Volume Title: Inquiries in the Economics of Aging

Volume Author/Editor: David A. Wise, editor

Volume Publisher: University of Chicago Press

Volume ISBN: 0-226-90303-6

Volume URL: http://www.nber.org/books/wise98-2

Publication Date: January 1998

Chapter Title: Measuring Poverty among the Elderly

Chapter Author: Angus S. Deaton, Christina Paxson

Chapter URL: http://www.nber.org/chapters/c7086

Chapter pages in book: (p. 169 - 204)

6 Measuring Poverty among the Elderly

Angus Deaton and Christina Paxson

6.1 Introduction

In the United States in 1992, there were four million elderly adults who were officially classified as poor. There were 31 million elderly in the United States in 1992, so that the poverty rate was just under 13 percent. Children were much more likely to be poor than the elderly; 22 percent, or 15 million children, were poor. This paper is about where such numbers come from and what (if anything) they mean. The data used to make the official calculations do not tell us anything about individual poverty. Instead they provide information on the income of *families*, information that is used to construct a set of poverty counts about individuals. The transformation from families to individuals makes many assumptions, about the allocation of resources within the household, about the differential needs of children, adults, and the elderly, and about the extent of economies of scale. Given the data, the effect of these assumptions on the poverty count depends on living arrangements, on how people combine to form families, on whether people are married or cohabit, on whether the elderly live by themselves or with other, younger adults.

In this paper we examine how living arrangements affect poverty measurement among the elderly in the United States, highlighting the importance of living arrangements. In the United States in 1992, 32 percent of those aged 65 or over lived alone, a further 42 percent lived in all-elderly families, and 26

Angus Deaton is the William Church Osborn Professor of Public Affairs and professor of economics and international affairs at Princeton University and a research associate of the National Bureau of Economic Research. Christina Paxson is professor of economics and public affairs at Princeton University and a research associate of the National Bureau of Economic Research.

Thanks are due to seminar participants at the conference, to the discussant, Doug Bernheim, to seminar participants at RAND, and to Bo Honoré. This research was supported by the National Institute on Aging through grants P01-AG05842 and R01-AG11957. The opinions in this paper do not necessarily reflect the views of the NBER or the sponsoring organization.

percent lived in families with at least one nonelderly person. Conclusions about the living standards of the elderly—absolute or relative—are determined not only by the data but also by assumptions about who gets what and how poverty lines vary with household composition. We demonstrate this fact by calculating the sensitivity of poverty counts to key assumptions in their construction, we examine the basis for the assumptions, and we explore whether the empirical evidence has anything useful to contribute.

There are two problems in passing from family resources to individual welfare, one of which is the main topic of this paper. The first issue, on which we shall have little to say, is the intrahousehold allocation of resources. The measurement of individual poverty requires a rule for assigning a welfare level to an individual, based on the consumption or income level of the family or household in which he or she lives. Any rule inevitably contains implicit assumptions about how resources are shared between different household members, for example, by age or sex. The issue is at its most acute in developing countries, where the elderly typically live in extended families. For example, in India where 90 percent of them live with other people, it is impossible to use household survey data to make simple statements about the resources available to elderly people. The issue is of rather less importance in the United States, where a much larger fraction of the elderly live alone or with other elderly adults.

The second problem is the one to which we give most attention here. Even if resources are distributed equitably across household members, the size and age structure of households affects the welfare levels of their members. The same level of income or income per capita does not give the same standard of living to a large family as to a small one, or to an all-adult household compared with one with children. Larger households may be able to take advantage of "scale economies" through shared consumption of public goods in the household, so that members of large households are likely to be better off than those of small households, even controlling for per capita income or total expenditure. Likewise, if children cost less than adults, then households with more children will require lower incomes to achieve a specified standard of living, given total household size. These issues are likely to be of particular importance when comparing poverty rates across age groups and are also likely to play out differently in countries with different living arrangements for elderly individuals. In the United States, where the elderly typically live in small households with few children, the treatment of child costs is unlikely to have large effects on the numbers of old people in poverty, although it can potentially have large effects on the poverty of the old relative to the young. Even when old people live alone, so that we can measure their resources from a household survey, we cannot classify them as poor or nonpoor without a standard of comparison, a standard that cannot be derived without assessing the needs of other nonelderly members of the population. The treatment of scale economies is likely to be an important issue for both absolute and relative poverty rates of the elderly.

Our first task is to examine the sensitivity of poverty measures in the United States to assumptions about child costs and scale economies. We then attempt to estimate the size of scale economies and child costs. We proceed as follows. In section 6.2 we begin by describing how official poverty measures are derived, and we present official poverty counts and rates for members of different age groups. The results are based on data from the 1993 Current Population Survey, which records information on 1992 incomes. We show that poverty measures for different age groups are quite sensitive to the treatment of scale economies and costs of children. In section 6.3 we focus on the measurement of scale economies and child costs, at a theoretical level and using data from the 1990 Consumer Expenditure Survey.

6.2 Poverty and Age

We start with a brief summary of how official poverty counts are derived in the United States. Further information can be obtained from Ruggles (1990) and National Research Council (1995).

In the United States, an individual is said to be poor if he or she lives in a family whose total income falls below a poverty line, where the poverty line depends on the size and age structure of the household. Thus, all individuals within a family are either poor or nonpoor: the implicit assumption is that resources are allocated fairly within households, so that all members have identical welfare levels. Official poverty statistics are based on information from the March Current Population Surveys (CPSs), which are conducted annually and contain information on the demographic characteristics and incomes of households. Income is before tax and includes cash transfer payments from the government but does not include nonmonetary transfers such as food stamps, health benefits, and subsidized housing. Several details deserve mention:

(i) The CPS has a structure that distinguishes between *households* and *families*. Although most households contain only one family, some contain multiple families, and others contain families that live with unrelated individuals. Of the 58,970 households surveyed in the 1993 CPS, 5,763 were not single-family households. For the purposes of measuring poverty, the Census Bureau combines the incomes of families and *related* subfamilies. For example, a married couple living with a son and daughter-in-law would be treated as a single family: their income would be combined and compared to the poverty line for a four-person family to determine whether all household members are in poverty. However, unrelated subfamilies and individuals are not considered part of the main family, and their poverty measures are based on their own incomes. For example, a person who boards with an unrelated family would be in poverty only if his or her personal income fell below the poverty line for a single person. The poverty line would not be adjusted for the fact that such a person, because he or she does not live alone, may benefit from scale economies in consumption.

This definition of families has some odd consequences. For example, unmarried couples who live together are treated as separate families. Each is assigned the poverty line of a single person, and each is defined to be poor if his or her personal income is below the line. This treatment of unmarried couples tends to increase poverty rates over what they would be if unmarried couples were treated the same as married couples, for two reasons. First, the U.S. poverty lines assume large economies of scale, so that the poverty line for a married couple is less than twice (in fact, only 1.29 times) the line for a single adult. Two unmarried adults who live together and are both in poverty can potentially move themselves out of poverty simply by getting married. Second, even with no scale economies built into poverty lines, treating unmarried couples as if they are married will reduce poverty counts. Combining the incomes of unmarried couples imposes a mean-preserving reduction in the spread of income per person, which can be expected to pull some observations out of the tails and thus lower poverty rates. This is exactly the same argument (in reverse) as in Haddad and Kanbur (1990), who show that poverty measures will be understated if it is assumed that resources are allocated equally across family members when in fact they are not. The fact that official U.S. poverty measures depend on legal marital status is worrisome, especially given the increase in the numbers of coresiding unmarried couples over the past several decades.

In what follows, we present poverty measures that are based on the official definition of a family, but we also present measures that are based on the total incomes of all members of a household (whether or not the members are related). For these latter computations we use the poverty line that would be applied if all household members were related. Since we are interested in how household scale economies affect living standards, it makes sense to use the household (i.e., individuals living together in the same quarters) rather than the family (i.e., those related by blood or marriage) as the unit of analysis. The switch from a family to a household basis is also one of the recommendations of the National Academy study on the poverty line.

(ii) The United States has used essentially the same set of poverty lines since the 1960s, adjusting them only for the effects of inflation. The lines depend on the size and age structure of households. They were originally calculated as the cost of the U.S. Department of Agriculture's "low-cost food plan" for households of different sizes and age structures, multiplied by three. The multiplier of three was selected because it was the inverse of the average budget share for food in a 1955 Department of Agriculture survey. Table 6.1 shows a subset of the poverty lines used in 1992 (and applied to the 1993 CPS data, which refer to 1992); to save space we have trimmed off poverty lines for families with more than six members and families with more than four children.

Table 6.1 U.S. Poverty Lines, 1992

	Related Children under 18 Years				
Size of Family	None	One	Two	Three	Four
One person					
Under 65 years	7,299				
65 years and over	6,729				
Two persons, with householder					
Under 65 years	9,395	9,670			
65 years and over	8,480	9,634			
Three persons	10,974	11,293	11,304		
Four persons	14,471	14,708	14,228	14,277	
Five persons	17,451	17,705	17,163	16,743	16,487
Six persons	20,072	20,152	19,737	19,339	18,747

Source: These poverty lines are reproduced from U.S. Bureau of the Census, *Poverty in the United States: 1992* (Washington, D.C., 1993), table A, p. vii.

Several features of the table are important for our analysis. First, the poverty lines for one- and two-person families depend on whether the reference person in the family is elderly.[1] The implicit assumption embodied in the table is that older people need less money to achieve a given welfare level, and the poverty line of an elderly person living alone is $570 less than that of a nonelderly person living alone, a "discount" equal to 7.8 percent of the poverty line for the nonelderly person. Note, however, that this adjustment to the poverty line is made only if the reference person is elderly, and families containing elderly who are not the reference person receive no adjustment to their poverty lines. Second, the poverty lines in the table make implicit adjustments for the costs of children and household size. As seems sensible, increases in family size result in less-than-proportional increases in the poverty line, holding the number of children constant. Increases in the number of children, holding family size constant, also affect the poverty line, but in neither a simple nor a reasonable fashion. For example, the poverty line for a three-person family with no children is less than that of a three-person family with one child. In this case the substitution of a child for an adult increases the poverty line by 2.9 percent. The implication is that children cost *more* than adults, which seems odd given that the poverty lines are based on baskets of food. The poverty lines for families with four people are even stranger; the line rises as we go from a four-adult family to a three-adult one-child family, falls as we go to a two-adult two-child family, and then increases for a one-adult three-child family. Overall,

1. In the CPS the term "reference person" is synonymous with the term "householder." This person is defined in the CPS documentation as "the person (or one of the persons) in whose name the housing unit is owned or rented (maintained) or, if there is no such person, any adult member, excluding roomers, boarders, or paid employees." If a married couple jointly owns or rents a home, either person may be designated as the reference person.

the poverty lines in table 6.1 contain the implicit assumptions that, given family size, old people are relatively cheap and children are relatively expensive. However, the poverty lines also imply economies of scale, so that the high expense of children may be offset by the fact that children tend to live in larger families. These assumptions about costs of children and old people, and about scale economies, will affect any conclusions about the poverty of the old relative to the young. We will return to these aspects of the table below, but they are useful to keep in mind when comparing official poverty rates across age groups.

