

Economies of Scale, Household Size, and the Demand for Food

Angus Deaton and Christina Paxson

Princeton University

Household scale economies are plausibly attributed to shared household public goods that make larger households better off at the same level of per capita resources. Larger households should therefore have higher per capita consumption of private goods, such as food, provided that they do not substitute too much toward the effectively cheaper public goods. The evidence shows exactly the opposite. Data from rich and poor countries indicate that, at constant per capita total expenditure, the per capita demand for food decreases with household size and that it does so most in the poorest countries, where substitution should be the least.

I. Introduction

Together with the measurement of child costs, the measurement of economies of scale within the household is an essential prerequisite for inferring individual living standards from household data. The

We are grateful to Jean Drèze, who posed some of the original questions that led us to this paper, and to the John D. and Catherine T. MacArthur Foundation for financial support. We have received useful comments from Orley Ashenfelter, Abhijit Banerjee, Martin Browning, David Card, Anne Case, Pierre-André Chiappori, Jean Drèze, Richard Eckhaus, Paul Gertler, Maitreesh Ghatak, Michael Kremer, Guy Laroque, Thierry Magnac, Jacques Mairesse, Jonathan Morduch, John Muellbauer, Julie Nelson, Ben Polak, Jean-Marc Robin, Amartya Sen, Peter Timmer, and two referees. We should also like to thank Jean-Michel Hourriez for help in accessing the French data. Material from the British Family Expenditure Survey is Crown Copyright and has been made available by the Central Statistical Office through the Economic and Social Research Council Data Archive and has been used by permission. Neither the CSO nor the ESRC Data Archive bears any responsibility for the analysis or the interpretation of the data reported here.

[*Journal of Political Economy*, 1998, vol. 106, no. 5]
© 1998 by The University of Chicago. All rights reserved. 0022-3808/98/0605-0001\$02.50

most plausible source of economies of scale is the presence of household public goods that can be shared and serve their function without needing to be replicated in proportion to the number of household members. Imagine bringing together two previously separate adults while retaining their original incomes, so that per capita income in the partnership is the same as it was for the average of the two separate units. In the presence of public goods, the couple are now better off. They can do everything that they did before, but new options are available. In particular, the resources released by sharing allow more to be spent on everything, public and private goods alike. There will also be substitution effects toward the shared goods, which are now effectively cheaper for members of the larger household; money spent on public goods is twice blessed. For the private goods that are *not* shared, income and substitution effects will operate in different directions. However, if we can identify a private good that is not easily substitutable, with low own- and cross-price elasticities, the income effect will dominate, and per capita consumption of the good will increase. Indeed the size of such an increase could potentially serve as a measure of economies of scale, parallel to Rothbarth's (1943) suggestion that the decline in the expenditure on adult goods consequent on an additional child be used to measure its cost.

The most obvious candidate for a private good that is not easily substitutable is food, particularly among poor households whose consumption of food is close to subsistence. With per capita resources held constant, food consumption per head should rise with household size. Failure of this prediction—if it occurs—is most likely when incomes are high and food needs are long satisfied. When some of food consumption is a luxury—meals in expensive restaurants—food can be substituted in favor of goods that can be shared among the members of a larger household. But for the poorest households, the increase in welfare from the economies of scale should turn into an increase in per capita consumption of food. Indeed, there are good general grounds for supposing that per capita food consumption is an excellent measure of individual welfare for households in developing countries so that, on these grounds alone, we should expect the per capita demand for food to increase with proportional increases in household size and total household resources.

In this paper, we demonstrate that the *empirical evidence is exactly the opposite of the theoretical predictions*. With total household expenditure per capita (PCE) held constant, expenditure per head on food *falls* with the number of heads. The result appears to be quite general: we find it not only in the United States, but also in Britain and

France and, more surprisingly, in Taiwan, Thailand, and Pakistan and among African households in South Africa. The *size* of the effect is also contrary to the theoretical analysis. In the United States, Britain, and France, food consumption falls by only a small amount as the scale of the household increases. The large effects come in Thailand, Pakistan, and South Africa. For these poorest countries, we estimate that, with PCE held constant, a unit increase in the logarithm of household size will decrease the budget share of food by five percentage points and decrease per capita expenditure on food by more than 10 percent. Such a result is entirely paradoxical. Larger households are better off at the same level of per capita resources since they have the option of decomposing themselves into smaller units. Yet by the most obvious indicator of their welfare, per capita food consumption, they are worse off. We have two apparently compelling arguments, one indicating that people are better off and one indicating that they are worse off.

The remainder of the paper is laid out as follows. Section II is concerned with theoretical issues concerning economies of scale. We use Barten's (1964) analogy between prices and household structure to sharpen the basic argument and to derive conditions on price and income elasticities under which food demand per person increases with household size in the presence of public goods. Section III contains empirical evidence from the United States, Great Britain, France, South Africa, Taiwan, Thailand, and Pakistan documenting the failure of the empirical prediction. We employ a range of parametric and nonparametric techniques in an attempt to eliminate various possible econometric explanations. The section concludes with a discussion of a number of possible explanations for our findings. None is particularly satisfactory or capable of explaining all the evidence; our results remain a puzzle.

