3 *Welfare, poverty, and distribution*

One of the main reasons for collecting survey data on household consumption and income is to provide information on living standards, on their evolution over time, and on their distribution over households. Living standards of the poorest parts of the population are of particular concern, and survey data provide the principal means for estimating the extent and severity of poverty. Consumption data on specific commodities tell us who consumes how much of what, and can be used to examine the distributional consequences of price changes, whether induced by deliberate policy decisions or as a result of weather, world prices, or other exogenous forces. In this chapter, I provide a brief overview of the theory and practice of welfare measurement, including summary measures of living standards, of poverty, and of inequality, with illustrations from the Living Standards Surveys of Côte d'Ivoire from 1985 through 1988 and of South Africa in 1993. I also discuss the use of survey data to examine the welfare effects of pricing and of transfer policies using as examples pricing policy for rice in Thailand and pensions in South Africa.

The use of survey data to investigate living standards is often straightforward, requiring little statistical technique beyond the calculation of measures of central tendency and dispersion. Although there are deep and still-controversial conceptual issues in deciding how to measure welfare, poverty, and inequality, the measurement itself is direct in that there is no need to estimate behavioral responses nor to construct the econometric models required to do so. Instead, the focus is on the data themselves, and on the best way to present reliable and robust measures of welfare. Graphical techniques are particularly useful and can be used to describe the whole distribution of living standards, rather than focussing on a few summary statistics. For example, the Lorenz curve is a standard tool for charting inequality, and in recent work, good use has been made of the cumulative distribution function to explore the robustness of poverty measures. For other questions it is useful to be able to display (univariate and bivariate) density functions, for example when looking at two measures of living standards such as expenditures and nutritional status, or when investigating the incidence of price changes in relation to the distribution of real incomes. While cross-tabulations and histograms are the traditional tools for charting densities, it is often more informative to calculate nonparametric estimates of densities using one of the smoothing methods that have recently been developed in the statistical literature. One of the purposes of this chapter is to explain these methods in simple terms, and to illustrate their usefulness for the measurement of welfare and the evaluation of policy.

The chapter consists of three sections. Section 3.1 is concerned with welfare measurement, and Section 3.3 with the distributional effects of price changes and cash transfers. Each section begins with a brief theoretical overview and continues with empirical examples. The techniques of nonparametric density estimation are introduced in the context of living standards in Section 3.2 and are used extensively in Section 3.3 This last section shows how regression functions—conditional expectations—can often provide direct answers to questions about distributional effects of policy changes, and I discuss the use of nonparametric regression as a simple tool for calculating and presenting these regression functions.

3.1 Living standards, inequality, and poverty

Perhaps the most straightforward way to think about measuring living standards and their distribution is a purely statistical one, with the mean, median, or mode representing the central tendency and various measures of dispersion—such as the variance or interquartile range—used to measure inequality. However, greater conceptual clarity comes from a more theoretical approach, and specifically from the use of social welfare functions as pioneered by Atkinson (1970). This is the approach that I follow here, beginning with social welfare functions, and then using them to interpret measures of inequality and poverty.

Social welfare

Suppose that we have decided on a suitable measure of living standards, denoted by x; this is typically a measure of per capita household real income or consumption, but there are other possibilities, and the choices are discussed below. We denote the value of social welfare by W and write it as a nondecreasing function of all the x's in the population, so that

$$W = V(x_1, x_2, \ldots, x_N) \tag{3.1}$$

where N is the population size. Although our data often come at the level of the household, it is hard to give meaning to household or family welfare without starting from the welfare of its members. In consequence, the x's in (3.1) should be thought of as relating to individuals, and N to the number of persons in the population. The issue of how to move from household data to individual welfare is an important and difficult one, and I shall return to it.

It is important not to misinterpret a social welfare function in this context. In particular, it should definitely not be thought of as the objective function of a government or policymaking agency. There are few if any countries for which the maximization of (3.1) subject to constraints would provide an adequate description

of the political economy of decisionmaking. Instead, (3.1) should be seen as a statistical "aggregator" that turns a distribution into a single number that provides an overall judgment on that distribution and that forces us to think coherently about welfare and its distribution. Whatever our view of the policymaking process, it is always useful to think about policy in terms of its effects on efficiency and on equity, and (3.1) should be thought of as a tool for organizing our thoughts in a coherent way.

What is the nature of the function V, and how is it related to the usual concepts? When V is increasing in each of its arguments, social welfare is greater whenever any one individual is better-off and no one is worse-off, so that Pareto improvements are always improvements in social welfare. For judging the effects of any policy, we shall almost always want this *Pareto condition* to be satisfied. However, as we shall see, it is often useful to think about poverty measurement in terms of social welfare, and this typically requires a social welfare function that is unresponsive to increases in welfare among the nonpoor. This case can be accommodated by weakening the Pareto condition to the requirement that V be *nondecreasing* in each of its arguments.

