mu} \left[(2H - 1) M^1 + (2H - 3) M^2 \right. \\ &\quad \left. + (2H - 5) M^3 + \dots + M^H \right]. \end{aligned} \quad (14.16)$$




Figure 14.6
Relating Gini to Lorenz

This second form of the Gini makes its computation simpler but hides the construction behind the measure.

The Gini also satisfies the Pigou–Dalton Principle. This can be seen by considering a transfer of income of size $\Delta > 0$ from household i to household j , with the households chosen so that $M^j > M^i$. From the ranking of incomes this implies $j > i$. Then

$$\Delta G = \frac{2}{H^2 \mu} (j - i) \Delta > 0, \quad (14.17)$$

as required. In the case of the Gini, the effect of the transfer of income on the measure depends only on the locations of i and j in the income distribution. For example, a transfer from the household at position $i = 1$ to the household at position $j = 11$ counts as much as one from position $i = 151$ to position $j = 161$. It might be expected that an inequality should be more sensitive to transfers between households low in the income distribution.

There is an important relationship between the Gini and the Lorenz curve. As shown in figure 14.6, the Gini is equal to the area between the Lorenz curve and the line of equality as a proportion of the area of the triangle beneath the line of equality. As the area of the triangle is $\frac{1}{2}$, the Gini is twice the area between the Lorenz curve and the equality line. This definition makes it clear that the Gini, in common with R , C , and D , can be used to rank distributions when the Lorenz curves cross, since the relevant area is always well defined. All these measures provide a stronger ranking of income distributions than the Lorenz curve; hence they must each impose additional restrictions that allow a comparison to be made between distributions even when their Lorenz curves cross.

A final statistical measure that displays a different form of sensitivity to transfers is the *Theil entropy measure*. This measure is drawn from information theory and is used in that context to measure the average information content of a system of information. The definition of the Theil entropy measure, T , is given by

$$\begin{aligned} T &= \frac{1}{\log H} \sum_{h=1}^H \frac{M^h}{H\mu} \left[\log \frac{M^h}{H\mu} - \log \frac{1}{H} \right] \\ &= \frac{1}{H \log H} \sum_{h=1}^H \frac{M^h}{\mu} \log \frac{M^h}{\mu}. \end{aligned} \tag{14.18}$$

The effect of an income transfer, $d\epsilon$, between households i and j on the entropy index is given by

$$\frac{dT}{d\epsilon} = \frac{1}{H \log H} \log \frac{M^j}{M^i} < 0, \tag{14.19}$$

so the entropy measure also satisfies the Pigou–Dalton Principle. For the Theil entropy measure, the change is dependent on the relative incomes of the two households involved in the transfer. This provides an alternative form of sensitivity to transfers.

14.4.3 Inequality and Welfare

The analysis of the statistical measures of inequality has made reference to “acceptable” criteria for a measure to possess. One of these was made explicit in the Pigou–Dalton Principle, while other criteria relating to additional desirable sensitivity properties have been implicit in the discussion. To be able to say that something is acceptable or not implies that there is some notion of distributive justice or social welfare underlying the judgment. It is then interesting to consider the relationship between inequality measures and welfare.

The first issue to address is the extent to which income distributions can be ranked in terms of welfare with minimal restrictions imposed on the social welfare function. To investigate this, let the level of social welfare be determined by the function $W = W(M^1, \dots, M^H)$. It is assumed that this social welfare function is symmetric and concave. Symmetry means that the level of welfare is unaffected by changing the ordering of the households. This is just a requirement that all households are treated equally. Concavity ensures that the indifference curves of the welfare function have

the standard shape with mixtures preferred to extremes. This assumption imposes a concern for equity on the welfare function.

The critical theorem relating the ranking of income distributions to social welfare is now given.

Theorem 14.2 (Atkinson) Consider two distributions of income with the same mean. If the Lorenz curves for these distributions do not cross, every symmetric and concave social welfare function will assign a higher level of welfare to the distribution whose Lorenz curve is closest to the main diagonal.

The proof of this theorem is very straightforward. Since the welfare function is symmetric and concave, it follows that $\frac{\partial W}{\partial M^i} \geq \frac{\partial W}{\partial M^j}$ if $M^i < M^j$. Hence the marginal social welfare of income is greater for a household lower in the income distribution. If the two Lorenz curves do not cross, the income distribution represented by the inner one (that closest to the main diagonal) can be obtained from that of the outer one by transferring income from richer to poorer households. Since the marginal social welfare of income to the poorer households is never less than that from richer, this transfer must raise welfare as measured by any symmetric and concave social welfare function.

The converse of this theorem is that if the Lorenz curves for two distributions cross, then two symmetric and concave social welfare functions can be found that will rank the two distributions differently. This is because the income distributions of two Lorenz curves that cross are not related by simple transfers from rich to poor. So, if the Lorenz curves do cross, the income distributions cannot be unambiguously ranked without specifying the social welfare function.

Taken together, the theorem and its converse show that the Lorenz curve provides the most complete ranking of income distributions that is possible without our making assumptions on the form of the social welfare function other than symmetry and concavity. To achieve a complete ranking when the Lorenz curves cross requires restrictions to be placed on the structure of the social welfare function. In addition any measure of inequality is necessarily stronger than the Lorenz curve because it generates a complete ranking of distributions. This is true of all the statistical measures, which is why it can be argued that they all carry implicit welfare judgments.

This argument can be taken a stage further. It is in fact possible to construct the social welfare function that is implied by an inequality measure. To see how this can be done, consider the Gini. Assume that the total amount of income available is constant. Any redistribution of this that leaves the Gini unchanged must leave the implied level of welfare unchanged. A redistribution of income will not affect the Gini if the term




Figure 14.7
Gini social welfare

$[2H - 1]M^1 + \dots + M^H]$ remains constant. The welfare function must thus be a function of this expression. Furthermore the Gini is defined to be independent of the total level of income, but a welfare function will increase if total income rises and distribution is unaffected. This can be incorporated by not dividing through by the mean level of income. Putting these arguments together, the welfare function implied by the Gini is given by

$$W_G(M) = \frac{1}{H^2} [2H - 1]M^1 + [2H - 3]M^2 + \dots + M^H]. \quad (14.20)$$

The form of $W_G(M)$ is interesting, since it shows that the Gini implies a social welfare function that is linear in incomes. It also shows a clear structure of increasing welfare weights for lower income consumers. The welfare function further has indifference curves that are straight lines above and below the line of equal incomes but kinked on this line. This is illustrated in figure 14.7.

In the same way a welfare function can be constructed for all the statistical measures. Therefore acceptance of the measure is acceptance of the implied welfare function. As shown by the linear social indifference curves and increasing welfare weights for the Gini social welfare function, the implied welfare functions can have a very restrictive form. We do not need to merely accept such welfare restrictions. The fact that each inequality measure implies a social welfare function suggests that the relationship can be inverted to move from a social welfare function to an inequality measure. By assuming

a social welfare function at the outset, it is possible to make welfare judgments explicit and, by deriving the inequality measure from the social welfare function, to ensure that these judgments are incorporated in the inequality measure.