II. Food and Economies of Scale: Theory

We begin by considering the demand for food and for public goods using a very simple version of the model first developed by Barten (1964). Suppose that there are two goods, one (food) that is entirely private and one (housing) that is shared and is entirely public. A household with n members allocates its total outlay or expenditure x across these two goods. (The distinction between total expenditure and income will be dealt with in the empirical Sec. III below.) For now, assume that there are only adults in the household. If q_f and q_h are the levels of *household* consumption of food and housing, respectively, and everything is shared equally, the household utility

function is $n\nu(q_f/n, q_h)$ because food, unlike housing, must be divided. The household budget constraint in per capita terms is

$$p_f \left(\frac{q_f}{n} \right) + \left(\frac{p_h}{n} \right) q_h = \frac{x}{n}. \quad (1)$$

When two one-person households join to form one two-person household, neither resources per head, x/n , nor the price of the private good, p_f , changes. But the price of the public good is halved, which induces a positive income effect and a negative substitution effect on the demand for the private good. If the private good is a necessity such as food, the substitution effect will be small, and its per capita consumption will rise.

This basic insight carries through to more realistic models. Suppose that we distinguish food and nonfood and that both are subject to some scale economies. We write utility in the form

$$u = n\nu \left(\frac{q_f}{\phi_f(n)}, \frac{q_h}{\phi_h(n)} \right). \quad (2)$$

where h is broadly interpreted to be nonfood, and the quantities are *household* purchases. Maximization of (2) subject to the budget constraint (see, e.g., Deaton and Muellbauer 1980, p. 200) gives a food demand function

$$\frac{q_f}{n} = \frac{\phi_f(n)}{n} g_f \left(\frac{x}{n}, \frac{p_f \phi_f(n)}{n}, \frac{p_h \phi_h(n)}{n} \right), \quad (3)$$

where $g_f(x, p_f, p_h)$ is the food demand function for a single-person household. Taking logs of (3) and differentiating with respect to $\ln n$, we see that per capita food consumption will increase with n at constant x/n if and only if

$$\sigma_h(\epsilon_{f_x} + \epsilon_{ff}) - \sigma_f(1 + \epsilon_{ff}) > 0, \quad (4)$$

where σ_i , $i = f, h$, is the commodity-specific economy of scale measure

$$\sigma_i = 1 - \frac{\partial \ln \phi_i(n)}{\partial \ln n}, \quad (5)$$

ϵ_{ff} and ϵ_{f_x} are the own-price and income elasticities of food, and we have used the zero-degree homogeneity of g_f . The validity of (4) is an empirical matter, and the magnitudes will surely vary with the level of development; the general supposition is that, in poor countries, the food share is high, the income elasticity of food is high,

and the price elasticity of food is low (in absolute value). When food is a pure private good and $\sigma_f = 0$, (4) is equivalent to the requirement that the absolute value of the price elasticity of food be less than its income elasticity, which should be true among poor consumers, if not the rich.

It is useful to rewrite (4) in terms of the compensated own-price elasticity $\tilde{\epsilon}_{ff}$ and to assume that the ratio σ_f/σ_h is small, in which case the condition for the per capita demand for food to increase with n is (approximately) that

$$\tilde{\epsilon}_{ff} > -\epsilon_{fx}(1 - w_f) + \frac{\sigma_f}{\sigma_h}(1 - \epsilon_{fx}). \quad (6)$$

If food is a necessity, has limited substitutes, and is a private good, (6) will hold. Failures, if they exist, are more likely among wealthy households, where the income elasticity is low. In the data that we use in the next section, the total expenditure elasticity ϵ_{fx} ranges, for example, from .75 in rural Thailand to .50 in Great Britain, with corresponding average food shares of .40 and .20, respectively. If σ_f is zero, (6) will be satisfied in rural Thailand if the compensated own-price elasticity of food is greater than $-.45$ and in Britain if it is greater than $-.40$; these are the minimal and maximal values of $-\epsilon_{fx}(1 - w_f)$ in the seven countries that we examine.

The Barten model can be criticized, not for its insight that demographics exert price-like effects, but for its supposition that they operate *only* in such a way, and extensions of the model can change the results. One such extension concerns fixed consumption costs, of which the most important is likely to be housing. If the fixed costs are already being met or if an additional person adds only modestly to fixed costs, additional household members will reduce fixed costs per head. This is a pure income effect and should increase food demand per capita, as in the analysis above. However, the opposite may occur if the increase in size makes it worthwhile to incur the fixed costs, in which case the income effect works in reverse. This may be relevant for the move from one-person to two-person households in countries such as the United States, where the setup costs of housing are high and young couples may skimp on all nonhousing expenditures in order to get a foothold in the housing market. It is likely to be less important in poor countries.