Table 6.2 provides estimates for the United States of the numbers and fractions of elderly people, nonelderly adults, and children less than 18 years of age in poverty, using the official poverty lines. The top panel uses family income to measure poverty, and the bottom panel uses household income. Individuals living in "group quarters" were excluded for these and all subsequent calculations. Using official poverty lines, there are 10.9 million more poor children than poor elderly adults. Poverty rates for children are also higher (22.14 percent vs. 12.87 percent), and poverty rates for elderly and nonelderly adults are very close. This is true whether we use the family or the household as the unit of analysis, although as is to be expected, poverty rates using household-level data are slightly lower for all age categories.

Given the arbitrariness and controversy surrounding the choice of poverty lines, it is important to examine the extent to which these results are sensitive

Table 6.2 **U.S. Poverty Rates for Elderly Adults, Nonelderly Adults, and Children**

	All People	Children under 18	Adults 18–64	Adults 65+
		Family-Level Data		
Sample size	154,977	42,869	93,069	19,039
Estimated population (thousands)	253,924	67,062	156,035	30,827
Number poor (thousands)	36,987	14,846	18,174	3,968
Percentage in poverty	14.57	22.14	11.65	12.87
		Household-Level Data		
Number poor (thousands)	33,311	13,800	15,696	3,815
Percentage in poverty	13.12	20.58	10.06	12.38

Source: The data used for this table are from the March 1993 Current Population Survey, which records 1992 income. The sample consists of all individuals who do not live in group quarters. The population weights from the "person" records were used to obtain population estimates. A person was defined to be poor if his or her family income (*top panel*) or household income (*bottom panel*) was below the relevant poverty line for his or her family (*top panel*) or household (*bottom panel*).

to how the lines are set. A useful way to address the question is to graph cumulative distribution functions for income for different groups (children, nonelderly adults, and elderly adults) and see whether the distribution of living standards of some groups first-order stochastically dominate those of other groups. If the distribution of one group stochastically dominates that of another, then the poverty rate for that group will be lower at any poverty line. An immediate difficulty is that there is no single poverty line for the United States. Rather, we have a *set* of poverty lines for families of different sizes and composition. To draw the cumulative distributions we must first "rebase" family income to make it comparable across families with different sizes and numbers of children. Specifically, we choose families with two adults and two children as the "base," with a poverty line of $14,228 from table 6.1. The living standard of person i in family f is then measured as

$$y_{if}^{*} = y_f(z_b/z_f),$$

where y_f is the income of family f, z_f is its poverty line, and z_b is the poverty line of the base family. We then graph cumulative distribution functions for y_{if}^{*} for children, nonelderly adults, and elderly adults.

The results of this exercise (which uses the family rather than the household as the unit of analysis) are shown in figure 6.1. The left-hand panel graphs the cumulative distributions for the three groups. Since we are interested in whether the poverty counts are sensitive to choice of poverty line, we graph the distribution functions only up to $20,000, so any conclusions about stochastic dominance are restricted to this range of income. The "base" poverty line of $14,228 is shown as a point of reference. The right-hand panel of the figure graphs two lines: the vertical difference between the cumulative distributions for elderly and nonelderly adults and the vertical difference between the distributions for children and nonelderly adults. The figure shows that the living standards of adults first-order stochastically dominate those of children. The distribution functions for the elderly and nonelderly adults cross at an income

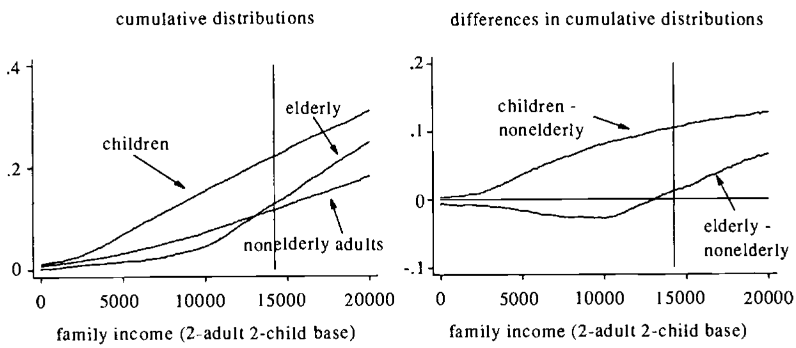

Fig. 6.1 Cumulative distributions of family income by age group

level slightly below the poverty level, indicating that at lower poverty lines the elderly would be less poor than the nonelderly. This may reflect the fact that social security provides an income floor for most elderly citizens.

The key question is how sensitive these results are to treatments of child costs and scale economies different from those implicit in the official poverty lines. One straightforward exercise is to examine how the numbers change if the old-age "discount" to poverty lines is removed, so that one- and two-person families with an elderly reference person receive the same poverty line as do those with younger reference people. The results of these tabulations are shown in table 6.3. When the unit of analysis is the family, the number of elderly in poverty increases by 716 thousand people, from 4 to 4.7 million people, and the rate of poverty among the elderly increases from 12.87 percent to 15.19 percent. There are only tiny increases in poverty among other groups, reflecting the fact that the old-age discount applies to only one- and two-person families with an elderly reference person and most elderly people in two-person families live with other elderly people.

The poverty lines in table 6.1 contain strong implicit assumptions about the costs of children and economies of scale, and it is useful to examine how poverty rates for different age groups change as the assumptions are modified. However, making ad hoc adjustments to the table of poverty lines does not seem especially desirable. The approach we take is to specify the poverty line as a function of the age and size composition of the family and fit this function to the official poverty lines. We can then examine how sensitive poverty measures are to child costs and scale economies by altering the parameters of the function. Specifically, we specify the poverty line for a family with A adults and K children as

Table 6.3 U.S. Poverty Rates with the Old-Age Discount Removed from Poverty Lines

	All People	Children under 18	Adults 18–64	Adults 65+
		Family-Level Data		
Sample size	154,977	42,869	93,069	19,039
Estimated population (thousands)	253,924	67,062	156,035	30,827
Number poor (thousands)	37,745	14,848	18,213	4,684
Percentage in poverty	14.86	22.14	11.67	15.19
		Household-Level Data		
Number poor (thousands)	34,062	13,800	15,736	4,526
Percentage in poverty	13.41	20.58	10.08	14.68

$$z(A, K) = z_b(A + \alpha K)^\theta, \tag{1}$$

where the term $(A + \alpha K)^\theta$ represents the number of "adult equivalents" in the household and z_b is the per capita poverty line for a base household. In what follows we refer to the term $A + \alpha K$ as the "effective household size." The parameter α measures the cost of children relative to adults, and the parameter θ reflects economies of scale. For example, a value of θ equal to 0.5 implies that doubling effective household size $A + \alpha K$ multiplies the poverty line by the square root of 2, or 1.41.

We take logarithms of equation (1) and fit the resulting equation to the actual data on poverty lines, numbers of adults, and numbers of children. Using the complete list of poverty lines, we obtain the estimates $\alpha = 1.3$ and $\theta = 0.47$. The high value of α reflects the fact that for smaller families the actual poverty lines increase as the numbers of children increase and the majority of families in our sample are small: 77.7 percent of families have fewer than four members, and less than 9 percent have more than four members. When we establish a baseline by recomputing age-specific poverty rates using the poverty lines estimated using equation (1), with $\alpha = 1.3$ and $\theta = 0.47$, we obtain estimates very similar to those using the actual poverty lines: the poverty rate among children is 22.0 percent, among nonelderly adults is 11.5 percent, and among elderly adults is 14.4 percent. That the poverty rate among elderly adults is somewhat higher than the "official" rate of 12.87 percent in table 6.2 is not surprising, given that we have not built an old-age discount into our estimated poverty lines.

We are now in a position to examine how changes in the costs of children, as measured by α, and changes in scale economies, as measured by θ, affect the poverty of the old relative to other age groups. We present these results graphically, by showing cumulative distributions of living standards for people in different age groups, where living standards are measured as

$$y_{if}^* = y_f \frac{z(2,2)}{z(A_f, K_f)} = y_f \frac{(2 + \alpha 2)^\theta}{(A_f + \alpha K_f)^\theta}; \tag{2}$$

as before, our "base" family has two adults and two children, so that for families of this type living standards are simply measured as family income.

Before proceeding, it is useful to think about how we might expect changes in α and θ to affect the living standards of individuals relative to those in base households. Taking derivatives of the logarithm of equation (2) with respect to α yields

$$\frac{\partial \ln y_{if}^*}{\partial \alpha} = \theta \left(\frac{2}{2 + \alpha 2} - \frac{K_f}{A_f + \alpha K_f} \right), \tag{3}$$

which is positive for households with smaller ratios of children to effective household size than the base household and negative for others. Increases in the cost of children raise the relative living standards of households with small fractions of members who are children and reduce the relative living standards of those with large fractions. The derivative with respect to θ is

$$\frac{\partial \ln y_{if}^*}{\partial \theta} = \ln\left(\frac{2 + \alpha 2}{A_f + \alpha K_f}\right), \tag{4}$$

which implies that increases in the scale parameter θ will raise the relative living standards of those in families with small numbers of effective members and lower the relative living standards of those in large families. Thus, the living arrangements of the elderly—the size of their families as well as the ratios of children to household size—will determine the effects that changes in child costs and scale economies have on living standards.

Table 6.4 presents evidence on the family sizes and fractions of members who are children for people of different ages. The elderly live in households that are on average smaller than those of children and nonelderly adults. The average household size for an elderly person is 1.94, as opposed to 3.04 for a nonelderly adult and 4.38 for a child. In consequence, poverty measures that assume greater household scale economies will increase the poverty of the old relative to the young, with the primary beneficiaries of scale economies being children. Likewise, the average ratio of children to family size for elderly people is only 0.017, far lower than the average ratios of 0.203 for younger adults and 0.533 for children. Poverty measures that assume smaller costs of children will also increase the poverty of the elderly relative to children and also increase their poverty relative to the younger adults with whom children reside.