To implement this approach, assume that the social welfare function is utilitarian with

$$W = \sum_{h=1}^H U(M^h). \quad (14.21)$$

The household utility of income function, $U(M)$, is taken to satisfy the conditions that $U'(M) > 0$ and $U''(M) < 0$. The utility function $U(M)$ can either be the households' true cardinal utility function or be chosen by the policy analyst as in the evaluation of the utility of income to each household. In this second interpretation, since social welfare is obtained by summing the individual utilities, the importance given to equity can be captured in the choice of $U(M)$. This is because increasing the concavity of the utility function places a relatively higher weight on low incomes in the social welfare function.

A measure of inequality can be constructed from the social welfare function by defining M_{EDE} as the solution to

$$\sum_{h=1}^H U(M^h) = HU(M_{EDE}). \quad (14.22)$$

M_{EDE} is called the *equally distributed equivalent* income and is the level of income that, if given to all households, would generate the same level of social welfare as the initial income distribution. Using M_{EDE} , the *Atkinson* measure of inequality is defined by

$$A = 1 - \frac{M_{EDE}}{\mu}. \quad (14.23)$$

For the case of two households the construction of M_{EDE} is illustrated in figure 14.8. The initial income distribution is given by $\{M^1, M^2\}$, and this determines the relevant indifference curve of the social welfare function. M_{EDE} is found by moving around this indifference curve to the 45 degree line where the two households' incomes are equal. The figure makes clear that because of the concavity of the social indifference curve, M_{EDE} is less than the mean income, μ . This fact guarantees that $0 \leq A \leq 1$. Furthermore, for a given level of mean income, a more diverse income distribution will achieve a lower social indifference curve and be equivalent to a lower M_{EDE} .

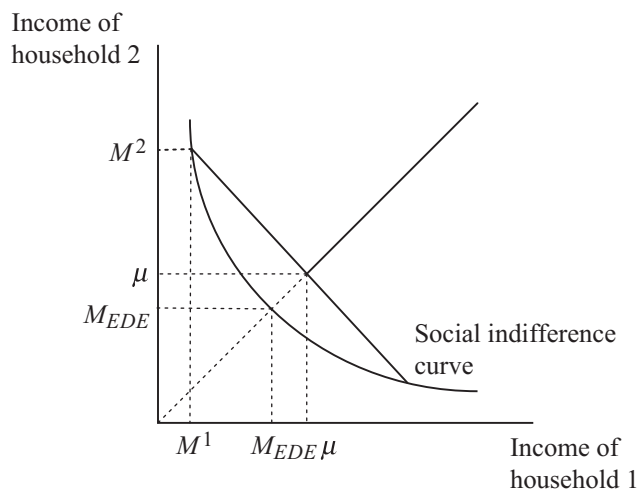


Figure 14.8
Equally distributed equivalent income

The flexibility in this measure lies in the freedom of choice of the household utility of income function. Given the assumption of a utilitarian social welfare function, it is the household utility that determines the importance attached to inequality by the measure. One commonly used form of utility function is

$$U(M) = \frac{M^{1-\varepsilon}}{1-\varepsilon}, \quad \varepsilon \neq 1, \quad (14.24)$$

which allows the welfare judgments of the policy analyst to be contained in the chosen value of the parameter ε . The value of ε determines the degree of concavity of the utility function: it becomes more concave as ε increases. An increase in concavity raises the relative importance of low incomes because it causes the marginal utility of income to decline at a faster rate. The utility function is isoelastic, and concave if $\varepsilon \geq 0$. When $\varepsilon = 1$, $U(M) = \log(M)$, and when $\varepsilon = 0$, $U(M) = M$.

The Atkinson measure can be illustrated using the example of the income distribution $\{1, 3, 6, 9, 11\}$. If $\varepsilon = \frac{1}{2}$ the household utility function is $U = 2M^{1/2}$ so the level of social welfare is

$$W = 2 \times 1^{1/2} + 2 \times 3^{1/2} + 2 \times 6^{1/2} + 2 \times 9^{1/2} + 2 \times 11^{1/2} = 22.996. \quad (14.25)$$

The equally distributed equivalent income then solves

$$5 \times \left[2 \times (M)^{1/2} \right] = 22.996, \quad (14.26)$$

so $M_{EDE} = 5.2882$. This gives the value of the Atkinson measure as

$$A = 1 - \frac{5.2882}{6} = 0.1186. \quad (14.27)$$

14.4.4 An Application

As has been noted in the discussion, inequality measures are frequently used in practical policy analysis. Table 14.2 summarizes the results of an OECD study into the change in inequality over time in a wide range of countries. This is undertaken by calculating inequality at two points in time and determining the percentage change in the measure. Inequality is calculated for income prior to taxes and transfers, and for income after taxes and transfers. The difference between the inequality levels in these two situations gives an insight into the extent to which the tax and transfer system succeeds in redistributing income.

Looking at the results, in all cases inequality is smaller after taxes and transfers than before, so the tax systems in the countries studied are redistributive. For instance, in Denmark inequality is 0.0420 when measured by Gini before taxes and transfers but only 0.0217 after. The second general message of the results is that inequality has tended to rise in these countries—only in three cases has it been reduced, and in every case this is after taxes and transfers.

It is also interesting to look at the rankings of inequality and changes in inequality under the different measures. If there is general agreement for different measures, then we can be reassured that the choice of measure is not too critical for what we observe. For the level of inequality all four measures are in agreement for both the before-tax and after-tax cases except for the *SCV* (squared coefficient of variation), which reverses the after-taxes and transfers ranking of Denmark and Sweden, and the Atkinson, which reverses the before-taxes and transfers ranking of Denmark and the United States. For these four measures there is a considerable degree of consistency in the rankings. Taking the majority opinion, observe that before taxes and transfers the ranking (with the highest level of inequality first) is Italy, Sweden, United States, Denmark, and Japan. After the operation of taxes and transfers this ranking becomes Italy, United States, Japan, Sweden, and Denmark. This change in rankings is evidence of the highly redistributive tax and transfer systems operated in the Nordic countries.

The rankings for the change in inequality are not quite as consistent across the four measures, but there is still considerable agreement. The majority order for the before

Table 14.2
Inequalities before and after taxes and transfers

Measure	SCV		Gini		Atkinson	
	Before	After	Before	After	Before	After
Denmark 1994	0.671	0.229	0.420	0.217	0.209	0.041
% Change 1983–1994	4.9	2.0	11.2	–4.9	25.3	–11.1
Italy 1993	1.19	0.584	0.570	0.345	0.299	0.105
% Change 1984–1993	59.6	44.7	20.8	12.8	43.8	33.1
Japan 1994	0.536	0.296	0.340	.265	0.124	0.059
% Change 1984–1994	33.7	21.7	14.0	4.9	47.3	10.9
Sweden 1995	0.894	0.217	0.487	0.230	0.262	0.049
% Change 1975–1995	49.1	36.9	17.2	–1.0	28.7	3.2
United States 1995	0.811	0.441	0.455	0.344	0.205	0.100
% Change 1974–1995	32.0	25.4	13.1	10.0	19.6	18.6

Source: OECD ECO/WKP(98)2.

Notes: The squared coefficient of variation (*SCV*) is defined by $SCV = [H - 1]C$. For the Atkinson measure, $\varepsilon = 0.5$.

taxes and transfers case (with the greatest increase in inequality first) is Italy, Sweden, Japan, United States, and Denmark. The Atkinson measure places Japan at the top and reverses Denmark and the United States. For the after-taxes and transfers ranking, the Gini and the Atkinson measures produce the same ranking, but the *SCV* places Sweden above the United States and Japan. But what is clear is the general agreement on an increase in inequality.