We can extend the analysis to accommodate children, who typically consume less than adults. For example, "effective" household size can be represented as $A + \alpha K$, where A and K are the numbers of adults and children, and α is a number between zero and one that indicates the fraction of an adult each child represents. The

deflation factors ϕ_f and ϕ_h in (2) can then be specified as functions of effective household size rather than n . (A more sophisticated approach could allow different values of α for different goods.) In such a model, increases in the fraction of children, with both x/n and n held fixed, are likely to *decrease* food demand per person. For example, if food is a pure private good, so that $\phi_f = A + \alpha K$, whereas for housing, $\phi_h = (A + \alpha K)^{(1 - \sigma_h)}$, per capita food consumption will decrease with the number of children provided that

$$(1 - \epsilon_{fx}) + \sigma_h(\epsilon_{fx} + \epsilon_{ff}) > 0. \quad (7)$$

The first term will be positive for a necessity such as food; substituting children for adults decreases effective household size, improves welfare per head, and shifts the budget away from necessities. The second term will also be positive as long as food has limited substitutes; children are less than fully equivalent to adults, so that the substitution of children for adults reduces the size of the scale economies, which also reduces the demand for food. As before, the presence of fixed costs can alter these predictions. For example, suppose that there are fixed costs for additional people and such costs are smaller for children than for adults. In this case, substituting children for adults will yield income effects that could have a positive effect on per capita food demand.

In summary, although it is possible to find anomalous cases, the existence of economies of scale has two predictions for the per capita demand for food: (a) that at constant PCE, demand will rise with household size, and (b) that at constant PCE and constant household size, demand will be smaller the larger the proportion of children.

Before we proceed to the empirical work, it is useful to contrast our treatment of economies of scale with that of the widely applied "Engel method." In 1895, Ernst Engel asserted that the food share correctly indicated welfare across households of different sizes and compositions, with lower food shares indicating higher welfare. This assertion, sometimes referred to as Engel's second law, has been used as the basis for measuring scale economies (for a recent example, see Lanjouw and Ravallion [1995]). Given that the food share enables us to identify households that are equally well off, scale economies can be calculated by comparing the total outlays of households with the same food shares but different numbers of members. As we shall see below, the empirical fact is that the food share is negatively related to both household size and PCE, so that Engel's method delivers positive estimates of economies of scale; as household size rises the food share falls, so that a decrease in PCE is required to restore the food share to its higher initial level.

Although Engel's method is internally consistent, it directly contradicts the model of scale economies and public goods presented above, which showed that economies of scale should cause the food share to *increase* with household size, not *decrease*. Because we are so familiar with the fact that the food share falls with increases in welfare as *resources* increase—Engel's (first) law—it at first seems reasonable that the food share should also decline when welfare increases as household size and resources expand in proportion. But because the food share is food expenditure per capita divided by total expenditure per capita, a decline in the food share with constant PCE can occur only if there is a decline in food expenditure per capita, which is not at all what we expect to happen when welfare increases. Reductions in food per person are inconsistent with the welfare gains that (PCE-constant) increases in household size should generate in the presence of economies of scale. In consequence, the estimates of economies of scale that are derived by Engel's method have no theoretical underpinning and are identified by an assertion that makes no sense.

III. Food and Economies of Scale: Evidence

A. Data

In this section we use data from a number of household surveys to examine the relationship between expenditures on food and household size. We work with data from seven countries: the United States, Great Britain, France, Taiwan, Thailand, Pakistan, and South Africa. While the specific choice of countries is conditioned by data availability, we want to look at a wide range of developing and developed countries. The measurement of economies of scale is important in poor and rich countries alike, and we wish to document the generality—or lack of it—of the relationship between food and household size. Even if the data were to conform to the theory of the previous section, there is no reason to expect the results to be identical across countries with different incomes and prices, preferences, and relative costs of adults and children. As we shall see, not only do the results contradict the general theoretical prediction that per capita food consumption increases with household size, but the differences across countries are exactly contrary to expectations. The theory says that it is the poorest countries that should have the strongest positive relationship between per capita food consumption and household size. However, we shall find that exactly the opposite is the case: in the poorest countries there is a large *negative* association between food consumption per person and household size.