Table 6.4 U.S. Average Family and Household Composition

	Children	Nonelderly Adults	Elderly Adults
	Family-Level Data		
Number of children	2.39	.86	.08
Number of nonelderly adults	1.95	2.10	.36
Number of elderly adults	.04	.08	1.50
Children + adults	4.38	3.04	1.94
Children/(children + adults)	.533	.203	.017
	Household-Level Data		
Number of children	2.42	.90	.09
Number of nonelderly adults	2.02	2.23	.38
Number of elderly adults	.037	.08	1.51
Children + adults	4.81	3.21	1.98
Children/(children + adults)	.524	.208	.018

Note: The table shows averages, across all people in the relevant age category, of family (*top panel*) and household (*bottom panel*) characteristics. For example, children live in families that contain, on average, 2.39 children.

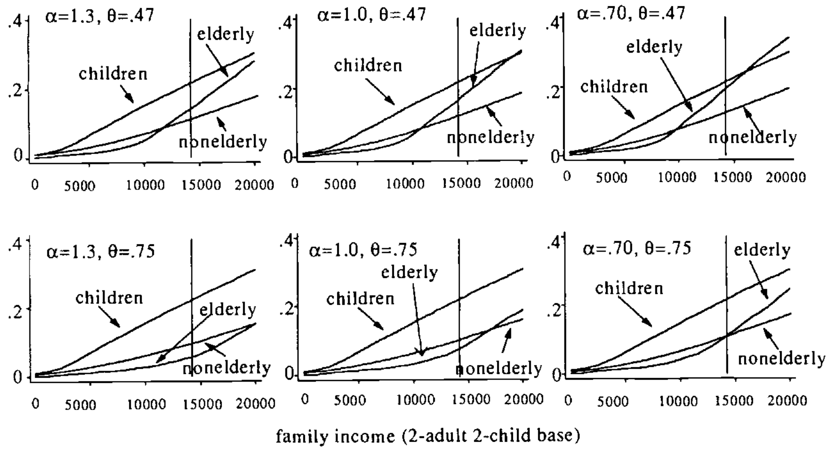

Fig. 6.2 Cumulative distributions of income with changes in child costs and scale economies

Figure 6.2 shows the effects of altering α and θ on the cumulative distributions of family income in the United States. The top left-hand graph shows the distributions with α = 1.3 and θ = 0.47, the values obtained using the official poverty lines. Consistent with figure 6.1, the distribution of living standards for children is stochastically dominated by that of elderly and nonelderly adults, and the lines for elderly and nonelderly cross near the official poverty line for the base family. As one moves from left to right, α is decreased, from 1.3 to 1.0 to 0.7. As expected, increases in child costs make the elderly less well off relative to children, and at α = 0.7 the cumulative distributions for the elderly and for children cross at an income level only slightly above the official poverty line. The three graphs in the bottom panel also let α vary from 1.3 to 0.7, but use a scale parameter θ equal to 0.75 instead of 0.47. Scale economies are less important at higher values of the scale parameter, and as expected the elderly become better off relative to others due to their smaller average family sizes. The information in figure 6.2 is summarized in figure 6.3, which shows the *differences* in the cumulative distributions similar to those shown in figure 6.1.

The effects of changes in α and θ can also be seen in table 6.5, which shows the fraction of individuals with values of y^*_{if} less than the base poverty line of $14,228. Holding θ fixed at 0.47, the poverty rate of the elderly increases from 14.40 to 19.60 and the poverty rate of children declines from 22.02 to 21.61 as the cost of children declines from 1.3 to 0.7. It may seem counterintuitive that when child costs are lowered, the poverty rates of children fall very little and the poverty rates of the elderly rise substantially. It should be kept in mind, however, that we have used a two-adult two-child family as the base and that the incomes of base families are unaltered by changes in α and θ. Since a large

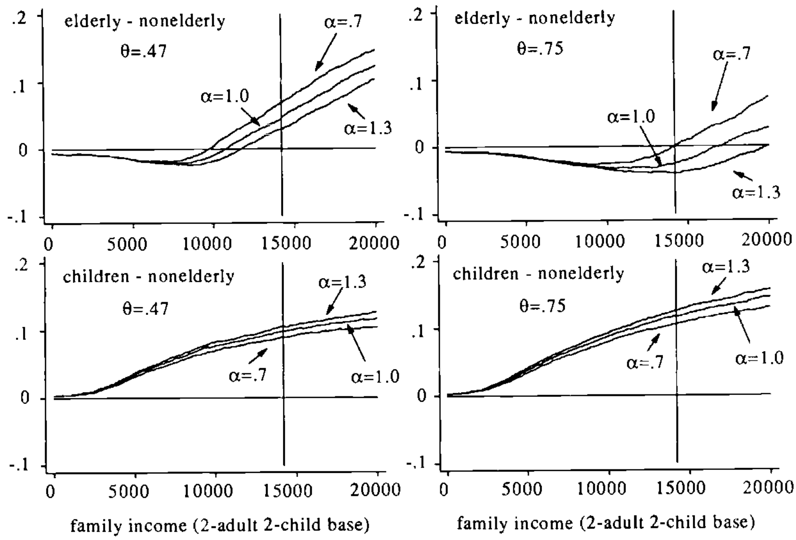

Fig. 6.3 Differences in cumulative distributions with different assumptions about child costs and scale economies

Table 6.5 U.S. Poverty Rates Using Different Base Household Types

	Two-Adult Two-Child Base			One-Adult Base		
Child Costs and Scale Economies	Children	Nonelderly Adults	Elderly Adults	Children	Nonelderly Adults	Elderly Adults
α = 1.3, θ = 0.47	22.02	11.46	14.40	23.04	12.28	16.08
α = 1.0, θ = 0.47	21.84	11.97	16.54	21.60	11.81	15.97
α = 0.7, θ = 0.47	21.61	12.66	19.60	19.49	11.11	15.90
α = 1.3, θ = 0.75	22.36	9.75	5.79	35.51	18.21	20.29
α = 1.0, θ = 0.75	22.07	10.29	7.72	31.86	16.89	20.22
α = 0.7, θ = 0.75	21.62	11.03	11.04	27.90	15.48	20.12

fraction of children live in two-adult two-child families, we are effectively guaranteeing that their poverty rates change little as α is increased. The more general point is that the choice of base family has large effects on the measured levels of poverty. This is illustrated by the right-hand panel of table 6.5, which repeats the exercise using a one-adult no-child base. In this case, increases in α reduce poverty rates for all age groups, since there are people in all age groups who live with children, but reduce the poverty rate for children the most.

Table 6.5 also shows the effects of changes in scale economies on relative poverty rates. In general, increases in θ from 0.47 to 0.75 (holding α fixed) increase the poverty of children relative to elderly adults. This is to be expected, given that children live in families that are on average larger than those of other age groups. Again, however, the effect of changes in the scale parameter on absolute poverty rates depends on the definition of the base family. When the base is a two-adult two-child family, the poverty rate of children barely changes, while adult poverty rates decline. Using a one-adult base, increases in the scale parameter result in large poverty increases for all groups, with the largest increase for children.

It is also useful to compare changes in poverty rates as one moves from the case in which both child costs and scale economies are high (i.e., $\alpha = 1.3$ and $\theta = 0.47$) to a possibly more realistic case of lower child costs and scale economies (i.e., $\alpha = 0.70$ and $\theta = 0.75$). These changes in α and θ have offsetting effects, since large families tend to have more children. Using a two-adult two-child base, the poverty rates of children and nonelderly adults are largely unchanged, and that of the elderly declines by less than 4 points. Thus, it is possible that modifying the official poverty lines so that they imply lower child costs *and* lower scale economies will leave aggregate poverty rates for different age groups largely unchanged. This does not imply, however, that the same *people* will be in poverty under the two scenarios.

The conclusion of this section is that the age distribution of poverty is quite sensitive to the treatment of child costs and economies of scale. Official estimates of the numbers and percentages of people in different age groups in poverty hinge on poverty lines that contain strong implicit assumptions about child costs and scale economies. Deviations from the assumptions implicit in official poverty lines result in very different poverty measures for different age groups.

In section 6.3, we examine the empirical evidence on household consumption patterns in a (not very successful) attempt to establish a less arbitrary basis for the costs of children and household economies of scale. As we shall see, there are good reasons to be skeptical of the results. Nevertheless, to the extent that there are results, they are consistent with a value of α around 0.75 and a value of θ of 0.85 or even higher. Such figures generate poverty counts close to those shown in the last row of table 6.5, with $\alpha = 0.70$ and $\theta = 0.75$.

6.3 Calculating Equivalence Scales

6.3.1 Introduction

The official poverty lines in the United States differ for different family sizes as shown in table 6.1. These differentials come from differences in food purchasing patterns in the 1955 survey that Mollie Orshansky used to construct the original poverty line. In consequence, the equivalence scales implicit in the table have some basis in behavior, although, as we have seen, the pattern of child costs and economies of scale is not easy to defend. Several alternatives to the Orshansky scales have been proposed from time to time. For example, Ruggles (1990) suggests that the scale be proportional to the square root of household size, so that two adults are equivalent to 1.41 single adults, three adults to 1.73 single adults, and so on. As discussed in section 6.2, this rule approximates the Orshansky scale while removing some of its stranger features. Larger families are also typically those with larger numbers of children, and the square root rule should be seen as reflecting not only economies of scale but also that children cost less than adults. It is also possible explicitly to separate adults and children, and the recent report on the poverty line by the National Academy (National Research Council 1995) recommends that scales be calculated according to the formula used above, namely, $(A + \alpha K)^{\theta}$ for A adults and K children, and recommends values around 0.75 for both α and θ. Alternatively, the Organization for Economic Cooperation and Development (1982) has suggested a scale in which the first adult counts as unity, other adults as 0.7, and children as 0.5, so that economies of scale are not allowed for explicitly but are reflected in the discount for additional adults and a relatively low cost of children.

In this section, we use data from the Consumer Expenditure Survey (CEX) in 1990 to try to tease out information about the relative costs of children and adults, and about the extent of economies of scale. That these data should be relevant for such an exercise seems obvious. The CEX collects information on purchases of goods that are differentially consumed by adults and by children and can be used to infer how these purchases change with the composition of the family. Economies of scale presumably arise because some goods are public and can be shared by several members of the household, and an expenditure survey can tell us how the balance between private and public goods shifts with household size and composition. Even so, attempts to infer equivalence scales from budget data have a long and discouraging history. Although expenditure data are *relevant* to the construction of scales, they are insufficient by themselves to *identify* scales. Additional identifying assumptions are required, and it has been difficult to find such assumptions that are widely acceptable. Worse still, many of the estimates in the literature have been derived under identifying assumptions that are not made explicit, so that it is difficult or impossible to know what the results mean.