Social welfare functions are nearly always assumed to have a *symmetry* or *anonymity* property, whereby social welfare depends only on the list of welfare levels in society, and not on who has which welfare level. This makes sense only if the welfare levels are appropriately defined. Money income does not translate into the same level of living at different price levels, and a large household can hardly be as well-off as a smaller one unless it has more money to spend. I shall return to this issue below, when I discuss the definition of x, and in Chapter 4, when I discuss the effects of household composition on welfare.

Finally, and perhaps most importantly, social welfare functions are usually assumed to prefer more equal distributions to less equal ones. If we believe that inequality is undesirable, or equivalently that a gift to an individual should increase social welfare in (3.1) by more when the recipient is poorer, then for any given total of x—and ignoring any constraints on feasible allocations—social welfare will be maximized when all x's are equal. (Note that policies that seek to promote equality will often have incentive effects, so that a preference for equality is not the same as a claim that equality is desirable once the practical constraints are taken into account.) Equity preference will be guaranteed if the function V has the same properties as a standard utility function, with diminishing marginal utility to each x, or more formally, when it is quasi-concave, so that when we draw social indifference curves over the different x's, they are convex to the origin. Quasi-concavity of V means that if x^1 and x^2 are two lists of x's, with one element for each person, and if $V(x^1) = V(x^2)$ so that the two allocations are equally socially valuable, then any weighted average, $\lambda x^1 + (1-\lambda)x^2$ for λ between 0 and 1, will have as high or higher social welfare. A weighted average of any two equally good allocations is at least as good as either. In particular, quasi-concavity implies that social welfare will be increased by any transfer of x from a richer to a poorer person, provided only that the transfer is not sufficiently large to reverse their relative positions. This is the "principle of transfers," originally proposed by Dalton (1920). It should be noted that the principle of transfers does not require quasi-concavity, but a weaker condition called "s-concavity" (see Atkinson 1992 for a survey and more detailed discussion).

Inequality and social welfare

For the purposes of passing from social welfare to measures of inequality, it is convenient that social welfare be measured in the same units as individual welfare, so that proportional changes in all x's have the same proportional effect on the aggregate. This will happen if the function V is homogeneous of degree one, or has been transformed by a monotone increasing transform to make it so. Provided the transform has been made, we can rewrite (3.1) as

$$W = \mu V\left(\frac{x_1}{\mu}, \ldots, \frac{x_N}{\mu}\right) \tag{3.2}$$

where μ is the mean of the x's. Equation (3.2) gives a separation between the mean value of x and its distribution, and will allow us to decompose changes in social welfare into changes in the mean and changes in a suitably defined measure of inequality. Finally, we choose units so that $V(1,1,\ldots,1) = 1$, so that when there is perfect equality, and everyone has the mean level of welfare, social welfare is also equal to that value.

Since social welfare is equal to μ when the distribution of x's is perfectly equal, then, by the principle of transfers, social welfare for any unequal allocation cannot be greater than the mean of the distribution μ. Hence we can write (3.2) as

$$W = \mu(1 - I) \tag{3.3}$$

where I is defined by the comparison of (3.2) and (3.3), and represents the cost of inequality, or the amount by which social welfare falls short of the maximum that would be attained under perfect equality. I is a measure of inequality, taking the value zero when the x's are equally distributed, and increasing with disequalizing transfers. Since the inequality measure is a scaled version of the function V with a sign change, it satisfies the principle of transfers in reverse, so that any change in distribution that involves a transfer from rich to poor will decrease I as defined by (3.2) and (3.3).

Figure 3.1 illustrates social welfare and inequality measures for the case of a two-person economy. The axes show the amount of x for each of the two consumers, and the point S marks the actual allocation or status quo. Since the social welfare function is symmetric, the point S', which is the reflection of S in the 45-degree line, must lie on the same social welfare contour, which is shown as the line SBS'. Allocations along the straight line SCS' (which will not generally be feasible) correspond to the same total x, and those between S and S' have higher values of social welfare. The point B is the point on the 45-degree line that has the same social welfare as does S; although there is less x per capita at B than at S, the equality of distribution makes up for the loss in total. The amount of x at B is