In each case, we use a single cross-sectional household survey: for the United States, the 1990 Consumer Expenditure Survey; for Great Britain, the 1992 Family Expenditure Survey; for France, the 1989 Family Budget Survey; for Taiwan, the 1990 Survey of Personal Income Distribution; for Thailand, the 1992 Socioeconomic Survey; for Pakistan, the 1991 Living Standards Survey; and for South Africa, the 1993 Living Standards Survey. For all countries except Thailand and South Africa, we use all households in our analysis. For Thailand, the survey is large enough and differences in living standards across urban and rural areas are substantial enough for us to analyze rural and urban households separately and omit households in the peri-urban "sanitary districts." For South Africa, we use data only on African households and exclude households that are defined (in the language of apartheid) to be "white," "colored," or "Indian." The African households represent the largest and the poorest group within South Africa. All the surveys contain data on total household expenditures on all consumption items; total expenditure on food, including meals taken away from home; and the ages and genders of household members, usually together with a wide range of other sociodemographic variables. For all surveys except the U.S. Consumer Expenditure Survey (CEX), each household is surveyed only once. The CEX is conducted quarterly, and each household is surveyed in up to four consecutive quarters, so that there are multiple observations on some households in the data. The Appendix contains more information on these surveys and on the construction of variables used for the analysis.

We use two approaches to analyze the relationship between per capita food expenditure and household size conditional on PCE. The first is to estimate nonparametric Engel curves that summarize the relationship between the share of expenditure allocated to food and PCE for households of different sizes and ratios of adults to children. The comparison of Engel curves for households of different sizes shows how food shares and per capita food expenditure change with household size, with PCE held fixed. Our nonparametric results are discussed in subsection *B*. As usual, although these nonparametric regressions provide very useful information, the escape from functional form comes at the expense of limiting the number of covariates that can be handled. The problem is that food budgets may be affected by other factors, such as the age composition of children, the fraction of elderly versus nonelderly adults, and the gender composition of the household. Because these factors are likely to be systematically related to household size, there is a risk of confounding scale effects with compositional effects. To handle

such problems, we estimate parametric and semiparametric Engel curves that allow us to control for other factors; the results are presented in subsection C.

B. Nonparametric Methods and Results

Our hypothesis concerns the behavior of per capita expenditure on food for various household types at constant levels of PCE, so that we typically want to test the inequality

$$E\left(\frac{p_f q_f}{n} \middle| i, \frac{x}{n}\right) > E\left(\frac{p_f q_f}{n} \middle| j, \frac{x}{n}\right), \quad (8)$$

where i and j refer to households of two different “types,” defined by household size and demographic composition; n is household size; $p_f q_f/n$ is per capita food expenditure; and x is total household expenditure. Since the budget share of food $w_f = p_f q_f/x$ is also the ratio of per capita food expenditure to per capita total expenditure, PCE, on which we are conditioning, (8) is equivalent to

$$E\left(w_f \middle| i, \frac{x}{n}\right) > E\left(w_f \middle| j, \frac{x}{n}\right), \quad (9)$$

and we work with the food shares and (9).

The inequality (9) is simple enough for it to be possible to make some progress without further parametric specification. Suppose that household type is defined with reference to the number of adults and the number of children in the household so that, for example, we might want to examine whether the food shares are larger for three-adult (3, 0) than two-adult (2, 0) households or for three-adult and three-child (3, 3) than two-adult and two-child (2, 2) households. In both of these cases the ratio of children to adults is the same for the two types of households being compared, so that we are comparing households of different sizes but identical adult/child compositions. Provided that we have enough households of each type, we can fit each side of (9) directly by a nonparametric regression of the food share on x/n or, as we do in practice, on the logarithm of x/n . This is done using Fan’s (1992) local regression smoother, which works as follows. Writing $z = \ln(x/n)$, we choose an interval $[z_0, z_1]$ and then construct a 50-point equally spaced grid over the interval. For each point along this grid, for example at the m th point denoted zz_m , we calculate a weighted regression of the




FIG. 1.—Nonparametric Engel curves, for households with no children and differing numbers of adults: urban Thailand (for bandwidths see table 1).

food share on $\ln(x/n)$ using all observations of the same “type,” where the i th observation gets weight

$$\omega_{im} = \frac{1}{h} \kappa \left(\frac{z_i - zz_m}{h} \right), \quad i = 1, \dots, N; m = 1, \dots, 50, \quad (10)$$

where N is the sample size. For all countries except Britain, survey weights (i.e., inflation factors) are provided in the data, and the weights in (10) are multiplied by the survey weights and then re-scaled to sum to one across households. The term h is a bandwidth, and $\kappa(t)$ is a kernel function, in this case the quartic

$$\kappa(t) = \frac{15}{16} (1 - t^2)^2 1(|t| \leq 1), \quad (11)$$

where $1(\cdot)$ is an indicator function. The bandwidth h is selected by trial and error on the basis of visual inspection; our choices are listed in the notes to the tables below. The results of each weighted regression are then used to compute the expected value of the food share for each point zz_m .