The review of this application has shown that the different measure can produce a fairly consistent picture about ranking by inequality, about the changes in inequality, and on the effect of taxes and transfers. Despite the differences emphasized in the analysis of the measures, when put into practice in this way, the differences need not lead to widespread disagreement among the measures. In fact a fairly harmonious picture can emerge.

14.5 Poverty

The essential feature of poverty is the possession of fewer resources than are required to achieve an acceptable standard of living. What constitutes poverty can be understood in the same intuitive way as what constitutes inequality, but similar issues about the correct measure arise again once we attempt to provide a quantification. This section

first discusses concepts of poverty and the poverty line, and then proceeds to review a number of common poverty measures.

14.5.1 Poverty and the Poverty Line

Before measuring poverty, it is first necessary to define it. It is obvious that poverty refers to a situation involving a lack of income and a consequent low level of consumption and welfare. What is not so clear is the standard against which the level of income should be judged. Two possibilities arise in this context: an absolute conception of poverty and a relative one. The distinction between these has implications for changes in the level of poverty over time and the success of policy in alleviating poverty.

The concept of *absolute poverty* assumes that there is some fixed minimum level of consumption (and hence of income) that constitutes poverty and is independent of time or place. Such a minimum level of consumption can be a diet that is just sufficient to maintain health and limited housing and clothing. Under the concept of absolute poverty, if the incomes of all households rise, there will eventually be no poverty. Although a concept of absolute poverty was probably implicit in early studies of poverty, such as that of Rowntree in 1901, the appropriateness of absolute poverty has since generally been rejected. In its place has been adopted the notion of relative poverty.

Relative poverty is not a recent concept. Even in 1776 Adam Smith was defining poverty as the lack of necessities, where necessities are defined as “what ever the custom of the country renders it indecent for creditable people, even of the lowest order, to be without.” This definition makes it clear that relative poverty is defined in terms of the standards of a given society at a given time and that the level that represents poverty rises as does the income of that society. Operating under a relative standard, it becomes much more difficult to eliminate poverty. Relative poverty has also been defined in terms of the ability to “participate” in society. Poverty then arises whenever a household possesses insufficient resources to allow it to participate in the customary activities of its society.

The starting point for the measurement of poverty is to set a poverty line that separates those viewed as living in poverty from those who are not. Of course, this poverty line applies to the incomes levels after application of an equivalence scale. Whether poverty is viewed as absolute or relative matters little for setting a poverty line at any particular point in time (though advocates of an absolute poverty concept may choose to set it lower). Where the distinction matters is whether and how the poverty line is adjusted over time. If an absolute poverty standard were adopted, then there would be no revision.

Conversely, with relative poverty the level of the line would rise or fall in line with average incomes.

In practice, poverty lines have often been determined by following the minimum needs approach that was discussed in connection with equivalence scales. As noted in section 14.3, this is the case with the US poverty line that was fixed in 1963 and has since been updated annually. As the package of minimum needs has not changed, the underlying concept is that of an absolute poverty measure. In the United Kingdom the poverty line has been taken as the level of income that is 120 or 140 percent of the minimum supplementary benefit level. As this level of benefit is determined by minimum needs, a minimum needs poverty line is implied. In addition benefits have risen with increases in average income, so causing the poverty line to rise. The UK poverty line thus represents the use of a relative concept of poverty.

The assumption that there is a precise switch between poverty and nonpoverty as the poverty line is crossed is very strong. It is much more natural for there to be a gradual move out of poverty as income increases. The precision of the poverty line may also lead to difficulty in determining where it should lie if the level of poverty is critically dependent on the precise choice. Both of these difficulties can be overcome by observing that often it is not the precise level of poverty that matters but changes in the level of poverty over time and across countries. In these instances the poverty value is not too important but only the rankings. This suggests the procedure of calculating poverty for a range of poverty lines. If poverty is higher today for all poverty lines than it was yesterday, then it seems unambiguous that poverty has risen. In this sense the poverty line may not actually be of critical importance for the uses to poverty measurement is often put. An application illustrating this argument is given below.

14.5.2 Poverty Measures

The poverty line is now taken as given, and we proceed to discuss alternative measures of poverty. The basic issue in this discussion is how best to combine two pieces of information (how many households are poor, and how poor they are) into a single quantitative measure of poverty. By describing a number of measures, the discussion will draw out the properties that are desirable for a poverty measure to possess.

Throughout the discussion the poverty line is denoted by the income level z such that a household with an income level below or equal to z is classed as living in poverty. For a household with income M^h the *income gap* measures how far their income is below the poverty line. Denoting the income gap for household h by g_h , it follows that $g_h = z - M^h$. Given the poverty line z and an income distribution $\{M^1, \dots, M^H\}$,

where $M^1 \leq M^2 \leq \dots \leq M^H$, the number of households in poverty is denoted by q . The value of q is defined by the facts that the income of household q is on or below the poverty line, so $M^q \leq z$, but that of the next household is above $M^{q+1} > z$.

The simplest measure of poverty is the *head-count ratio*, which determines the extent of poverty by counting the number of households whose incomes are not above the poverty line. Expressing the number as a proportion of the population, the head-count ratio is defined by

$$E = \frac{q}{H}. \quad (14.28)$$

This measure of poverty was first used by Rowntree in 1901 and has been employed in many subsequent studies. The major advantage of the head-count ratio is its simplicity of calculation.

The head-count ratio is clearly limited because it is not affected by how far below the poverty line the households are. For example, with a poverty line of $z = 10$ the income distributions $\{1, 1, 20, 40, 50\}$ and $\{9, 9, 20, 40, 50\}$ would both have a headcount ratio of $E = \frac{2}{5}$. A policy maker may well see these income distributions differently, since the income required to alleviate poverty in the second case (2 units) is much less than that required for the first (18 units). The head-count ratio is also not affected by any transfer of income from a poor household to one that is richer if both households remain on the same side of the poverty line. Even worse, observe that if we change the second distribution to $\{7, 11, 20, 40, 50\}$ the head-count ratio falls to $E = \frac{1}{5}$, so a regressive transfer has actually reduced the head-count ratio. This will happen whenever a transfer takes the income of the recipient of the transfer above the poverty line.

Only one of the two pieces of information on poverty are used in the head count. A measure that uses only information on how far below the poverty line are the incomes of the poor households is the *aggregate poverty gap*. This is defined as the simple sum of the income gaps of the households that are in poverty. Recalling that it is the first q households that are in poverty, the aggregate poverty gap is

$$V = \sum_{h=1}^q g_h. \quad (14.29)$$

The interpretation of this measure is that it is the additional income for the poor that is required to eliminate poverty. It provides some information but is limited by the fact that it is not sensitive to changes in the number in poverty. In addition the aggregate poverty gap gives equal weight to all income shortfalls regardless of how far they are from the

poverty line. It is therefore insensitive to transfers unless the transfer takes one of the households out of poverty. To see this latter point, for the poverty line of $z = 10$ the income distributions $\{5, 5, 20, 40, 50\}$ and $\{1, 9, 20, 40, 50\}$ have an aggregate poverty gap of $V = 10$. The distribution between the poor is somewhat different in the two cases.