We devote subsection 6.3.2 to the theoretical issues: to the identification issue and what it does and does not imply and to the basis for our calculations of child costs and economies of scale. The former rests on familiar ground and calculates the costs of *children* by comparing the behavior of expenditure on *adult goods* across households with different numbers of children. The latter uses a new method, suggested to one of us by Jean Drèze, that uses expenditures on *private goods* to identify the economies of scale to household size that operate through *public goods*.

6.3.2 Identifying Equivalence Scales

We begin by explaining what cannot be done, and the consequences of trying. The theory of equivalence scales parallels that of cost-of-living index numbers, and both begin from a description of household preferences in terms of the cost or expenditure function. If total household expenditure is x, and if this is efficiently spent to maximize the collective utility of a family of A adults and K children, say, then we can write

$$c(u, p, A, K) = x, \tag{5}$$

where u is the collective utility level. By analogy with cost-of-living numbers, the number of single adult equivalents in a family with A adults and K children is (see Deaton and Muellbauer 1980, 205)

$$E(A, K; u, p) = \frac{c(u, p, A, K)}{c(u, p, 1, 0)}, \tag{6}$$

an expression that, in general, will depend on prices and on the level of real income of the family. Since the cost function provides a convenient summary of the family's preferences, and since once specified it can be used to yield a set of demand functions, there is an apparently straightforward way of calculating the equivalence scale in equation (6). Write down some suitable cost function (eq. [5]), for example, a translog or almost ideal system, recover the parameters by estimating the associated demand functions, and fill in equation (6). This is what is done, for example, by Slesnick (1993) in his recalculation of poverty in the United States.

Pollak and Wales (1979) have shown that such calculations of equivalence scales ignore the central identification problem that the data do not identify the equivalence scales. In particular, Pollak and Wales point out that all cost functions of the form

$$c(\phi(u, A, K), p, A, K) \tag{7}$$

have the same demand functions irrespective of the function ϕ provided only that it is monotone increasing in u. Relabeling indifference curves in a way that depends on family structure has no effect on demand functions, although it clearly changes how we should measure welfare. As a result, the data do not allow us to calculate equivalence scales without additional assumptions that pin down the function ϕ.

At first blush, it is tempting to associate the identification issue with the fact that parents get utility from children and to try to separate "economic" welfare as represented from the structure of the cost function from the "psychic" welfare that is captured by the function ϕ. But this hope is not only imprecise, it is wrong, as the following example shows. Consider families that contain only adults, and suppose that we model economies of scale by writing

$$c(u, p, A) = A^{\theta} c(u, p, 1), \tag{8}$$

so that for $\theta \leq 1$ there are economies of scale; we want to estimate θ. To fix ideas, suppose that the cost function for the single-adult household takes the "quasi-homothetic" form in which costs are a linear function of utility:

$$(9) \qquad c(u, p, A) = A^{\theta}c(u, p, 1) = A^{\theta}a(p) + A^{\theta}b(p)u \,,$$

where $a(p)$ and $b(p)$ are linearly homogeneous functions of the vector of prices p. The demand functions are derived from equation (9) in the usual way; the budget shares take the form

$$(10) \qquad w_i = \frac{p_i q_i}{x} = \alpha_i(p)\frac{a(p)A^{\theta}}{x} + \beta_i(p)\left(1 - \frac{a(p)A^{\theta}}{x}\right),$$

where $\alpha_i(p)$ and $\beta_i(p)$ are the elasticities with respect to the ith price of $a(p)$ and $b(p)$, respectively. (That the budget shares for a family of A adults with outlay x are the same as the budget shares for a family with one adult and outlay x/A^{θ} is a general feature of cost functions of the multiplicative form [8].) According to equations (9) and (10), preferences and the associated behavior are each weighted averages of preferences at subsistence $a(p)$ with associated budget shares $\alpha_i(p)$ and preferences at bliss $b(p)$ with associated budget shares $\beta_i(p)$. Given specifications for $a(p)$ and $b(p)$, the demand functions (10) can be taken to the data, and θ estimated together with the other parameters; see, for example, Lanjouw and Ravallion's (1995) work for Pakistan, where the parameter θ is estimated to be around 0.6.

Economies of scale in equation (8) operate multiplicatively through the cost function, so that in equation (9) costs at subsistence and costs at bliss are subject to the same economies. But because of fixed costs or other effects, the economies of scale might be more or less effective at higher levels of living. One way of generalizing equation (10) to allow this would be to write

$$(11) \qquad c(u, p, A) = A^{\theta}a(p) + A^{\psi}b(p)u \,,$$

where θ and ψ are not necessarily equal. If we now derive the demand functions from these new preferences, either by elementary calculation or by application of the Pollak and Wales (1979) theorem, we once again retrieve equation (10). Hence, as far as their empirical implications are concerned, equations (9) and (11) are indistinguishable. As a corollary, if we take equation (10) to the data and obtain an estimate of θ, it has no particular claim on our attention as a measure of the extent of economies of scale. The scales used in Slesnick's (1993) reworking of the poverty count are obtained by a generalization of the technique outlined above and are subject to the same criticism. Slesnick's and Lanjouw and Ravallion's estimates are identified by assuming that the particular cardinalization of utility that they chose is correct, rather than the infinite number of other cardinalizations that are indistinguishable on the data but that would give different results for the scales. Indeed, since Slesnick's results are perhaps even more bizarre than those in the official scales, it is comforting to know that they can be arbitrarily "corrected" without consequences for the empirical evidence from which they were obtained.

By itself, the empirical evidence on expenditure patterns cannot generate equivalence scales. Instead, we need additional assumptions, typically in the form of prior information that links welfare to behavior. It is to these we now turn.

6.3.3 Alternative Identification Schemes for Child Costs and Economies of Scale

The most famous and venerable of the schemes for linking behavior and welfare is Engel's supposition that families with the same share of food in their budgets are equally well off, irrespective of size and composition. The food share assumption serves as an all-purpose scale identifier, allowing us to measure both child costs and economies of scale. For the former, we start from a two-adult family (say) and calculate for any given level of outlay the additional amount that would be required to bring a two-adult one-child family to the same food share, which by assumption would be at the same level of welfare. We can equally well equate food shares for a large and a small family so as to calculate the relative outlays that make them equally well off and so estimate any economies of scale. If the cost function is multiplicatively separable in family characteristics on the one hand and prices and utility on the other—as is equation (8)—the food share, or any other budget share, does indeed indicate welfare, so the method is consistent with a well-defined theoretical structure. Furthermore, the assumption that the food share indicates welfare solves the identification impasse; for example, it is easy to check that the food share indicates welfare if the cost function is equation (9), but not when the cost function is equation (11). The Engel assumption ties down the function ϕ in equation (7).

Engel's assumption is the *kind* of assumption that we need, but it is not the right one. As first argued by Nicholson (1976) and elaborated by Deaton and Muellbauer (1986), Engel's assumption is quite implausible. The addition of a child with full compensation would normally *increase* the food share, not leave it unchanged, in which case the Engel compensation is too large. A better procedure for measuring child costs uses expenditure on adult goods as an indicator of adult welfare and calculates the compensation that would be required after the addition of the child in order to restore adult expenditures to their previous level. This method was first proposed by Rothbarth (1943) and has been used in the United States by Betson (1990) and (more or less) by Lazear and Michael (1988). The Rothbarth procedure is not without its problems. In particular, it takes no account of possible substitution toward adult goods in the presence of additional children, or of broader rearrangements of consumption patterns that might follow the addition of a child. There is also a shortage of well-measured adult goods in the data. The procedure inevitably uses expenditures on adult clothing and footwear, alcohol, and tobacco, and the last two are much underestimated in the CEX—as in other expenditure surveys around the world—see Gieseman (1987). Even so, the Engel method is clearly wrong, and the transparency of Rothbarth's method and its identification assumptions are in sharp and favorable contrast with the mechanical estimation of equivalence scales from demand systems augmented by demographic variables. We present some results using Rothbarth's method below.

The calculation of economies of scale has a less well trodden history. Although the Engel method provides estimates of scale economies, there is no reason to credit them without some theory of why, in the presence of public goods, the food share should correctly indicate welfare between families with different numbers of people. Indeed, there are good reasons for supposing that the Engel method will give the wrong answer. The argument follows much the same lines as Nicholson's argument about Engel's method for calculating child costs, and we shall adopt a solution that is similar to the Rothbarth method in that case.

Suppose that there exists a pure private good that is not substitutable for public goods, food being the obvious example. Consider what happens when family size increases, for example, by combining two single adults into a family of two adults. Suppose too that compensation is paid, presumably negative compensation since the scale economies in the public good will make the family better off if the two original incomes are simply combined. Because less of the public good is needed than in the two separate single-person families, and because there is no substitution from the private to the public good, the budget share of the private good will increase. But according to Engel, a family with a higher food share must be worse off, and less money should be taken away. In consequence, Engel's method will overstate the amount of money needed for the larger household and thus understate the extent of and discount from economies of scale.

In our own empirical work in this paper, we adopt an approach to economies of scale that parallels Rothbarth's procedure for measuring child costs. We measure child costs by examining adult goods, and we measure economies of scale by examining the behavior of private goods as household size changes. Consider again the example of the previous paragraph in which two people, previously in single-person families, come together to form a two-person family. As before, they no longer need as much heat and light, kitchens, bathrooms, and (possibly) bedrooms so that they have more for private goods, whose share of the budget can be expected to increase. If nothing else happens, we could calculate the extent of the economies of scale, not by restoring the *share* of private goods to its previous value, but by calculating the reduction in total income that would restore the previous *per capita consumption* of private goods. If the reduction is (say) 20 percent of total outlay, we have established that there is a 20 percent discount for two people over one. The major caveat is that public goods are effectively cheaper in larger households; an oriental rug or a painting costs the same but provides pleasure to twice as many people. In consequence, there will be substitution away from private toward public goods, so that a fully compensated two-person family will have lower per capita consumption of private goods than will the two single-person families combined. As a result, the reduction in income that restores the per capita consumption of private goods is less than the utility-preserving reduction and so understates the extent of economies of scale. This effect will be small if private goods are not substitutable for public goods, as, for example, food for housing, or we can calculate the income reduction for a range of private goods and select the one that gives the largest reduction. Although both this method and the Engel (food share) method will understate economies of scale, the latter will do so even in the absence of substitution between private and public goods and will therefore understate by more.