Figure 3.1. Measuring inequality from social welfare

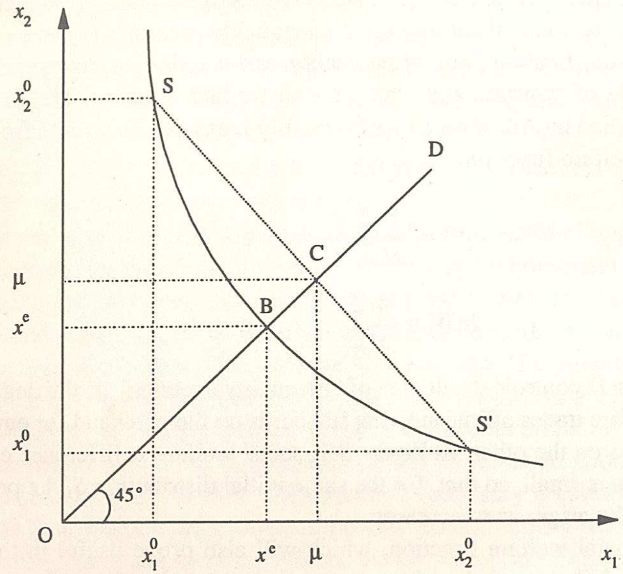

denoted x^e, and is referred to by Atkinson as "equally distributed equivalent x." *Equality* is measured by the ratio OB/OC, or by x^e/μ, a quantity that will be unity if everyone has the same, or if the social welfare contours are straight lines perpendicular to the 45-degree line. This is the case where "a dollar is a dollar" whoever receives it so that there is no perceived inequality. Atkinson's measure of inequality, defined by (3.3), is shown in the diagram as the ratio BC/OC.

One of the advantages of the social welfare approach to inequality measurement, as embodied in (3.3), is that it precludes us from making the error of interpreting measures of inequality *by themselves* as measures of welfare. It will sometimes be the case that inequality will increase at the same time that social welfare is increasing. For example, if everyone gets better-off, but the rich get more than the poor, inequality will rise, but there has been a Pareto improvement, and most observers would see the new situation as an improvement on the original one. When inequality is seen as a component of social welfare, together with mean levels of living, we also defuse those critics who point out that a focus on inequality misdirects attention away from the living standards of the poorest (see in particular Streeten et al 1981). Atkinson's formulation is entirely consistent with an approach that pays attention only to the needs of the poor or of the poorest groups, provided of course that we measure welfare through (3.3), and not through (negative) I alone. Just to reinforce the point, we might define a "basic-needs" social welfare function to be the average consumption of the poorest five percent of society, μ^p say. This measure can be rewritten as $\mu(1 - I)$, where I is the inequality measure $1 - \mu^p/\mu$.

Measures of inequality

Given this basic framework, we can generate measures of inequality by specifying a social welfare function and solving for the inequality measure, or we can start from a standard statistical measure of inequality, and enquire into its consistency with the principle of transfers and with a social welfare function. The first approach is exemplified by Atkinson's own inequality measure. This starts from the additive social welfare function

$$W = \frac{1}{N}\sum_{i=1}^{N}\frac{x_i^{1-\epsilon}}{1-\epsilon}, \quad \epsilon \neq 1 \tag{3.4a}$$

$$\ln W = \frac{1}{N}\sum_{i=1}^{N}\ln x_i, \quad \epsilon = 1. \tag{3.4b}$$

The parameter $\epsilon \geq 0$ controls the degree of "inequality aversion" or the degree to which social welfare trades off mean living standards on the one hand for equality of the distribution on the other. In Figure 3.1, social welfare indifference curves are flatter when ϵ is small, so that, for the same initial distribution S, the point B moves closer to the origin as ϵ increases.

Atkinson's social welfare function, which will also prove useful in the tax reform analysis of Chapter 5, has the property that the ratio of marginal social utilities of two individuals is given by the reciprocal of the ratio of their x's raised to the power of ϵ:

$$\frac{\partial W/\partial x_i}{\partial W/\partial x_j} = (x_j/x_i)^{\epsilon}. \tag{3.5}$$

Hence, if ϵ is zero so that there is no aversion to inequality, marginal utility is the same for everyone, and social welfare is simply μ, the mean of the x's. If ϵ is 2, for example, and i is twice as well-off as j, then the marginal social utility of additional x to i is one-fourth the marginal social utility of additional x to j. As ϵ tends to infinity, the marginal social utility of the poorest dominates over all other marginal utilities, and policy is concerned only with the poorest. When social welfare is the welfare of the poorest, which is what (3.4) becomes as ϵ tends to infinity, social preferences are sometimes said to be maximin (the object of policy is to maximize the minimum level of welfare) or Rawlsian, after Rawls (1972). Thinking about relative marginal utilities according to (3.5) is sometimes a convenient way of operationalizing the extent to which one would want poor people to be favored by policies or projects.