One direct extension of the aggregate poverty gap is to adjust the measure by taking into account the number in poverty. The *income gap ratio* does this by calculating the aggregate poverty gap and then dividing by the number in poverty. Finally the value obtained is divided by the value of the poverty line, z , to obtain a measure whose value falls between 0 (the absence of poverty) and 1 (all households in poverty have no income):

$$I = \frac{1}{z} \frac{\sum_{h=1}^q g_h}{q}. \quad (14.30)$$

For the income distribution $\{1, 9, 20, 40, 50\}$, the income gap ratio when $z = 10$ is

$$I = \frac{1}{10} \frac{9 + 1}{2} = 0.5. \quad (14.31)$$

However, when this income distribution changes to $\{1, 10, 20, 40, 50\}$, so that only one household is now in poverty, the measure becomes

$$I = \frac{1}{10} \frac{9}{1} = 0.9. \quad (14.32)$$

This example reveals that the income gap ratio has the unfortunate property of being able to report increased poverty when the income of a household crosses the poverty line and the number in poverty is reduced.

These observations suggest that it is necessary to reflect more carefully on the properties that a poverty measure should possess. In 1976 Amartya K. Sen suggested that a poverty measure should have the following properties:

- Transfers of income between households above the poverty line should not affect the amount of poverty.
- If a household below the poverty line becomes worse off, poverty should increase.
- The poverty measure should be anonymous, meaning it should not depend on who is poor.
- A regressive transfer among the poor should raise poverty.

These are properties that have already been highlighted by the discussion. Two further properties were also proposed:

- The weight given to a household should depend on their ranking among the poor, meaning more weight should be given to those furthest below the poverty line.
- The measure should reduce to the headcount if all the poor have the same level of income.

One poverty measure that satisfies all of these conditions is the *Sen measure*

$$S = E \left[I + (1 - I) G_p \left(\frac{q}{q + 1} \right) \right], \quad (14.33)$$

where G_p is the Gini measure of income inequality among the households below the poverty line. This poverty measure combines a measure of the number in poverty (the head-count ratio), a measure of the shortfall in income (the income gap ratio), and a measure of the distribution of income among the poor (the Gini). Applying the Sen measure to the income distribution $\{1, 9, 20, 40, 50\}$, when $z = 10$, we have $E = \frac{2}{5}$ and $I = 0.5$. The Gini is calculated for the distribution of income of the poor $\{1, 9\}$, so $G_p = 1 - \frac{1}{2^2 \times 5} [3 \times 1 + 9] = \frac{4}{10}$. These values give

$$S = \frac{2}{5} \left[0.5 + (1 - 0.5) \frac{4}{10} \left(\frac{2}{2 + 1} \right) \right] = 0.2533. \quad (14.34)$$

In contrast, for the distribution $\{1, 10, 20, 40, 50\}$ that was judged worse using the income gap ratio, there is no inequality among the poor (since there is a single poor person), so the Sen measure is

$$S = \frac{1}{5} \left[0.9 + (1 - 0.9) 0 \left(\frac{1}{1 + 1} \right) \right] = 0.18, \quad (14.35)$$

which is simply the head-count ratio and records a lower level of poverty.

There is a further desirable property that leads into an alternative and important class of poverty measures. Consider a population that can be broken down into distinct subgroups. For instance, imagine dividing the population into rural and urban dwellers. The property we want is for the measure to be able to assign a poverty level for each of the groups and to aggregate these group poverty levels into a single level for the total society. Further we will also want the aggregate measure to increase if poverty rises in one of the subgroups and does not fall in any of the others. So, if rural poverty rises while urban poverty remains the same, aggregate poverty must rise. Any poverty measures that satisfies this condition are termed *subgroup consistent*.

Before introducing a form of measure that is subgroup consistent, it is worth providing additional discussion of the effect of transfers. The measures discussed so far have all had the property that the effect of a transfer has been independent of the income levels of the loser and gainer (except when the transfer was between households on different sides of the poverty line or changed the number in poverty). In the same way as in inequality measurement, where we argued for magnifying the effect of deviations far from the mean, we can argue that the effect of a transfer in poverty measurement should be dependent on the incomes of those involved in the transfer. For example, a transfer away from the lowest income household should have more effect on measured poverty than a transfer away from a household close to the poverty line. A poverty measure will satisfy this *sensitivity to transfers* if the increase in measured poverty caused by a transfer of income from a poor household to a poor household with a higher income is smaller, the larger is the income of the lowest income household.

Let the total population remain at H . Assume that this population can be divided into Γ separate subgroups. Let g_h^γ be the income gap of a poor member of subgroup γ and q^γ be the number of poor in that subgroup. Using this notation, a poverty measure that satisfies the property of subgroup consistency is the Foster–Greer–Thorbecke (FGT) class, given by

$$P_\alpha = \frac{1}{H} \sum_{\gamma=1}^{\Gamma} \left[\sum_{h=1}^{q^\gamma} \frac{g_h^\gamma}{z} \right]^\alpha. \quad (14.36)$$

The form of this measure depends on the value chosen for the parameter α . If $\alpha = 0$, then

$$P_0 = \frac{\sum_{\gamma=1}^{\Gamma} q^\gamma}{H} = E, \quad (14.37)$$

the head-count ratio. If instead $\alpha = 1$, then

$$P_1 = \frac{1}{H} \sum_{\gamma=1}^{\Gamma} \left[\sum_{h=1}^{q^\gamma} \frac{g_h^\gamma}{z} \right] = EI, \quad (14.38)$$

the product of the head-count ratio and the income gap ratio. Note that P_0 is insensitive to transfers, while the effect of a transfer for P_1 is independent of the incomes of the households involved. For higher values of α the FGT measure satisfies sensitivity to transfers, and more weight is placed on the income gaps of lower income households.

14.5.3 Two Applications

The use of these poverty measures is now illustrated by reviewing two applications. The first application, taken from Foster, Greer, and Thorbecke (1984), shows how subgroup consistency can give additional insight into the sources of poverty. The second application is extracted from an OECD working paper and illustrates how a range of poverty lines can be used as a check on consistency. It also reveals that there can be a good degree of agreement between different measures of poverty.

Table 14.3 reports an application of the FGT measure. The data are from a household survey in Nairobi and groups the population according to their length of residence in Nairobi. The measure used is the P_2 measure, so $\alpha = 2$. As already discussed, the use of the FGT measure allows the contribution of each group to total poverty to be identified. For example, those living in Nairobi between 6 and 10 years have a level of poverty of 0.0343 and contribute 12.1 percent to total poverty—this is also the percentage by which total poverty would fall if this group were all raised above the poverty line. The division into groups also allows identification of where the major contribution to poverty arises. In this case the major contribution is made by those in the 21 to 70 group. Although the actual poverty level in this group is low, the number of households in this group causes them to have a major effect on poverty.