Alternative methods and estimates of economies of scale have been provided by Nelson (1988) and in the National Academy report. The latter starts from the estimates of child costs obtained by the Rothbarth procedure in Betson (1990), who made separate calculations according to the number of adults in the family. His results can therefore be used to give the costs of an additional child relative to an adult in one-adult families, in two-adult families, and so on. If it is supposed that the number of adult equivalents is approximated by the formula $(A + \alpha K)^{\theta}$, with parameters the same for different household types, then the differences in child costs across different family types reflect economies of scale and can be used to measure them. In particular, in a family with A adults and K children, the cost of an additional child relative to an adult is given by the formula $[(A + \alpha(K + 1))/(A + \alpha K)]^{\theta}$, which can be fitted to Betson's results to obtain estimates of α and θ. The calculations in the report suggest values of around 0.75 for both α and θ. Unlike the Engel procedure, this method appears to be soundly based, at least if we accept the Rothbarth method of estimating child costs. Even so, it can provide only a summary of the scale economies and yields no insight into the process by which public goods work within the household. Nor does it ensure that we are measuring economies of scale rather than (for example) the possibility that larger households devote less to additional children for some other reason.

Nelson assumes that different goods exhibit different economies of scale, and she uses CEX data on all adult families to estimate a demand system that, following Barten (1964), allows explicitly for the patterns of substitution that come from the changes in effective relative prices that are induced by the differential economies of scale. Although the identification of welfare is left implicit, Nelson's procedures are theoretically consistent with the method adopted below. Her model assumes that economies of scale work so that the fractions of family consumption of each good received by each person can add up to a total different from unity, and that the sum of these fractions is different for different goods, being larger than unity for goods with economies of scale. Because she makes no allowance for fixed costs, economies of scale can only operate this way and must result in substitution effects in line with the price effects that the scaling factors are designed to mimic. She works with five broad groups of goods and finds that all display economies of scale; for housing there are implausibly large economies, with each member of a two-person household receiving twice as much in housing services as would each in a one-person household. Nelson attributes the result to the fact that, in her data, two-person households spend almost the same amount on housing as one-person households, in spite of having incomes that are 50 percent higher. It is also possible that the Barten model should bear some of the blame; high economies of scale act so as to diminish needs and thus consumption, but they also cause an offsetting substitution effect toward the good. When this price effect is substantial, and the data show that two people spend much the same on housing as one person, the model responds by choosing a very large value for the economies of scale. But it is also possible that the model is incorrect, and that economies of scale operate differently, as would be the case if there are fixed costs.

6.3.4 Using Adult Goods to Estimate Child Costs

We used data from the 1990 CEX to examine the behavior of expenditure on adult goods over families with differing numbers of children. Expenditure on adult goods is defined, as is usual, as the sum of expenditures on alcohol, tobacco, and adult clothing and footwear. The 1990 CEX has 20,517 observations in total, and we use all of these in the analysis. The CEX is a rolling panel, in which each household is (in theory) interviewed five times at three-month intervals and asked about consumption in the previous three months. New households are added as other households exit the survey. Because each interview generates an observation, there are only 10,127 different households in the sample. Of these, 3,893 are interviewed once, 3,053 twice, 2,206 three times, and 975 four times. The lower numbers of households with three and four observations are accounted for by the fact that some households miss interviews or drop out of the survey before completing all five interviews: 953 of our households do not complete the full series of interviews or have "gaps" between interviews.

We define total expenditure as the sum of expenditures on food, both away and at home, alcohol, tobacco, clothing, housing, transportation excluding purchases of new cars, entertainment, personal care, medical expenditures, reading materials, and educational expenditures. The Bureau of Labor Statistics computes these semiaggregated expenditure items from the more finely detailed expenditure information collected by the survey and puts them in the "summary expenditure section" of the CEX data tape. It should be noted that "housing" is broadly defined to include utilities, household furnishings, and household equipment. Furthermore, the cost of an owner-occupied house is measured as actual interest payments on mortgages and not as the rental value of the house. This definition of total expenditure is the conditioning variable in our Engel curves, and when we calculate the compensation required to restore adult expenditures to prechild levels, it is the increase in this measure of the total that is used. The validity of the procedure requires that additional children have no effect on total expenditure, so that we are ignoring effects of children on income through changes in labor supply, as well as on saving. To the extent that children *decrease* this definition of total expenditure—which seems the more probable outcome—the costs of children will be *understated* by the amount of the decrease.

In this section and the next, we start with a nonparametric analysis, to explore whether the data are at all consistent with the underlying hypothesis, and then use the results to suggest the more parametric forms that allow the inclusion of a wider range of covariates. For the analysis of adult goods and child costs, the nonparametric analysis adopts the obvious procedure of fitting adult good Engel curves for different family types, comparing households with and without children. The CEX sample will support such an analysis for a limited number of family types; table 6.6 lists the number of families by combinations of adults and children. Given these, it makes sense to fit nonparametric regressions for adult-child combinations of (1,0), (1,1), (1,2), (2,0), (2,1), (2,2), (2,3), (3,1), and (3,2). We do this by calculating the regression function of adult expenditures (in dollars per month) conditional on the logarithm of total annual expenditure using a version of Fan's (1992) local regression smoother. In order to exclude areas of low density at the extremes of rich and poor, we trim the sample to the range of log total outlay shown and then erect a 100-point uniform grid over log outlay. At each point x on this grid, we run a weighted regression with weights given by the quartic kernel:

$$w_i = \tfrac{15}{16}\lambda_i(1 - t_i^2)^2, \quad |t_i| \leq 1, \qquad (12)$$
$$= 0, \quad |t_i| > 1,$$

where t_i is the distance of x_i from x in units of the bandwidth h,

$$t_i = (x - x_i)/h. \qquad (13)$$

The bandwidth h is chosen by visual inspection of alternative estimates in an attempt to preserve genuine structure while eliminating random fluctuations. The multiplier λ_i is an inflation factor, provided in the CEX, that accounts for the fact that the CEX is not a simple random sample. For each observation i,

Table 6.6 Number of Households by Composition, CEX 1990

	Children				
Adults	0	1	2	3	4
1	5,567	554	371	165	48
2	5,452	1,943	2,047	953	258
3	1,195	463	255	126	47
4	406	146	57	52	9
5	94	48	21	13	5

λ_i is equal to the number of households in the U.S. population that the observation represents, so observations that represent more households are given more weight in the regressions.

The results are shown in figure 6.4 for one-adult, two-adult, and three-adult nonpoor families. We restrict our sample to families above the poverty line because families in poverty have very different consumption patterns, as we shall see in subsection 6.3.6. Although the ordering of the curves is not uniformly monotone over the range of total outlay, the ordering is generally in the right direction, and certainly so in the middle of the total expenditure distribution. Holding total expenditure constant, the addition of children reduced expenditure on adult goods. Note also that the slopes of the regression functions are positive, a condition that must be satisfied if we are to calculate the variations in total expenditure that would have effects on adult expenditure equivalent to the addition of children. It is also a condition that is often not satisfied in practice; see, for example, Cramer (1969).

The curves in figure 6.4 are close to one another so that the estimates of child costs that they imply are relatively modest. The mean of log outlay for (1,0) households is 8.06; conditional on this, adult expenditure is predicted to be \$252 a month. To have the same predicted adult expenditure, a (1,1) household would need a log outlay of 8.34, and a (1,2) household 8.52. One child costs 28 percent of a single adult, and two children (together) 46 percent of an adult. For a (2,0) family, an additional child is 32 percent of one adult, two children 72 percent, and three children 76 percent. For three-adult families, the first child costs 54 percent of an adult, two children 90 percent of an adult. While these numbers might make sense for India or a similarly poor country, they seem unreasonably small for the United States.

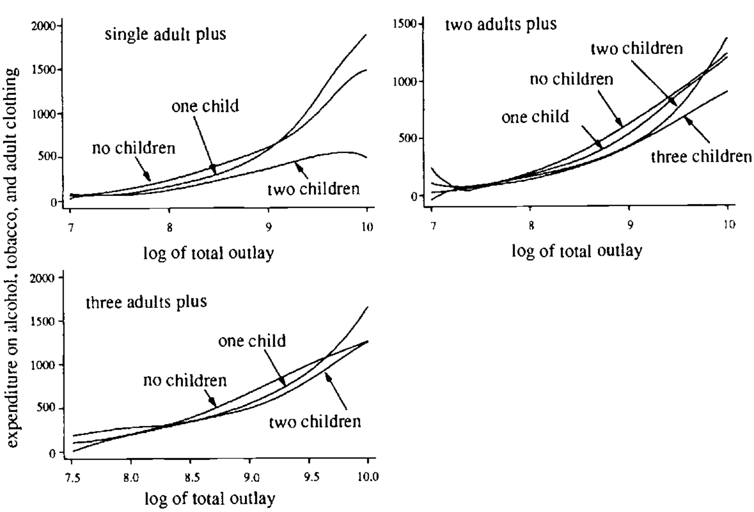

Fig. 6.4 Adult good Engel curves for different family types

Alcohol and tobacco are two items that are notoriously underestimated in household expenditure surveys, including the CEX, so it is wise to reestimate the results using only adult clothing and footwear. But the results are very similar; if anything child costs are estimated to be somewhat lower. This echoes Betson's (1990) findings that the results are not much affected by the choice of the broad or narrow definition of adult expenditures. Another potential source of underestimation is that children consume some of the adult goods. If so, the addition of the child will have a direct positive effect on consumption as well as the negative income effect, so that expenditure on adult goods will be reduced by too little, and child costs underestimated. For alcohol, tobacco, and adult clothing, this is unlikely to be a problem for young children, but it may well be serious for teenagers. However, it is difficult to allow for child age in the nonparametric regressions, so we move to a flexible parametric model.

Our general approach, for the investigation of both child costs and economies of scale, is to allow total expenditure to enter the regression function in a general way, while other variables enter linearly. There is a great deal of information in the data about the shape of the Engel curve, and it is desirable to use it. Because family size and compositional variables are correlated with total expenditure, a failure to model the shape of the Engel curve is likely to compromise the estimates in which we are interested. The model we estimate is written

$$e_A = f(\ln x) + \beta_1 \ln n + \sum_{k=1}^{K} \gamma_k \frac{n_k}{n} + u, \qquad (14)$$

where e_A is expenditure on adult goods, n is household size, x is total outlay, and n_k is the number of people in the kth age-sex category in the household. The right-hand side is motivated by the wish to approximate the "obvious" but nonlinear form

$$e_A = \phi\left(\frac{x}{(A + \sum a_j k_j)^\theta}\right), \qquad (15)$$

which, with an extension to different types of children, is the form with which we began.