Figure 1 displays the resulting nonparametric Engel curves, using for illustrative purposes urban Thai households with no children




FIG. 2.—Nonparametric Engel curves and confidence bands, for households with no children and one and two adults: urban Thailand (for bandwidths see table 1).

and from one to five adults. The numbers of households of each type range from 156, for five-adult households, to 886, for one-adult households. As expected, the food share declines with PCE for each household type. What is not expected is that, at constant PCE, the food share (and therefore per capita expenditure on food) generally *declines* with household size. In fact, the Engel curve for one-adult households lies above the Engel curves for all larger households over the full range of $\ln(x/n)$. Furthermore, for most values of $\ln(x/n)$, the difference in the food share between one-adult and larger households is statistically significant. Figure 2 shows the Engel curves and their associated confidence bands for one-adult and two-adult households. The standard errors used to construct these confidence bands were obtained from a bootstrap procedure that is detailed in the Appendix. The Engel curves are less precisely estimated at the smallest and largest values of $\ln(x/n)$, where fewer households are located.

In figure 1, the regression lines for the larger Thai households sometimes cross; this is also true of the results that follow for other countries. Because crossing is common and also because it is useful to have summary measures of expected food shares for different household types, we compute weighted averages of the expected

food shares over all values of z for each household type. Specifically, we estimate

$$\int_{z_0}^{z_1} E(w_{f|i}, z) f(z) dz \quad (12)$$

for each household type i . We calculate a kernel density estimate of $f(z)$ over the same grid used for the Fan regressions and then use the results to compute a weighted average of the regression values over the 50 grid points. The important feature of (12) is that the density $f(z)$ is common across all household types so that when we compare (12) across household types, PCE is being held fixed.

To compute (12) we select all households with types (defined by the numbers of adults and children) for which there at least 100 households. (For Pakistan, there were many adult/child cells with between 95 and 100 observations, and we used 95 as the cutoff point in order to use them.) The number of household types with at least 100 observations varies considerably across countries and depends on total sample sizes as well as the patterns of living arrangements that prevail in different countries. Large households, containing both more children and more adults, are generally more common in the poorer countries. The range of grid points is selected such that z_0 is greater than or equal to the first percentile of $\ln(x/n)$ for each household type and is less than or equal to the ninety-ninth percentile of $\ln(x/n)$ for each household type. Fan regressions are then calculated over this range for each household type, and the kernel density estimate $f(z)$ is computed by pooling the data for all households. Because the density is estimated over all cells of adult/child combinations, there are many more data points than for the regressions, which are estimated cell by cell. In consequence, the bandwidth for the density estimation is smaller; in practice, a ratio of 0.5 or 0.7 to 1 gave good results. Exact bandwidths used are in the notes to the tables.

The procedure outlined above results in estimates of (12) that can be compared across all household types. However, because the support of the density of PCE differs from one household type to another, the choice of a common $[z_0, z_1]$ results in the elimination of many households; for example, in urban Thailand, more than 10 percent of households are lost. An alternative that preserves more observations is to make pairwise comparisons of household types, for example between one-adult, zero-child households and two-adult, zero-child households, and evaluate

$$\int_{z_0}^{z_1} [E(w_{f|i}, z) - E(w_{f|j}, z)] f(z) dz. \quad (13)$$

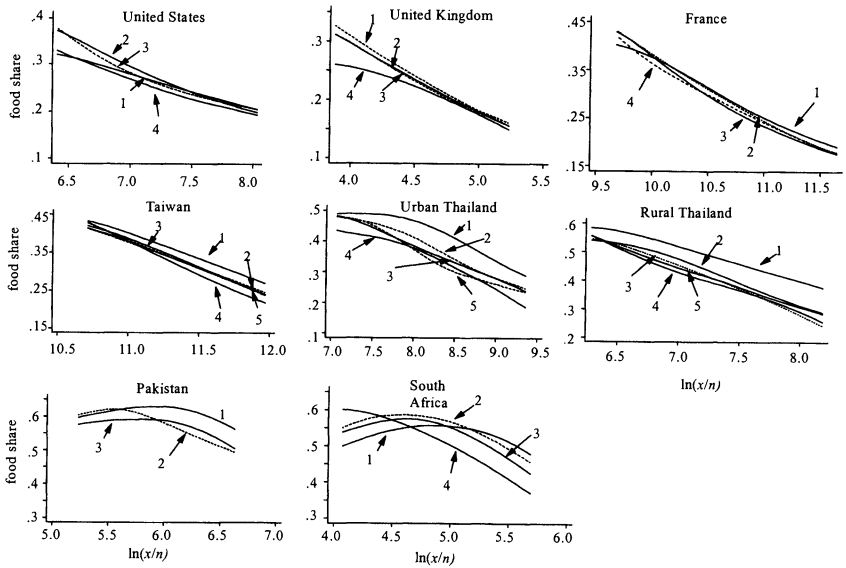


FIG. 3.—Nonparametric Engel curves for all countries, for households with no children and differing numbers of adults (for bandwidths see table 1).

A positive value for (13) shows that, with PCE held fixed, the food share is, on average, larger for those in group i than in group j . In this case, the density $f(z)$ is computed using only households in the two groups i and j , and the intersection of the supports of the two densities is larger than the intersection over all household types.