The second application is reported in table 14.4. This OECD analysis studies the change in poverty over (approximately) a ten-year period from the mid-1980s to the

Table 14.3
Poverty using the FGT P_2 measure

Years in Nairobi	Level of poverty	Contribution to total poverty (%)
0	0.4267	5.6
0.01–1	0.1237	6.5
2	0.1264	6.6
3–5	0.0257	5.1
6–10	0.0343	12.1
11–15	0.0291	9.4
16–20	0.0260	6.6
21–70	0.0555	23.8
Permanent resident	0.1659	8.7
Don't know	0.2461	15.5
Total	0.0558	99.9

Source: Foster, Greer, and Thorbecke (1984).

Table 14.4
Evolution of poverty (% change in poverty measure)

Poverty line (% of median income)	40%	50%	50%	50%	60%
Measure	Head count	Head count	Income gap	Sen index	Head count
Australia, 1984–1993/94	0.0	–2.7	5.0	–4.2	–1.4
Belgium, 1983–1995	–1.4	–2.8	1.1	–27.1	–2.3
Germany, 1984–1994	1.8	2.9	2.5	20.8	3.8
Japan, 1984–1994	0.6	0.8	2.5	23.1	1.0
Sweden, 1983–1995	0.9	0.4	7.9	23.7	0.4
United States, 1985–1995	–1.2	–1.2	0.2	–4.9	–0.1

Source: OECD ECO/WKP(98)2.

mid-1990s. The numbers given are therefore the percentage change in the measure and not the value of the measure. What the results show is that the direction of change in poverty as measured by the head-count ratio is not sensitive to the choice of the poverty line—the only inconsistency is the value for Australia with the poverty line as 40 percent of median income. In detail, there has been a decrease in poverty in Australia, Belgium, and the United States but an increase in Germany, Japan, and Sweden. The results in the three central columns report the calculations for three different poverty measures. These show that the Sen measure and the head count are always in agreement about the direction of change. This is not true of the income gap, which disagrees with the other two for Australia and the United States.

14.6 Unequal Opportunities

The main focus of the literature on equality of opportunity is on separating sources of inequality of outcomes that are morally acceptable from those that are morally unacceptable. In the seminal book of Roemer, *Equality of Opportunity* (1998), it is shown that some inequality of outcomes is morally acceptable. Unequal outcomes that are a consequence of factors for which individuals are judged to be responsible—referred to as *effort*—are morally acceptable and should not be compensated for. Only inequality that is outside the realm of individual choice—referred to as *circumstances*—should be compensated for. Typical examples of circumstances that may affect individual outcomes are family background and individual attributes such as race, gender, and place of birth. To explore these ideas, it is first necessary to define equality of opportunity.

14.6.1 Defining Equality of Opportunity

When are opportunities equal? The central idea is to observe how the opportunity sets of people vary with their family background, race, gender, and so on. Equality of opportunity is achieved when no particular set of circumstances is preferred to another set of circumstances by all individuals. Given that people facing similar circumstances may produce different outcomes, the problem amounts to the comparison of distributions of outcomes, conditional on circumstances.

Consider a situation where individuals are allowed to choose their circumstances, s —which we refer to as their *type*—before they know their level of effort. Equality of opportunity prevails between circumstances s and s' if s is not preferred to s' by all individuals, and vice versa. In other words, people do not unanimously order the opportunity sets s and s' . The opportunity set of an individual can be represented by the conditional distribution function $F(x|s)$ denoting the probability of producing an outcome less than or equal to x for given type s . For example, this could describe the distribution of educational attainments for pupils of low-income families (say type s) compared to pupils from high-income families (say type s').

Under the (weak) assumption that preferences satisfy the criteria of first-order stochastic dominance (FSD) and second-order stochastic dominance (SSD), stochastic dominance tests can be used to rank conditional distribution functions. Formal definitions of first-order stochastic dominance (FSD) and second-order stochastic dominance (SSD) are given in, respectively, (14.39) and (14.40). Suppose inequality of opportunity where circumstance s is preferred to circumstance s' by all individuals. Inequality of opportunity defined as first-order stochastic dominance between s and s' means that the distribution of outcome x conditional on s is for all x below the distribution of x conditional on circumstance s' :

$$s \succeq_{FSD} s' \quad \text{iff} \quad F(x|s) \leq F(x|s'), \quad \forall x \in \mathfrak{R}_+. \quad (14.39)$$

However, it can easily be shown that this is a very weak definition of equality of opportunity. Indeed suppose a situation where the outcome distribution of type s always dominates the outcome distribution of type s' , except at the top (possibly, when they exert maximal effort). Under the definition of first-order stochastic dominance, equality of opportunity is not rejected in this case. But it is unfair because type s' must exert maximal effort to get a chance to outperform type s .

Second-order stochastic dominance provides extra restrictions. Under second-order stochastic dominance, equality of opportunity prevails when the expected value derived from distribution $F(y|s)$ is not greater than the one derived from $F(y|s')$:

$$s \succeq_{SSD} s' \text{ iff } \int_0^x F(y|s)dy \leq \int_0^x F(y|s')dy, \forall x \in \mathfrak{R}_+. \quad (14.40)$$

14.6.2 Measuring Equality of Opportunity

Econometric stochastic dominance techniques can be used to test FSD and SSD of conditional distributions. In relation to classical measure of inequality, we can also estimate inequality of opportunity with a Gini-type index. This index is based on the equivalence between SSD and Lorenz dominance. The Gini opportunity (GO) index with k types is defined as

$$GO(x) = \frac{1}{\mu} \sum_{i=1}^k \sum_{j>i}^k p_i p_j (\mu_j(1 - G_j) - \mu_i(1 - G_i)), \quad (14.41)$$

where μ is the mean of the population, μ_k the mean of group k , p_k the population weight of group k , and G the Gini coefficient. The GO index computes the sum of all pairwise differences of the opportunity sets of all types, where the opportunity sets are defined as twice the area under the generalized Lorenz curve, $\mu_s(1 - G_s)$, for type s . The GO index is in the interval $[0, 1]$. A value of 0 indicates perfect equality of opportunity.

14.6.3 Equal-Opportunity Policy

How does the theory of equal opportunity translate into policy? The distinctive feature of the equal-opportunity policy, compared to classical welfare policy, is that it is a nonwelfarist policy. Welfarism is the view that only the set of vectors of outcome (welfare) possibilities matters for choosing public policy. To be precise, if we represent individual preferences over social alternatives by utility functions, then the choice of a social alternative should depend only on the information that is recoverable from the utility possibilities sets. In this sense, welfarism is a consequentialist view.

The equal-opportunity approach says that one cannot judge the goodness of a social outcome by knowing only the distribution of outcomes; one must also know how hard people tried in order to evaluate that goodness. To put it differently, one must know the role of effort in achievement to pass judgment on the goodness of a public policy. From this perspective income taxation may not be the instrument of choice to equalize opportunities for income (as recommended by the tax principle): one naturally

thinks of using educational finance policy as a method for compensating children from disadvantaged families.

To illustrate, there is a large difference between an equal-resource policy, which invests the same amount in all children, and an equal-opportunity policy, which invests in children so as to compensate for different social circumstances. The United States, with its system of locally financed public education, is in most places less equitable even than the equal resource policy would be: that is, usually more is invested in the public education of advantaged children than of disadvantaged children. Most of the affirmative action policies are grounded in an equal-opportunity approach.

14.7 Intergenerational Inequality

Social mobility occurs when different generations of a family have differing social status. Formally, social mobility refers to a lack of correlation between the educational attainments and earnings of parents and those of their children. To what extent is low educational attainment and consequent poverty transmitted by parents to their children? How can this be measured? And how can it be explained?