We proceed as follows. Equation (14) is first estimated with per capita outlay entered linearly, followed by the Fourier specification

$$f(t) = \sigma_0 + \sigma_1 t + \sigma_2 t^2 + \sum_{i=1}^{M} (\sigma_{ci} \cos(it) + \sigma_{si} \sin(it)), \qquad (16)$$

for which we have generally found that setting M equal to 2 is adequate. The equation is estimated for both narrow (clothing) and broad (clothing, alcohol, and tobacco) definitions of adult goods and is estimated separately for one-adult, two-adult, and three-adult households. Family composition is disaggregated by age and sex, so that we have males and females aged 0–5, 6–11, 12–17, 18–64, and 65 or older. Since the 10 ratios sum to one, we omit that associated with females aged 65 or over, which therefore becomes the reference group.

The regression results are not shown in detail, but a few points are worth noting since they affect the calculations of child costs. The Fourier approximation always improves the fit, but the effects on the other parameter estimates are modest, so we can focus on the more easily interpreted linear form. Clothing and footwear has an expenditure elasticity that is greater than unity, in contrast to the broad aggregate, which is dominated by alcohol and tobacco, the share of which declines with total outlay. The compositional effects work more or less as anticipated. As is required for the model to make any sense, additional prime-age adults increase expenditure on adult goods, although it is adult women who spend most on adult clothing, with little effect for men. Additional children decrease expenditure on adult goods, or at least they do until they are 12 years old, at which point they smoke and drink and, if they are girls, buy adult clothing. We clearly cannot use expenditure on these goods to estimate the costs of older children.

Table 6.7 presents the estimates of child equivalences, calculated separately according to (1) whether we are using all adult expenditure or just clothing, and (2) whether we are using the Fourier approximation or the linear model. The numbers are calculated as follows. In the top panel, we consider the addition of a child to a family containing a single adult female aged 18–64. By equation (14) if $f(\cdot)$ has the linear form $\alpha + \beta \ln x$, adult expenditure in the new larger household will be the same as in the original smaller household if, when the original household has expenditure x, the new one has expenditure x^* given by

$$(17) \qquad \ln \frac{x^*}{x} = \beta^{-1}\left(\beta_1 \ln \frac{1}{2} + \frac{\gamma_{af}}{2} - \frac{\gamma_j}{2}\right),$$

where γ_{af} is the coefficient of the adult female ratio and γ_j is the coefficient of the child ratio in which we are interested. The top panel in the table shows ln $(x^*/x) - 1$, the proportional additional cost of the child. For the two-adult family, the corresponding formula is

$$(18) \qquad \ln \frac{x^*}{x} = \beta^{-1}\left(\beta_1 \ln \frac{2}{3} + \frac{\gamma_{af}}{6} + \frac{\gamma_{am}}{6} + \frac{\gamma_j}{3}\right),$$

where γ_{am} is the adult male coefficient and the numbers presented in the table are scaled—that is, $2(\ln(x^*/x) - 1)$ in order to present the child cost relative to a single adult rather than a couple. The three-adult household is taken to be an adult male, an adult female, and an elderly female, and the costs are calculated in the corresponding way. When the Fourier approximation to the Engel curve is used, it is no longer possible to derive a closed-form expression for

Table 6.7 Estimated Child Costs Relative to Single Adults for Various Family Types by Child Age and Sex (percent)

	All Adult Expenditures		Adult Clothing and Footwear	
Family Type	Linear Form	Flexible Form	Linear Form	Flexible Form
Adult female plus				
Boy 0–5	69	65	90	81
Girl 0–5	70	67	98	91
Boy 6–11	62	69	53	58
Girl 6–11	54	63	72	81
Boy 12–17	68	71	66	66
Girl 12–17	33	45	27	39
Adult couple plus				
Boy 0–5	70	82	53	60
Girl 0–5	72	82	62	68
Boy 6–11	76	83	63	66
Girl 6–11	80	91	76	85
Boy 12–17	36	42	22	24
Girl 12–17	15	19	−5	−6
Three adults plus				
Boy 0–5	60	54	65	58
Girl 0–5	60	57	73	67
Boy 6–11	52	59	28	34
Girl 6–11	43	52	48	57
Boy 12–17	58	61	42	43
Girl 12–17	19	33	−1	13

Notes: The three-adult household is an adult male, an adult female, and a female aged 65 or over. The calculations for the linear form are exact; those for the Fourier flexible form are approximations.

child costs. Instead, we use as an approximation equations (17) and (18) with β replaced by the derivative of the Engel curve with respect to log expenditure evaluated at the mean of log total expenditure. We make no attempt to calculate child costs for children other than the first. While such calculations are straightforward in principle, the differences from the cost of the first child are effectively determined by the choice of functional form (14), rather than by any genuine feature of the data.

In terms of the broad orders of magnitude with which we are concerned, the child cost estimates in table 6.7 are sensitive neither to the choice of functional form for the Engel curve nor to the broad or narrow definition of adult goods. As expected, the estimates for the older children are much too low; "children" between ages 12 and 17 clearly consume these "adult" goods. But the differentiation by age has solved the underestimation problem in the nonparametric results, so that children under age 12 appear to cost around two-thirds to three-quarters of an adult. Comparing the results for single-adult and two-adult families, there is not much evidence of the reduction in child costs that would be expected from economies of scale, although the estimates based on clothing are lower for the smallest children. However, child costs do seem to be lower for the three-adult households. Even so, we could not obtain useful results by following the National Academy procedure and identifying economies of scale by matching the numbers in the table to the ratio of $(A + \alpha K)^{\theta}$ to A^{θ} with $K = 1$. To measure economies of scale, we must turn to more direct methods.

6.3.5 Using Private Goods and Economies of Scale

In this subsection we report our attempts to obtain estimates of economies of scale by looking at the relationship between expenditure on private goods, expenditure per head, and family size. The basic idea is illustrated in figure 6.5, which shows the relationship between expenditure per capita on private goods—or on any single private good or group of private goods—and income (or total expenditure) per head. The lower curve is the private good Engel curve for the smaller household with n_1 members, all of whom are assumed to be adults. Consider an increase to n_2 adults with per capita income held constant, for example, when two single adults get married. Since they have the same total resources as before, it is possible for them to keep their consumption pattern unchanged. However, we would not expect this to happen since the larger family thereby ignores the potential economies of scale associated with public goods. Since the nature of public goods is that they do not have to be duplicated for each household member, the larger household need spend less on them, freeing more resources for private goods, so that the Engel curve

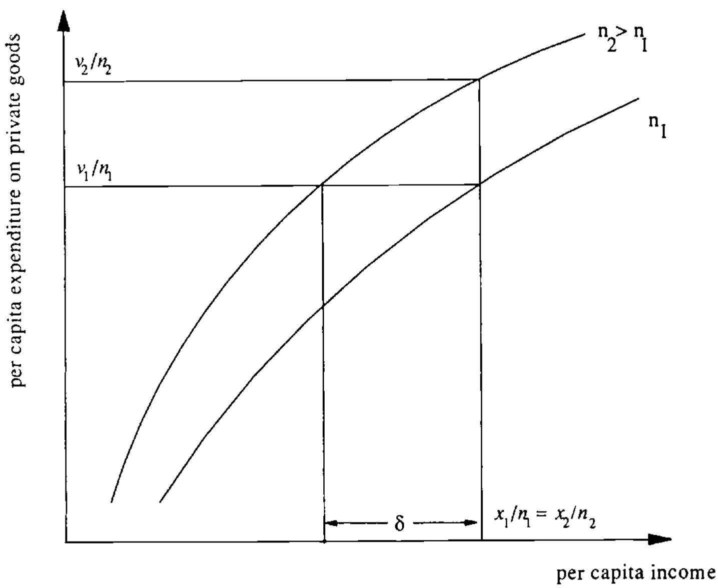

Fig. 6.5 Private good expenditure and the discount to family size

moves upward as shown. The graph shows the amount δ by which per capita income must be reduced so as to restore private expenditures to their previous per capita level. If this amount were deducted from the income of the larger family, that family would still be better off than the smaller family, since a correctly compensated family would presumably take advantage of the relative cheapness of public goods to substitute in their favor and would therefore spend less on private goods. The calculation of δ therefore *understates* the discount from economies of scale. It would be possible to correct for this underestimation by estimating a full demand system that allows for the pricelike effects of economies of scale as in Nelson (1988); the understatement can be minimized by selecting a private good that is a relatively poor substitute for the public goods in the household. For example, people are unlikely to substitute housing and utilities for food just because utility is cheaper through the former the larger is household size.

We therefore began our investigations with food. Although there may be some economies of scale associated with food—bulk buying, reduction of wastage, and saving on items such as cooking oil that are used less than proportionately with scale—the effects seem likely to be modest. There is the further advantage that food is a poor substitute for the most likely public goods. Nevertheless, to our considerable surprise, expenditure on food does not behave as illustrated in figure 6.5. Holding constant per capita outlay, per capita expenditure on food *decreases* with family size, which suggests that the larger the family, the less food each member needs to remain equally well off. Our preliminary investigations for Britain, France, Taiwan, Thailand, and South Africa suggest that this result is true much more broadly than for the United States; the results in Lanjouw and Ravallion (1995) and Ayadi et al. (1995) appear to show the same result in Pakistan and Tunisia, respectively. The detailed evidence would take us too far afield from our main purpose, and a fuller report is contained in Deaton and Paxson (forthcoming). Here we confine ourselves to the main results for the United States, to some possible explanations, and to the search for economies of scale through other potential private goods.

Figure 6.6 shows the locally weighted nonparametric regressions for the food share conditional on the logarithm of per capita expenditure for households with one to four nonelderly adults and zero children. Although the argument in figure 6.5 was presented in terms of per capita food expenditure increasing with family size, the same argument applies to the food share in the budget, which is the ratio of per capita food expenditure to per capita total expenditure, and the latter is being held constant as family size is changed. Although the relationship between the four curves changes with the level of per capita outlay, the one-adult curve is highest through most of the range, and the four-adult curve is the lowest. The curves for two and three adults are close to one another and cross more than once. At the same level of per capita outlay, one-adult households spend more per capita on food than do two- or three-adult households, who in turn spend more than four-adult households. While

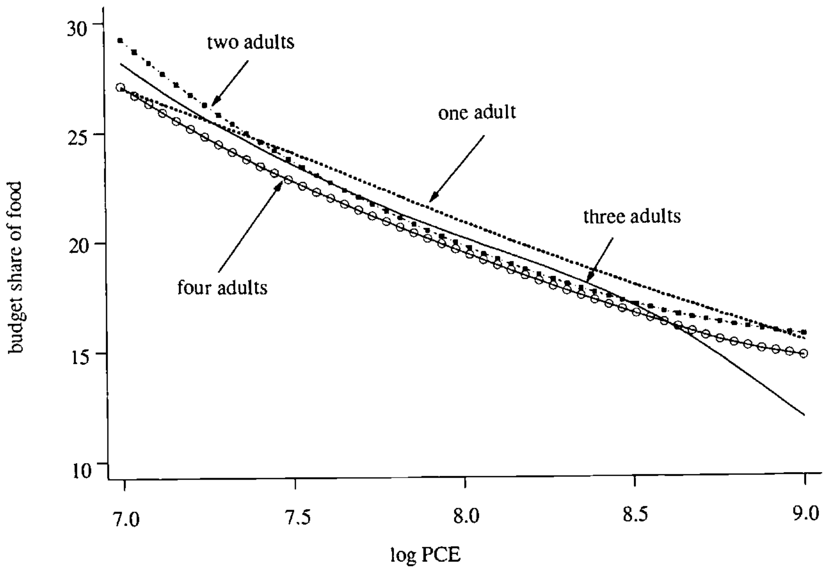

Fig. 6.6 Food Engel curves for nonelderly adult households

it is possible to think of reasons why one-adult families are special—they eat out a great deal, or they buy expensive precooked food, or they waste a great deal of food—it is hard to reconcile the results for the other family types with the presumption that food is a private good and that economies of scale operate through nonfood public goods.