A selection of the nonparametric Engel curves for each country is shown in figure 3. There are too many household types for each country to show all the results, and we have selected those for all-adult households. Because children are excluded, differences in the Engel curves should not reflect differences in needs between adults and children. The results are easily summarized. In no country is there clear evidence that food share rises with household size, when PCE is held fixed, and in several countries it falls. The downward shifts are most noticeable in urban and rural Thailand and in South Africa, provided that we exclude one-adult households. We also estimated Engel curves (not shown) for households with two children and varying numbers of adults. Here we can also expect an increase in the food share with household size, both because there are more people and because the ratio of children to adults declines as household size increases. However, the general pattern, to which France is the only clear exception, is that the food share falls with the number of adults at constant PCE.

These Engel curves are cumbersome to present, and because they often cross they can be difficult to interpret. In table 1 we present the summary statistics (12) for each of the Engel curves we estimated; these are weighted averages of the food share, where the weights are chosen so as to correct for differences in PCE across types. The standard errors for the average food shares are computed using a bootstrap technique that is detailed in the Appendix. The first six rows of table 1 shows results for childless households with different numbers of adults. The simplest summary of these results is that they fail to show any uniform increase in the food share with the number of adults in the household. There are a few instances in which the food share rises with household size, most noticeably for the change from one-adult to two-adult childless households in the United States, Britain, and South Africa. Results (not shown) for two-child households are similar. In the rest of the table, we show how average food shares change with household size, holding fixed the ratio of adults to children. These patterns are even less supportive of the idea that household size increases the food share. In only one case does the food share strictly increase with household size: the movement from one-adult, one-child households to two-adult, two-child households in France. The table shows exactly the opposite of what the theory predicts: among households with the same adult/child composition and at constant PCE, the per capita expenditure on food *falls* with household size.

Table 2 shows estimates of (13), the pairwise comparisons of food shares across households of different types. The pairs being compared in table 2 have been ordered so that the differences should be *negative* if the food share rises with household size, a hypothesis for which there is little support. Of the 47 differences in table 2, only seven are negative, and in only one case (for the difference between one- and two-adult households in the United States) is the rise in the food share with household size statistically different from zero. While the standard errors are often large (they would be smaller for comparisons of household types that were further apart), so that of the 40 differences with *positive* values only 16 have *t*-values of two or more, the results provide more support for the hypothesis that per capita food expenditure declines with household size than for the original hypothesis, that it rises.

C. *Parametric and Semiparametric Methods and Results*

The results described above do not fully account for differences across households in factors that may affect food expenditures and

TABLE 1
FOOD SHARES: AVERAGES WEIGHTED BY DENSITY OF PER CAPITA EXPENDITURE

Adults, Children	United States	Great Britain	France	Taiwan	Urban Thailand	Rural Thailand	Pakistan	South Africa
1, 0	.262 (.004)	.230 (.002)	.290 (.004)	.355 (.004)	.425 (.006)	.496 (.010)534 (.016)
2, 0	.274 (.003)	.237 (.002)	.287 (.003)	.332 (.002)	.376 (.005)	.439 (.006)	.614 (.020)	.551 (.011)
3, 0	.269 (.005)	.230 (.004)	.279 (.004)	.335 (.003)	.364 (.006)	.421 (.006)	.587 (.017)	.533 (.010)
4, 0	.252 (.006)	.214 (.006)	.279 (.004)	.331 (.002)	.338 (.006)	.411 (.008)	.577 (.020)	.509 (.014)
5, 0333 (.003)	.352 (.009)	.423 (.013)
6, 0321 (.006)
1, 1	.253 (.008)	.213 (.006)	.236 (.010)	.319 (.007)429 (.012)607 (.012)
2, 2	.238 (.003)	.191 (.002)	.246 (.003)	.319 (.002)422 (.005)	.601 (.012)	.534 (.009)
3, 3	.242 (.006)311 (.005)393 (.031)	.601 (.012)	.489 (.014)
4, 4569 (.017)	.482 (.024)
1, 2	.255 (.008)582 (.015)
2, 4	.233 (.003)504 (.016)
2, 1322 (.003)	.369 (.006)	.439 (.005)	.615 (.014)	.560 (.008)
4, 2314 (.003)	.287 (.010)	.387 (.014)	.571 (.015)	.471 (.014)

NOTE:—These numbers are weighted averages of the expected food shares conditional on different levels of PCE, with weights given by the kernel estimate of the density of per capita total household expenditure. The weights are common across household types, so PCE is held constant as one moves down each column. The bootstrapping procedure used to compute the standard errors is described in the Appendix. We used a bandwidth of 1 for the regressions. For the estimates of the density of PCE we used a bandwidth of .5 for the United States, Taiwan, Pakistan, and South Africa and .7 for the United Kingdom, France, and Thailand. The same bandwidths are used when bootstrapping.