14.7.1 Measuring Issues

There are basically two sets of measures of the transmission of income and education across generations. The first set relates the earnings of parents to those of children, or the education of parents to that of children. Both the *correlation* and *elasticities* of these variables between parents and children give a basic measure of the intergenerational transmission of income and education. The difference between correlations and elasticities reflects differences in the variance of income/education for parents' and children's generations. The correlation is always between 0 and 1, but the elasticity can be greater than 1 if there is an increase in inequality across generations. The complement to the correlation ($1 - \text{correlation}$) is a measure of intergenerational mobility.

The central issues are to measure permanent income by averaging over a sufficient number of years (to accommodate temporary shocks and lifetime income progression) and to compare the income of the parents to that of children at the same point in their life cycle. As a matter of illustration, Nicoletti and Ermisch (2007) obtain income elasticities for the United States of around 0.5 to 0.6, and for the United Kingdom of around 0.3. Elasticities for the Nordic countries are always lower than 0.3. The comparison across countries and over time of these elasticities requires the use of

the same estimation methods, same definitions, and the same sample selection rule. There are also interesting nonlinearities in the elasticities. Bratsberg et al. (2007) find strong nonlinearities in the Nordic countries but not in the United States or the United Kingdom. For example, in Denmark the elasticities are 0.06 at the bottom of the income distribution and 0.31 at the top of the distribution. They suggest that this nonlinearity is related to the strong public education systems that exist in the Nordic countries.

The second set of measures use *transition matrices* to estimate mobility at each point in the distribution. Transition matrices enable us to compare the mobility across the full distribution rather than just around the mean. Jantti et al. (2006) split the distribution into quintiles and study mobility across quintiles. They find that more than 40 percent of sons in the United States who are born to fathers in the lowest quintile are in the lowest quintile themselves. Mobility from the lowest quintile is found to be much higher in Nordic countries. They find that much of the difference in mobility rates between the United States and the Nordic countries are in fact attributable to difference in the tails of the income distributions. Mobility across the middle three quintiles is very similar across all countries.

Some authors define upward mobility as the probability that a child's percentile rank in the distribution of children exceeds the father's percentile rank in the distribution of fathers. In this case more weight is placed on small moves in position compared to the quintile or quartiles approach. It has been shown that this distinction matters as the degree of upward mobility of blacks in the United States is similar to that of whites when the finer mobility metric is used.

14.7.2 Causal Mechanisms

Understanding the determinants of intergenerational correlations is crucial for the development of appropriate public policy. Without uncovering the driving forces behind intergenerational transmission, it is impossible to figure out how to promote changes. This is, of course, a very difficult task, as it is often the case that any particular parental attribute (e.g., as education or earnings) is correlated with a variety of parental nonobservable characteristics.

Nature and Nurture

Strong earnings and educational correlations among siblings reveal the importance of shared genetic and environmental factors, but they are not very helpful for pinning down causal mechanisms. Bjorklund et al. (2005) make use of correlations across

identical twins, fraternal twins, full siblings, half siblings, and adopted siblings, when they are both raised together to distinguish between nature and nurture. They show that genetics are more important than shared environment. Surprisingly, the identical twin correlation is only 0.36.

We must interpret this decomposition between nature and nurture with great prudence because it is possible that twins are treated differently by parents and, more important, because with assortative mating people tend to marry people with similar characteristics and so likely share genetic characteristics. It is just as likely there is positive interaction between environmental and genetic factors that complicates the separation of the effects of the two. Also twins are hardly representative of the population.

Using a sample of adopted children in Wisconsin, Plug (2004) finds a regression coefficient of about 0.28 on adoptive mother's education. However, the interpretation of this is not clear because adoption is not random and may result in positive selection whereby high-ability adoptees end up in more educated families. Overall, this literature on twins', siblings', and adoptees' correlations suggest that both environmental and genetic factors are important. Even so, we cannot be sure which factors matter most. Taken as a whole, the findings in the literature are very inconsistent. Some studies find the effect of fathers to be more important, other studies find the effects of mothers to be more important.

School System

It has been shown that the school system matters. A standard argument is based on *credit constraints*. Solon (2004) presents a basic model in which families are credit constrained and must reduce current consumption to invest in human capital. If there are no credit constraints, and thus parents can borrow from their children's future earnings, each family will optimally invest in the human capital of their children (assuming perfect altruism). The optimality conditions will recommend investing more in high-ability children, and so we should expect intergenerational transmission of inequality only if child ability and parental earnings are correlated. If there are credit constraints, however, poor families cannot invest optimally in their children's human capital. As a result a higher income level means an increase human capital investment.

The model also predicts that the intergenerational elasticity is increasing in the genetic transmission of ability and the rate of return on human capital investment, but it is decreasing with the public investment in education. The implication is that the low elasticities in the Nordic countries could be explained either by low returns on training investment (the compressed earnings distributions) or the public education system that

tends to equalize educational opportunities for children. Suggestive evidence is contained in Ichino et al. (2009) who report a negative correlation of -0.54 between public expenditure on education and intergenerational income elasticities. The correlation is even stronger with public expenditure on primary education.

Attitudes and Social Behaviors

What do parents transmit to their children? Other social scientists have claimed that attitudes matter more than the socioeconomic status of the family. This is particularly relevant for low-income groups and minorities. Indeed Cunha and Heckman (2008) survey a large number of studies that show that nonpecuniary factors (psychic costs, motivations, etc.) play a major role in explaining why minorities and persons from low-income families do not attend college even though it is financially profitable to do so. Because returns to schooling are lower for people less likely to attend college, these groups may “rationally” choose to underinvest in education. Obviously economists have much to learn from other social scientists to better grasp the nonpecuniary factors that shape future economic and educational attainments.

Segregation

Segregation is important for understanding the dynamics of group inequality. Members of different groups with identical distributions of cognitive abilities will invest differently in the education of their children when the segregation is sufficiently great. Furthermore the relationship between group equality and social segregation can exhibit a discontinuity if there exists a critical level of segregation such that convergence to group equality occurs if and only if segregation lies below this threshold. Hence a small increase in social integration, if it takes the economy across the threshold, may have large effects on long-run group inequality, while a large increase in integration that does not cross the threshold may have no persistent effect. This suggests that economic inequality across social groups might arise endogenously under certain conditions, without preexisting discrimination or group differences in ability or wealth.

Dynamic systems with positive feedback can exhibit multiple types of self-reinforced behavior at the group level. Positive feedback means that while there is tendency for members of a group to make similar decisions, such feedback does not say which is generally made. The implication of positive feedback is that different groups with identical characteristics can eventually adopt different norms of behaviors. This is certainly a fruitful direction of investigation but a suitable method of empirical testing is not obvious. Indeed there is no agreement in the literature on how to measure segregation

(with nearly 20 different indexes of segregation representing different concepts such as evenness, exposure, concentration, centralization, and clustering), and as a result it is hardly possible to verify whether segregation is a leading cause of racial and social differences on economic outcomes.

14.8 Conclusions

The need to quantify is driven by the aim of making precise comparisons. What economic analysis contributes is an understanding of the bridge between intuitive concepts of inequality and poverty, and specific measures of these phenomena. Analysis can reveal the implications of alternative measures and provide principles that a good measure should satisfy.