In order to investigate whether the same phenomenon holds for more complex family types and survives the conditioning on a wider range of variables, we once again have to move to a more parametric treatment. In parallel with our earlier discussion of child costs and adult goods, we adopt a flexible functional form for the food budget share of the form

$$w_f = f\left(\ln\frac{x}{n}\right) + \beta_1 \ln n + \sum_{k=1}^{K} \gamma_k \frac{n_k}{n} + \psi\pi + u, \tag{19}$$

where w_f is the food share and π is a dummy indicating whether the family falls below the official poverty line, a dummy that is provided on the CEX tape. Although not shown in equation (19), we also interact household size and the demographic compositional variables with the poverty dummy. This is done, not because poverty itself changes demand patterns—and if it did, the effect would be picked up by a sufficiently flexible function $f(\cdot)$—but because families in poverty, depending on their demographic composition, are eligible for food stamps, housing benefits, and other programs that are likely to change their consumption patterns.

As before, we treat the function $f(\cdot)$ in more than one way: (a) using the linear specification $\alpha + \beta_0 \ln x/n$, (b) using the Fourier flexible form (16), and

(c) following a nonparametric technique of Estes and Honoré (1995). The last works by first sorting the data according to the values of x/n and then fitting equation (19) by ordinary least squares to the first-difference of all the right-hand–side variables excluding x/n. As the sample size increases, and provided that the function $f(\cdot)$ is continuous, the distance between any two successive values decreases at rate n^{-1} so that, in the limit, the influence of x/n is purged from the regression as if it were a (local) fixed effect. The advantage of this procedure—apart from its elegance and simplicity—is that we do not have to specify any functional form. The disadvantage is that we do not obtain an estimate of the slope of the Engel curve, something that is of no importance when our aim is to find the sign of β_1, the coefficient on family size, but that would become necessary if we ever got as far as calculating the compensation as in figure 6.5.

The three sets of results are presented in table 6.8. In the event, getting the shape of the Engel curve right has little effect on the other coefficients, and in

Table 6.8 Food Share Regressions Using Three Alternative Techniques

Variable	Linear Form	Fourier Flexible Form	Estes and Honoré Nonparametric Form
ln x/n	−7.25 (39.3)	[a] [a]	[b] [b]
ln n	−0.71 (2.9)	−0.67 (2.7)	−0.95 (4.6)
rm05	−7.28 (7.3)	−7.48 (7.6)	−7.58 (7.8)
rm611	−1.08 (1.0)	−1.27 (1.2)	−0.85 (0.8)
rm1217	2.86 (2.5)	2.78 (2.4)	2.21 (2.1)
rm1864	2.36 (5.2)	2.10 (4.8)	2.13 (5.7)
rm65	3.34 (4.6)	3.29 (4.5)	3.20 (5.5)
rf05	−9.22 (9.1)	−9.41 (9.2)	−9.54 (9.5)
rf611	−2.73 (2.7)	−2.70 (2.8)	−1.32 (1.2)
rf1217	1.14 (1.0)	0.90 (0.8)	0.24 (0.2)
rf1864	−1.22 (2.9)	−1.22 (2.9)	−0.99 (2.7)
π*rm05	11.27 (3.9)	9.62 (3.3)	10.14 (4.6)
π*rm611	6.16 (2.1)	4.77 (1.6)	5.14 (2.3)
π*rm1217	2.87 (1.0)	1.14 (0.4)	−0.80 (0.3)
π*rm1864	−1.73 (1.4)	−1.62 (1.3)	−1.08 (1.3)
π*rm65	−3.90 (2.3)	−3.93 (2.3)	−3.63 (2.7)
π*rf05	6.80 (2.5)	5.27 (2.0)	6.76 (3.1)
π*rf611	6.97 (1.8)	5.31 (1.4)	4.66 (2.0)
π*rf1217	−1.38 (0.5)	−2.15 (0.7)	1.64 (0.7)
π*rf1864	−1.42 (1.4)	−1.31 (1.3)	−1.53 (1.9)
π	1.01 (2.2)	1.61 (2.0)	1.03 (1.7)

Notes: All coefficients multiplied by 100. Numbers in parentheses are absolute values of t-statistics. For the first two columns, the standard errors on which the t-statistics are based have been corrected for random household effects. (Recall that the CEX has a panel element, and a single household may contribute up to four observations.) The Estes and Honoré standard errors have not been corrected for the first-differencing or for household random effects.

[a]Estimated by Fourier flexible form and coefficients not shown.

[b]Estimated as a local fixed effect and eliminated from the regression.

all three specifications, the budget share of food is reduced by between a full point and two-thirds of a point in response to a unit increase in the logarithm of family size. This is not a large effect, but the negative sign is statistically significant, and once again we have the reverse of what would be predicted by our description of economies of scale. The other coefficients in table 6.8 are of some independent interest. Controlling for per capita outlay, children, whether boys or girls, reduce per capita food consumption, and men spend more on food than do women. The poverty dummy has strong positive interactions with the presence of children under the age of 12, presumably through the operation of food stamps and AFDC.

We have no good explanation for the failure of food to behave like a private good. Perhaps there are substantial economies of scale associated with food, or perhaps increases in family size cause more substitution than seems plausible from food toward housing and other public goods. Whatever the story, the conclusion that a negative discount understates the true discount to family size, although hard to challenge, is not very useful. We therefore turn to other goods as potential private goods, although we confess that our confidence in the general methodology is somewhat shaken by the failure of the most promising good to yield useful results.

As for food, we estimate equation (19) for a range of goods from the CEX. We calculate the change in the logarithm of per capita total outlay that keeps per capita expenditure constant in the face of a change in household size with household composition held constant. This is done through the per capita expenditure-constant elasticity, from equation (19):

$$\frac{d \ln x/n}{d \ln n} = \frac{\beta_1}{w + f'(\ln x/n)}, \tag{20}$$

which is evaluated at the sample means.

The results, in table 6.9, are given for a list of commodity groups that sums to total outlay, and they are almost as bewildering as those for food. Of course,

Table 6.9 **Family Size Percentage Discounts for Various Items of Expenditure**

Commodity Group	Discount Using Linear Form	Discount Using Flexible Form
Food	−4.5	−4.4
Alcohol	−82.4	−69.0
Tobacco	−60.6	−88.6
Clothing	5.6	5.9
Transportation	20.3	19.9
Housing and utilities	−10.6	−10.3
Entertainment	9.8	9.8
Personal care	9.1	8.9
Medical expenditures	31.0	36.6
Reading material	−8.6	−7.7
Educational expenditures	17.2	16.7

goods like transportation and housing and utilities are certainly not private goods, and it is no surprise to find that expenditure per capita on housing falls with family size. But commodities like alcohol and tobacco, which might seem to be candidates for private goods, have less spent on them as family size increases. Perhaps these are examples of *negative* public goods, which generate negative externalities for other members of the family, and which therefore become more expensive as family size increases. Transportation and medical expenditures—the latter strongly associated with the elderly in the estimates—are items for which expenditure increases with family size. The former would appear to have a public element, in which case its expenditure would be expected to decrease, while it is hard to see why family size should increase medical expenditures when the age composition of the household is being controlled for. There is perhaps some consolation in the estimates for clothing, personal care, and educational expenditures. If we use these, and ignore the anomalous results elsewhere, we obtain modest economies of scale, with the elasticity of needs to size between 85 and 95 percent. But it is hard to escape the conclusion that expenditure patterns respond to family size in ways that are a good deal more complex than the simple story of public and private goods that we have considered in this paper. Constructing better models of this process remains a challenge for the future.

References

Ayadi, Mohamed, Rafiq Baccouche, Mohamed Goaid, and Mohamed Matoussi. 1995. Spatial variations of prices and household demand analysis in Tunisia. Tunis: University of Tunis III, Faculté des Sciences Economiques et de Gestion de Tunis, January. Processed.

Barten, Anton P. 1964. Family composition, prices, and expenditure patterns. In *Economic analysis for national economic planning,* ed. Peter Hart, Gordon Mills, and John Whitaker. London: Butterworth.

Betson, David M. 1990. Alternative estimates of the cost of children from the 1980–86 Consumer Expenditure Survey. Special Report no. 51. Madison, Wisc.: Institute for Research on Poverty.

Cramer, Jan S. 1969. *Empirical econometrics.* Amsterdam: North-Holland.

Deaton, Angus S., and John Muellbauer. 1980. *Economics and consumer behavior.* New York: Cambridge University Press.

———. 1986. On measuring child costs. *Journal of Political Economy* 94:720–44.

Deaton, Angus S., and Christina H. Paxson. Forthcoming. Economies of scale, household size, and the demand for food. *Journal of Political Economy.*

Estes, Eugena M., and Bo E. Honoré. 1995. Partial regression using one nearest neighbor. Princeton, N.J.: Princeton University. Manuscript in preparation.

Fan, Jianqing. 1992. Design-adaptive nonparametric regression. *Journal of the American Statistical Association* 87:998–1004.

Gieseman, Raymond. 1987. The Consumer Expenditure Survey: Quality control by comparative analysis. *Monthly Labor Review* 110 (3): 8–14.