TABLE 2
DIFFERENCES IN FOOD SHARES ACROSS HOUSEHOLDS OF DIFFERENT TYPES: AVERAGES WEIGHTED BY DENSITY OF $\ln(x_i/n)$

(Adults, Children)	United States	Great Britain	France	Taiwan	Urban Thailand	Rural Thailand	Pakistan	South Africa (Africans)
(1, 0) - (2, 0)	-.009 (.004)	-.004 (.003)	.001 (.004)	.021 (.004)	.051 (.007)	.061 (.010)	...	-.008 (.013)
(2, 0) - (3, 0)	.004 (.005)	.009 (.005)	.010 (.005)	.000 (.003)	.011 (.009)	.023 (.009)	.053 (.097)	.029 (.013)
(3, 0) - (4, 0)	.019 (.009)	.015 (.006)	.000 (.009)	.004 (.003)	.029 (.009)	.006 (.010)	.040 (.087)	.003 (.021)
(4, 0) - (5, 0)	-.002 (.004)	-.018 (.011)	-.012 (.016)
(5, 0) - (6, 0)012 (.005)
(1, 1) - (2, 2)	.011 (.011)	.016 (.007)	-.010 (.010)	.008 (.007)004 (.016)036 (.023)
(2, 2) - (3, 3)	.001 (.012)012 (.005)024 (.054)	.026 (.019)	.015 (.018)
(3, 3) - (4, 4)063 (.034)	.008 (.022)
(1, 2) - (2, 4)	.050 (.017)041 (.021)
(2, 1) - (4, 2)010 (.005)	.076 (.011)	.056 (.015)	.084 (.043)	.033 (.016)

NOTE.—Each row contains estimates of the difference in food shares between households with the first (adult, child) pair and the second (adult, child) pair listed. For example, the row denoted "(1, 0) - (2, 0)" is the food share for one-adult, zero-child households minus the food share for two-adult, zero-child households. The estimates for each row are based on an estimate of the density of per capita expenditure only for households with one of the two adult/child combinations listed. See the note to table 1 for the bandwidths used.

are correlated with household size. For example, there may be differences between the food consumption needs of men and women, older and younger adults, older and younger children, and working and nonworking individuals (especially when the work involves heavy manual labor). To deal with these factors, we must move to a parametric or semiparametric regression. The simplest that we consider has the form

$$w_f = \alpha + \beta \ln \frac{x}{n} + \gamma \ln n + \sum_{k=1}^{K-1} \eta_k \frac{n_k}{n} + \zeta \cdot \mathbf{v} + u, \quad (14)$$

where n_k/n is the ratio to household size of household members who fall in one of K groups defined by age and sex. This specification is designed to separate the effects of household size, n , from household composition as represented by the ratios n_k/n . The parameter γ thus corresponds to the conceptual experiment of making a household larger by replication of both people and resources.

The vector \mathbf{v} includes a variety of variables. For all countries it includes measures of the fraction of adults in the household who work. Households with more working members could have different food budget shares for several reasons. Workers may eat more of their meals in restaurants, and expenditure on restaurant meals includes a service component that may increase the food share. On the other hand, the increased commuting expenses associated with work could affect the food share in the opposite direction. Changes in caloric requirements due to the intensity of work must also be accounted for, especially in poor countries in which agricultural employment is common. For the poorer countries we include controls for the fractions of adults who work in a variety of activities, including work on a family farm and work as agricultural wage labor. In addition, we control for such things as farm status, rural/urban status, region of residence, and (where appropriate) the calendar quarter in which the survey was conducted. These variables differ somewhat across countries with differences in survey instruments; see the Appendix for definitions. For the United States, we also include indicators of whether the household received food stamps or public housing. The value of the subsidy on public housing is not included in expenditure. Therefore, households that receive public housing may have higher measured food shares since true housing expenditure is understated. This is not an issue with food stamps since the full value of food purchased with food stamps is included in the measure of food expenditure. There is an active debate about whether food stamps exert more than a standard income effect on food expenditure (see Moffitt 1989; Fraker, Martini, and Ohls 1995), and the food stamp dummy allows for the possibility.

Our main concern with the simple form in equation (14) is the possible misspecification of the response of the food share to PCE, x/n , given the strong (negative) correlation in all our data sets between PCE and household size n . We therefore consider two other specifications in which the shape of the Engel curve is made increasingly flexible. The first uses a Fourier flexible functional form that includes not only the level and square of $\ln(x/n)$ but also $\sin[j \ln(x/n)]$ and $\cos[j \ln(x/n)]$ for values of $j = 1, 2, 3$ (see, e.g., Eubank and Speckman 1990). In theory, the Fourier specification requires that $\ln(x/n)$ be rescaled so that its range is less than 2π , but this restriction was automatically satisfied by our data. The second method allows x/n to enter nonparametrically, while maintaining the linear parametric structure for the other variables in (14). Semi-parametric models of this type are discussed in Engle et al. (1986) and Robinson (1988). The procedure we use, due to Estes and Honoré (1995), sorts the data on x/n and then first-differences prior to an ordinary least squares (OLS) regression of the change in the food share on the change in all variables *except* x/n . Provided that x/n appears in the regression additively through some continuous function $h(x/n)$, as the sample size increases, successive values will be forced closer together so that the differencing will eventually remove the effect.