The first problem we challenged in this chapter was the comparison of incomes between households of different compositions. It is clearly more expensive to support a large family than a small family, but exactly how much more expensive is more difficult to determine. Equivalence scales were introduced as the analytical tool to solve this problem. These scales were initially based on the cost of achieving a minimum standard of living. Although simple, such an approach does not easily generalize to higher income levels, nor does it take much account of economic optimization. In principle, equivalence scales could be built directly from utility functions, but to do so, issues must be addressed of how the preferences of the individual members of a household are aggregated into a household preference order.

Inequality occurs when some households have a higher income (after the incomes have been equivalized for household composition) than others. The Lorenz curve provides a graphical device for contrasting income distributions. Some income distributions can be ranked directly by the Lorenz curve, in which case there is no ambiguity about which has more inequality, but not all distributions can be. Inequality measures provide a quantitative assessment of inequality by imposing restrictions beyond those incorporated in the Lorenz curve. The chapter investigated the properties of a number of measures of inequality. Of particular importance was the observation that all inequality measures embody implicit welfare judgments. Given this, the Atkinson measure is constructed on the basis that the welfare judgments should be made explicit and the inequality measure constructed on these judgments. In principle, alternative measures can generate different rankings of income distributions, but in practice, as the application showed, they can yield very consistent rankings.

In many ways the measurement of poverty raises similar issues to those of inequality. The additional feature of poverty is the necessity to determine whether households can be classed as living in poverty. The poverty line, which provides the division between the two groups, plays a central role in poverty measurement. Where and how to locate this poverty line is important, but more fundamental is how it should be adjusted over time. At stake here is the key question of whether poverty should be viewed in absolute or relative terms. The practice in developed countries is to use relative poverty. The chapter reviewed a number of poverty measures from the head-count ratio to the Foster–Greer–Thorbecke measure. These measures are also distinguished by a range of sensitivity properties. The applications showed how they could be used and that the different measures could provide a consistent picture of the development of poverty despite their different conceptual bases.

The chapter has revealed how economic analysis is able to provide insights into what we are assuming when we employ a particular inequality or poverty measure. It has also revealed how we can think about the process of improving our measures. Inequality and poverty are significant issues, and better measurement is a necessary starting point for better policy.

Further Reading

The relationship between inequality measures and social welfare was first explored in:

Atkinson, A. B. 1970. On the measurement of inequality. *Journal of Economic Theory* 2: 244–63.

A comprehensive survey of the measurement of inequality is given by:

Sen, A. K. 1997. *On Economic Inequality*. Oxford: Oxford University Press.

A textbook treatment is in:

Lambert, P. 1989. *The Distribution and Redistribution of Income: A Mathematical Analysis*. Oxford: Basil Blackwell.

Issues surrounding the definition and implications of the poverty line are treated in:

Atkinson, A. B. 1987. On the measurement of poverty. *Econometrica*. 55: 749–64.

Callan, T., and Nolan, B. 1991. Concepts of poverty and the poverty line. *Journal of Economic Surveys* 5: 243–61.

The derivation of the Sen measure, and a general discussion of constructing measures from a set of axioms is given by:

Sen, A. K. 1976. Poverty: An ordinal approach to measurement. *Econometrica* 44: 219–31.

The FGT measure was first discussed in:

Foster, J. E., Greer, J., and Thorbecke, E. 1984. A class of decomposable poverty measures. *Econometrica* 52: 761–67.

An in-depth survey of poverty measures is:

Foster, J. E. 1984. On economic poverty: A survey of aggregate measures. *Advances in Econometrics* 3: 215–51.

The seminal contribution on the equality of opportunity is

Roemer, J. 1998. *Equality of Opportunity*. Cambridge: Harvard University Press.

Social mobility issues including the references presented in the chapter are treated in:

Bratsberg, B., Roed, K., Raum, R., Naylor, A., Jntti, M., Osterbacka, E. E. 2007. Nonlinearities in intergenerational earnings mobility: consequences for cross country comparisons, *Economic Journal* 117: C72–92,

Jantti, M., Bratsberg, B., Roed, K., Raaum, O., Naylor, R., Osterbacka, E., Bjorklund A., Eriksson T. 2006. America exceptionalism in a new light: A comparison of intergenerational earnings mobility in the Nordic countries, the U.S. and the U.K. IZA discussion paper 1938. Bonn.

Nicoletti, C., and Ermisch, J. 2007. Intergenerational earnings mobility: Changes across cohorts in Britain. *BE Journal of Economic Analysis and Policy* 7: Article 9.

The issue of nature and nurture is discussed by:

Bjorklund, A., Jantti, M., and Solon, G. 2005. Influences of nature and nurture on earnings variation: A report on a study of various sibling types in Sweden. In S. Bowles, H. Gintis, and M. Osborne, eds. *Unequal chances: Family Backgrounds and Economic Success*. Princeton: Princeton University Press, 145–64.

Plug, E. 2004. Estimating the effect of mother's schooling on children's schooling using a sample of adoptees. *American Economic Review* 94: 358–68.

The link between social inequality and the school system:

Cunha, F., and Heckman, J. 2008. Formulating, identifying and estimating the technology of cognitive and noncognitive skill formation. *Journal of Human Resources* 43(4): 238–82.

Ichino, A., Karabarbounis, L. and Moretti, E. 2009. The political economy of intergenerational income mobility. IZA Discussion paper 4767. Bonn.

Solon, G. 2004. A model of intergenerational mobility variation over time and place. In M. Corak, ed., *Generational Income Mobility in North America and Europe*. Cambridge: Cambridge University Press, 38–47,

Exercises

- 14.1** In many countries lottery prizes are not taxed. Is this consistent with Hicks's definition of income?
- 14.2** Let the utility function be $U = d^{1/2} \log(M)$, where d is family size. Construct the equivalence scale for the value of $U = 10$. How is the scale changed if $U = 20$?