Haddad, Lawrence, and S. Ravi Kanbur. 1990. How serious is the neglect of intra household inequality? *Economic Journal* 100:866–81.
Lanjouw, Peter, and Martin Ravallion. 1995. Poverty and household size. *Economic Journal* 105:1415–34.
Lazear, Edward P., and Robert T. Michael. 1988. *The allocation of income within the household.* Chicago: University of Chicago Press.
National Research Council. 1995. *Measuring poverty: A new approach.* Washington, D.C.: National Academy Press.
Nelson, Julie A. 1988. Household economies of scale in consumption: Theory and evidence. *Econometrica* 56:1301–14.
Nicholson, J. Leonard. 1976. Appraisal of different methods of estimating equivalence scales and their results. *Review of Income and Wealth* 22:1–18.
Organization for Economic Cooperation and Development. 1982. *The OECD list of social indicators.* Paris: Organization for Economic Cooperation and Development.
Pollak, Robert A., and Terence J. Wales. 1979. Welfare comparisons and equivalence scales. *American Economic Review Papers and Proceedings* 69 (2): 216–21.
Rothbarth, Erwin. 1943. Note on a method of determining equivalent income for families of different composition. In *War time pattern of saving and spending,* ed. Charles Madge. Cambridge: Cambridge University Press.
Ruggles, Patricia. 1990. *Drawing the line.* Washington, D.C.: Urban Institute Press.
Slesnick, Daniel T. 1993. Gaining ground: Poverty in the postwar United States. *Journal of Political Economy* 101:1–38.

Comment B. Douglas Bernheim

This provocative and ambitious paper wrestles with the difficult problem of measuring poverty consistently across diverse groups of individuals. The paper is divided into two parts. The first part describes current practices regarding the measurement of poverty in the United States, and it examines the sensitivity of official poverty measures to alternative assumptions about the size of appropriate adjustments for differences in household composition. The second part outlines and implements a potentially superior methodology for establishing the size of these adjustments.

Section 6.2 focuses on three specific household composition adjustments that are imbedded in existing poverty thresholds. These adjustments are intended to account for family size, the presence of children, and the presence of elderly households. The paper effectively makes the case that the official methods of making these adjustments are arbitrary, and moreover that the measurement of relative rates of poverty among different population subgroups is quite sensitive to plausible changes in the magnitudes of the adjustments.

The analysis in this part of the paper is, without question, very useful. Still, it should hardly come as a surprise that there is inherent arbitrariness attached to the use of a single summary statistic for income distribution, such as a pov-

B. Douglas Bernheim is the Lewis and Virginia Eaton Professor of Economics at Stanford University and a research associate of the National Bureau of Economic Research.

erty rate (the authors' idea of looking for stochastic dominance relations between income distributions for different population subgroups therefore strikes me as particularly good). One is left at the end of this paper wondering why economists should be engaged in the possibly futile exercise of trying to define meaningful poverty thresholds in the first place. I suspect that we have focused on poverty rates primarily to satisfy the demands of politicians and the press, who generally seem to limit their attention to single numbers. To the extent that economists wish to affect the policy process, it may be necessary to cater to the demand for oversimplification; thus, one justifies the exercise in this paper by arguing that, if politicians insist on using a single number, we should make sure that it is the best number possible. But then the ultimate point of this paper seems to be that the construction of a truly "good" poverty number is impossible.

That said, it should be emphasized that there may well be some conceptually legitimate applications of poverty statistics. Although it may be very difficult, and possibly even meaningless, to compare rates of poverty across population subgroups, it may be justifiable to make comparisons across time for the same subgroup. That is, by comparing poverty rates for the elderly in, say, 1990 and 1965 using some constant-real-dollar poverty threshold, one may be able to learn something important about changes in the lower tail of the income distribution, even if the absolute level of the threshold is inherently arbitrary.

In section 6.3, the authors attempt to develop and implement a methodology for improving two of the household composition adjustments: family size and presence of children. The adjustment for the presence of elderly individuals is conspicuously absent from this list, which is unfortunate; one would expect the measurement of poverty among the elderly (the focus of the paper) to be most sensitive to the elderly adjustment.

This portion of the paper does an extremely thorough job of criticizing itself. The authors have already acknowledged, in some form, most (if not all) of the problems with their approach. My critique differs from the authors' own critique more in emphasis than in substance.

As the authors note, there are two problems associated with moving from family resources to individual measures of welfare. The first problem concerns the intrahousehold allocation of resources. The second problem concerns the effects of household size and age structure on the standard of living. The authors write, somewhat euphemistically, that they "have little to say" about the first problem. Yet is difficult for me to see how one can come up with a coherent resolution to the second problem without resolving the first one.

Consider in particular the authors' method of determining the appropriate poverty line adjustment for children. This involves partitioning consumption into "adult goods" and other goods, on the premise that children do not consume adult goods. But this premise inherently raises the issue of distribution. The underlying hypothesis is simply inconsistent with the view that resources are divided equally within the household. The authors have tried to finesse this problem linguistically, by assuming that resources are distributed "equitably," rather than "equally," among household members. But the meaning of "equitably" in this context is unclear, since equity necessarily involves the comparison of unequal consumption bundles across consumers with presumptively different preferences.

Even if one could articulate a notion of equity that would correspond to equality of standard of living across heterogeneous family members, it is difficult to see why this notion would be held forth as a positive theory of household behavior, rather than as a normative standard. If the authors proceed on the basis of the assumption that intrahousehold allocation is governed by (as yet unspecified) egalitarian principles, while in fact resources are distributed unequally within households, then it is highly unlikely that their methodology will yield appropriate adjustment factors.

Consider the following example. Suppose that all household decisions are made by a completely selfish, dictatorial head. The head consumes all of the household's resources and instructs other family members to beg on the street for sustenance. In that case, it would appear to me that Deaton and Paxson would estimate $\alpha = 0$. However, it seems incorrect to conclude from this finding that policymakers should regard children as costless. Alternatively, imagine that the dictatorial head is not completely selfish but cares less about each member that is added to the family. In that case, Deaton and Paxson would significantly overestimate the true extent of economies to scale.

These examples make it clear that, for most positive theories of household behavior, it is not going to be possible to define a single poverty threshold for the entire household. It makes more sense to define a separate poverty level for each household member, where this threshold is chosen to represent the level of household income at which that member would achieve some measure of subsistence. Thus, in the first example above, there would be some finite poverty threshold for the head, and the formula for this threshold would indeed involve $\alpha = 0$, but the poverty thresholds for other household members would be infinite (since no amount of household income would bring them to subsistence). In the second example, the poverty threshold for the head would rise at a declining rate with the number of household members. However, for any given household composition, younger children would have higher thresholds than older children, and the head would have the lowest threshold of all.

These examples also make the general point that the authors' methodology is likely to produce little more than a measure of the preferences of the household decision maker. To drive this point home, consider the following whimsical suggestion: why not adjust the poverty line for the presence of household pets? This is potentially important in the context of aging, since the elderly are probably significantly more likely to be in pet-less households, and therefore less likely to have to share household resources with pets. How would one calculate the appropriate adjustment? Following Deaton and Paxson's approach, one would divide expenditures into "human goods" and all other goods (including nonhuman goods). This might, of course, prove difficult in practice, since household pets do consume many human goods ("people food" being the leading example). Abstracting from this problem, one could then determine the levels of household income that would equate spending on human goods between households with pets and pet-less households. One would then use this to set an adjustment factor for pets, analogous to α. Thus, the poverty threshold would be higher for households with pets.

I would, however, question the appropriateness of the pet adjustment. First, the ownership of pets is voluntary. The human members of the household would not have acquired a pet unless they thought that they would be better off as a result of owning the pet. Consequently, there is no reason to set a higher poverty threshold for the humans merely because they choose to own a pet. Second, with respect to the welfare of the pet (rather than the welfare of the humans), the division of expenditures between human and nonhuman goods is likely to tell us more about the humans' willingness to spend money on the pet than about the costs of maintaining a pet at subsistence.

These observations are not completely whimsical. Children are, at least to some extent, voluntary choices. Thus, if one were properly defining poverty thresholds for individuals, rather than for households, it is not at all obvious that it would be appropriate to increase the poverty threshold for parents based on number of children. With respect to poverty thresholds for children, the Deaton-Paxson method might measure some aspect of parents' willingness to spend money on children, but it probably does not tell us much about the welfare of children.

Generally, it seems to me that the Deaton-Paxson estimates teach us about preferences, rather than about subsistence standards. Consider again the strategy of differentiating between adult goods and other goods. Even if children do not consume adult goods, their presence may affect (or be correlated with) the preferences of adult household members. For example, adults may prefer to engage in fewer activities involving the consumption of alcohol after having children; alternatively, children may drive their parents to drink. People who have children may simply have different tastes than those without children. If different adult goods have different income expansion paths, then the methodology may simply identify, in a convoluted way, the income elasticities of the adult goods that have been selected for the analysis.

I am also skeptical about the kinds of comparisons that the authors make in their attempt to identify household economies of scale, and I am less surprised by their "puzzling" results concerning food. I agree that food is, almost certainly, a private good. However, the preparation and storage of food are public goods. Economies of scale in preparation and storage may induce smaller families to eat out more often, raising the cost per meal and giving rise to apparent economies of scale in food consumption. Alternatively, the taste for eating at home may change when family size grows, as the social aspects of family meals become better substitutes for the social aspects of eating out.

As I have emphasized, the authors are acutely aware of the limitations of their methodology. The great challenge here is to come up with a better approach. Although I have failed to rise to this challenge, I have two suggestions.

First, it might be possible to design a less arbitrary poverty line if the purpose of the poverty line were more clearly defined. One could, for example, decide to define the poverty line in terms of the adequacy of nutrition. In that case, meaningful poverty thresholds could be constructed as follows. First, set standards of adequacy for nutritional outcomes. Second, collect data on nutritional outcomes, income, and household characteristics. Third, relate nutritional outcomes to income and household characteristics separately for each household position (husband, wife, first male child, first female child, elderly parent, etc.). This would allow one to determine, for each household position, the level of household income necessary to assure nutritional adequacy with some prespecified level of probability.

Second, instead of comparing consumption across different households in cross sections, one might compare consumption for the same household across time, as household composition changes. The life cycle hypothesis, in effect, argues that households should smooth their standards of living through time. Thus, one could examine changes in consumption associated with household additions to determine the amount of resources required to compensate for an addition. One could also examine the manner in which consumption responds to departures from households. This has some advantages, since certain departures (e.g., of grown children) are usually anticipated. Life insurance purchases, which anticipate departures, can also be used to infer household economies of scale.

While this approach has some attractive conceptual features, it also presents some new problems. First, one must subscribe rather unreservedly to the hypothesis of rational intertemporal optimization. Second, the applicability of the method to lower income households may be limited, to the extent that these households are liquidity constrained. The possibility of binding liquidity constraints suggests a somewhat more ad hoc approach: measure the level of income at which the typical household begins to save, and examine how this varies with household composition.