The inequality measure associated with (3.4) are, when $\epsilon \neq 1$,

$$I = 1 - \left(\frac{1}{N}\sum_{i=1}^{N}(x_i/\mu)^{1-\epsilon}\right)^{1/(1-\epsilon)} \tag{3.6a}$$

and, when $\epsilon = 1$, the multiplicative form

$$I = 1 - \prod_{i=1}^{N}(x_i/\mu)^{1/N}. \tag{3.6b}$$

These expressions are obtained by raising social welfare to the power of $1/(1-\epsilon)$, which makes the function homogeneous of the first degree, and then following through the procedures of the previous subsection. In line with the interpretation of ϵ as an aversion or perception parameter, there is no (perceived) inequality when ϵ is zero, no matter what the distribution of the x's. Conversely, if $\epsilon > 0$ and one person has all but a small amount α, say, with α spread equally over the others, then I tends to one as the number of people becomes large. Values of ϵ above 0 but below 2 appear to be useful, although in applications, it is often wise to look at results for a range of different values.

We may also choose to start from the standard measures of inequality. Provided these satisfy the principle of transfers, they will be consistent with Atkinson's approach, and will each have an associated social welfare function that can be recovered by applying (3.3). Some statistical measures of inequality do not satisfy the principle of transfers. The interquartile ratio—the 75th percentile less the 25th percentile divided by the median—is one such. Transferring x from a richer to a poorer person in the same quartile group will have no effect on inequality, and a transfer from someone at the bottom quartile to someone poorer will lower the bottom quartile and so will actually increase inequality. Less obviously, it is also possible to construct cases where a transfer from a better-off to a poorer person will increase the variance of logarithms. However, this can only happen when both people are far above the mean—which may not be relevant in some applications—and the other conveniences of the log variance may render it a competitive inequality measure in spite of this deficiency.

Other standard measures that *do* satisfy the principle of transfers are the Gini coefficient, the coefficient of variation, and Theil's "entropy" measure of inequality. The Gini coefficient if often defined from the Lorenz curve (see below), but can also be defined directly. One definition is the ratio to the mean of half the average over all pairs of the absolute deviations between people; there are $N(N-1)/2$ distinct pairs in all, so that the Gini is

$$\gamma = \frac{1}{\mu N(N-1)}\sum_{i>j}\sum_{j}|x_i - x_j|. \tag{3.7a}$$

Note that when everyone has the same, μ, the Gini coefficient is zero, while if one person has $N\mu$, and everyone else zero, there are $N-1$ distinct nonzero absolute differences, each of which is $N\mu$, so that the Gini is 1. The double sum in (3.7a) can be expensive to calculate if N is large, and an equivalent but computationally more convenient form is

$$\gamma = \frac{N+1}{N-1} - \frac{2}{N(N-1)\mu}\sum_{i=1}^{N}\rho_i x_i \tag{3.7b}$$

where ρ_i is the rank of individual i in the x-distribution, counting from the top so that the richest has rank 1. Using (3.7b), the Gini can straightforwardly and rapidly be calculated from microeconomic data after sorting the observations. I shall give examples below, together with discussion of how to incorporate sample weights, and how to calculate the individual-level Gini from household-level data.

Not surprisingly in view of (3.7b), the social welfare function associated with the Gini coefficient is one in which the x's are weighted by the ranks of each individual in the distribution, with the weights larger for the poor. Since the Gini lies between zero and one, the value of social welfare in an economy with mean μ and Gini coefficient γ is $\mu(1-\gamma)$, a measure advocated by Sen (1976a) who used it to rank of Indian states. The same measure has been generalized by Graaff (1977) to $\mu(1-\gamma)^{\sigma}$, for σ between 1 and 0; Graaff suggests that equity and efficiency are separate components of welfare, and that by varying σ we can give different weights to each (see also Atkinson 1992 for examples).

The coefficient of variation is the standard deviation divided by the mean, while Theil's entropy measure is given by

$$I_T = \frac{1}{N}\sum_{i=1}^{N}\frac{x_i}{\mu}\ln\left(\frac{x_i}{\mu}\right). \tag{3.8}$$

I_T lies between 0, when all x's are identical, and $\ln N$, when one person has everything. This and other measures are discussed at much greater length in a number of texts, for example, Cowell (1995) or Kakwani (1980).

The choice between the various inequality measures is sometimes made on grounds of practical convenience, and sometimes on grounds of theoretical preference. On the former, it is frequently useful to be able to decompose inequality into "between" and "within" components, for example, between and within regions, sectors, or occupational groups. Variances can be so decomposed, as can Theil's entropy measure, while the Gini coefficient is not decomposable, or at least not without hard-to-interpret residual terms (see, for example, Pyatt 1976). It is also sometimes necessary to compute inequality measures for magnitudes—such as incomes or wealth—that can be negative, which is possible with the Gini or the coefficient of variation, but not with the Theil measure, the variance of logarithms, or the Atkinson measure. Further theoretical refinements can also be used to narrow down the choice. For example, we might require that inequality be more sensitive to differences between the poor than among the rich (see Cowell 1995), or that inequality aversion be stronger the further we are away from an equal allocation (see Blackorby and Donaldson 1978). All of these restrictions have appeal, but none has acquired the universal assent that is accorded to the principle of transfers.