- 14.3** What economies of scale are there in family size? Are these greater or smaller at low incomes?
- 14.4** Take the utility function $U = \log\left(\frac{x_1}{d}\right) + \log\left(\frac{x_2}{d}\right)$, where d is family size and good 1 is food.
- What proportion of income is spent on food? Can this provide the basis for an equivalence scale? Calculate the exact equivalence scale. Does it depend on U ?
 - Repeat part a for the utility function $U = \left[\frac{x_1}{d}\right]^{1/2} + \left[\frac{x_2}{d}\right]^{1/2}$.
- 14.5** If children provide utility for their parents, show on a diagram how an equivalence scale can decrease as family size increases.
- 14.6** Consider a community with ten persons.
- Plot the Lorenz curve for the income distribution (2, 4, 6, 8, 10, 12, 14, 16, 18, 20)
 - Consider an income redistribution that takes two units of income from each of the four richest consumers and gives two units to each of the four poorest. Plot the Lorenz curve again to demonstrate that inequality has decreased.
 - Show that the Lorenz curve for the income distribution (2, 3, 5, 9, 11, 12, 15, 17, 19, 20) crosses the Lorenz curve for the distribution in part a.
 - Show that the two social welfare functions $W = \sum M^h$ and $W = \sum \log(M^h)$ rank the income distributions in parts a and c differently.
- 14.7** What is the Gini index, and how can it be used to determine the impact of taxes and transfers on income inequality?
- 14.8** Calculate the Gini index for the income distributions used in parts a through c of exercise 14.6. Discuss the values obtained.
- 14.9** For a utilitarian social welfare function construct M_{EDE} for the distributions used in exercise 14.6 if the utility of income is logarithmic. Find the Atkinson inequality measure. Repeat the exercise for the Rawlsian social welfare function. Compare and discuss.
- 14.10** What drawbacks are there to eliminating inequality?
- 14.11** Should we be concerned with inequality if it is due to differences in ability? What if it is due to differences in effort levels?
- 14.12** Define inequality aversion. Explain how it is related to the concept of risk aversion.
- 14.13** Discuss the following quote from Cowell (1995, p. 23): “The main disadvantage of G[ini] is that an income transfer from a rich to a poorer man has a much greater effect on G if the men are near the middle rather than at either end of the parade.” Do you agree? Why or why not? (*Hint*: Use the formula for the Gini coefficient to determine the effect of a fixed transfer at different points in the income distribution.) Does the Gini have other “disadvantages”?
- 14.14** Consider a hypothetical island with only ten people. Eight have income of \$10,000, one has income of \$50,000, and one has income of \$100,000.
- Draw the Lorenz curve for this income distribution. What is the approximate value of the Gini coefficient?

- b. Suppose that a wealthy newcomer arrives on this island with an income of \$500,000. How does it change the Lorenz curve? What is the impact on the Gini coefficient?
- 14.15** Have a look at actual income distribution in the United States available on the website <<http://www.census.gov/hhes/income/histinc/histinctb.html>. Select Households and then Table H-2>.
- Plot the Lorenz curve for 1981 and 2001. Clearly label each curve. What can you say about the evolution of inequality over time?
 - Based on your diagram, can we conclude that the Gini coefficient was higher in 1981 or 2001? Explain. Check your answer by consulting Table H-4 on the website.
 - Can we conclude from the diagram that the poor were necessarily worse off in either 1981 or 2001? Why or why not? Use Table H-1 on the website to refine your answer.
 - Now suppose that people with similar incomes are more likely to get married than people with dissimilar incomes. How would this change affect the Lorenz curve drawn in part a?
- 14.16** There are two senior advisors to the government, *A* and *B*, both of whom agree that the poverty line is at \$4,000 for a single person. However, they have different equivalence scales. Mr. *A* believes that the scale factor in determining equivalent income should be 0.25 for each additional family member. Mrs. *B* suggests that the scale factor should be 0.75.
- Find the poverty line for families of two, three, and four under both values of the scale factors 0.25 and 0.75.
 - Explain how Mr. *A* and Mrs. *B* must have very different views about income sharing within a family to end up with such different answers.
 - Suppose that the government is committed to providing welfare eligibility to every family below the poverty line. If this government wishes to keep total spending to a minimum, which of the two views should it support?
- 14.17** Given the income distributions
(1, 2, 2, 5, 5, 5, 7, 11, 11, 12, 20, 21, 22, 24),
(2, 3, 3, 4, 4, 5, 7, 7, 11, 11, 12, 20, 21, 24),
and a poverty line of $z = 6$, calculate the Sen poverty measure. Explain the values obtained for the two distributions.
- 14.18** Use the two income distributions in exercise 14.17 to evaluate the Foster–Greer–Thorbecke poverty measure for $\alpha = 2$. Pool the distributions to evaluate the poverty measure for the total population. Show that the measure is a weighted sum of the measures for the individual distributions.
- 14.19** (Decoster) The Pareto distribution is a popular functional form for describing income distributions. It is a two-parameter specification for which the frequency density function reads as follows: $f(x) = \alpha x_0^\alpha x^{-[1+\alpha]}$ for $x \geq x_0$, where $x_0 > 0$ is the lowest income level and $\alpha > 1$ is a parameter.
- Show that the mean income for the Pareto distribution is $\bar{x} = \frac{\alpha x_0}{\alpha - 1}$.
 - Show that the distribution function for the Pareto distribution is $F(x) = 1 - \left[\frac{x_0}{x}\right]^\alpha$ for $x \geq x_0$. Discuss the effect of changing the parameter α .

c. The Pareto distribution parameterized by α can easily be used to construct a very simple inequality measure, which is defined as follows: Take an arbitrary income level, say x . Calculate the mean income of the subpopulation of all income earners who have an income larger than x . The ratio of this mean income to the income x is given by $I = \frac{\alpha}{\alpha-1}$. Calculate the values of this inequality index for some different values of α (e.g., $\alpha = 1.5, 2, 3$). Does α represent equality or inequality? What is the limiting value of I for very large α ? Interpret this result.

d. Show that the Lorenz curve for the Pareto distribution is $L(p) = 1 - [1 - p]^{\alpha-1/\alpha}$, where $p = F(x)$ and $p \in [0, 1]$. What is the shape of the curve for very large α ?

e. Draw the Lorenz curve for two values $\alpha_1 > \alpha_2$, and verify that the two Lorenz curves will never cross.

f. Show that the Gini coefficient for the Pareto distribution (with parameter α) is $G = \frac{1}{2\alpha-1}$. How does it compare with your answer in part e?

14.20 How is it possible for a government that pursues a poverty reduction target to increase inequality? Does this possibility invalidate poverty comparisons?

14.21 How are intergenerational persistence and mobility measured?

14.22 (Solon 2004) Consider the following model of intergenerational transmission. Define Y_{it} as the earnings of generation t in family i , and I_{it-1} investment by generation $t-1$ in the human capital of generation t . Human capital accumulation is given by

$$H_{it} = \theta \log(I_{it-1}) + e_{it},$$

where θ is a (uniform) productivity parameter, and e_{it} denotes individual ability and follows an AR(1) process:

$$e_{it} = \delta + \lambda e_{it-1} + v_{it},$$

where λ denotes genetic transmission of ability. Human capital translates into earnings on the labor market as follows:

$$\log(Y_{it}) = \mu + \rho H_{it},$$

where ρ denotes the returns to human capital. Family i have Cobb–Douglas preferences over own consumption, C_{it-1} , and child's income, Y_{it} ,

$$U(C, Y) = C_{it-1}^{1-\alpha} Y_{it}^{\alpha}.$$

The budget constraint is

$$C_{it-1} = Y_{it-1} - I_{it-1}.$$

- What is the optimal investment in human capital?
- What is the resulting equation of intergenerational earnings transmission?
- What is the resulting intergenerational elasticity?
- Discuss the solutions.

14.23 To test the predictions of the Solon model described in the exercise above, use the cross-sectional PISA data on education test scores for pupils at 15 years old in OECD countries. (Available from OECD at www.pisa.oecd.org). Pick a country of your choice and assemble

for each pupil of the country of your choice their test score and their ESCS index (i.e., the economic and sociocultural index of their parent).

- a. Plot the pupils test scores and their ESCS index. Compute the regression line.
- b. Regroup pupils by ESCS quartile with the first quartile representing the 25 percent of pupils with the lowest ESCS and the fourth quartile representing the 25 percent of pupils with the highest ESCS index. Draw the cumulative distribution function of test scores for each ESCS quartile and check whether the first-order stochastic dominance condition is satisfied.