A second concern is measurement error in $\ln(x/n)$. In the linear specification (14), standard measurement error in $\ln(x/n)$ would bias the estimate of β toward zero. However, since the food share and PCE are constructed from the same expenditure information, errors in w_f and $\ln(x/n)$ are almost inevitably correlated. The correlation can be positive or negative, depending on whether food or nonfood contains more measurement error and depending on the relative importance of each in the budget, so that the direction of the bias of β cannot be determined a priori. Since PCE and household size are negatively correlated, it is likely that measurement error in the former will bias estimates of γ , the parameter of interest, although it is impossible to predict in which direction. Fortunately, each survey contains information on household income as well as expenditure. Since income is highly correlated with actual total expenditure but is measured independently, it is a good candidate to be an instrument for expenditure. We present instrumental variables estimates of equation (14), using the logarithm of per capita income as the instrument for PCE. A problem with this procedure, of particular importance in the poorer countries, is that the imputed value of home-produced consumption items is included in measures of both income and expenditure, and this can introduce common errors in the two variables. To counter this problem, we measure

“income” as cash income, excluding all items that appear in both income and expenditure. Note that this instrumentation of consumption by income would also deal with the endogeneity of consumption relative to income, for example in a simple consumption function, provided that we were prepared to maintain the exogeneity of income.

Table 3 presents the estimated coefficients on the logarithm of household size (multiplied by 100): they are the estimated percentage point changes in the food share if household size were to be doubled and PCE and household composition were held constant. In panel A, the regressions are run for all households. In panel B, the regressions are run for households with at least two adults. The motivation for excluding one-adult households comes from the results in table 2, which show large differences across countries in how the food share changes as we move from one-adult to two-adult households. Specifically, in the United States, Britain, and South Africa, one-adult households have lower average food shares than two-adult households, whereas in Taiwan and Thailand, one-adult households have substantially higher average food shares than two-adult households. By excluding one-adult households, we can examine whether the results of panel A are dominated by the differences between one-adult and other households. For the linear, Fourier, and instrumental variables estimates we present *t*-statistics based on standard errors that have been corrected for heteroskedasticity, for correlation between households in the same geographical survey cluster (for countries that provide clustering information), and for correlation across multiple observations for the same household (for the United States). For the Estes-Honoré results, we present the OLS standard errors, which are certainly too low; we have not found any practical way of allowing for the clustering when calculating the standard errors for this estimator, but the estimates remain valuable as a cross-check on the Fourier and linear specifications.

The results in table 3 indicate that, for all countries except Britain, larger household size is associated with a lower food share. The coefficients for Britain are positive in panel A and negative when one-adult households are excluded, but they are generally not significantly different from zero. The differences between panels A and B for all countries are relatively small. Furthermore, the coefficients are not greatly affected by the choice of functional form for PCE, with the linear, Fourier, and Estes-Honoré specifications yielding similar results. The instrumental variables estimates are typically not much different from the OLS results, except for two countries (Pakistan and South Africa) in which instrumenting PCE yields an even larger negative coefficient on the logarithm of household size.

TABLE 3

REGRESSION COEFFICIENTS ON THE LOGARITHM OF HOUSEHOLD SIZE IN FOOD SHARE REGRESSIONS ($\times 100$)

Model	United States	Great Britain	France	Taiwan	Urban Thailand	Rural Thailand	Pakistan	South Africa (Africans)
A. All Households								
Linear	-77 (2.8)	.05 (2)	-.75 (2.3)	-1.39 (7.1)	-5.37 (14.6)	-5.48 (12.7)	-5.65 (8.9)	-5.98 (10.9)
Fourier	-.80 (2.9)	.26 (1.0)	-.69 (2.1)	-1.40 (7.2)	-5.40 (14.6)	-5.43 (12.6)	-5.50 (8.9)	-6.46 (12.4)
Estes-Honoré	-.90 (3.5)	.09 (.3)	-.66 (2.0)	-1.43 (6.1)	-5.13 (9.4)	-5.64 (8.8)	-4.97 (7.15)	-7.22 (9.1)
IV	-1.12 (4.9)	.14 (.6)	-.42 (1.4)	-1.84 (12.3)	-5.25 (16.2)	-6.18 (13.1)	-7.15 (6.2)	-9.37 (17.2)
Sample size	20,504	7,415	8,485	16,434	5,188	5,604	4,720	6,476
B. Households with More than One Adult								
Linear	-1.01 (1.9)	-.64 (1.3)	-1.18 (2.1)	-.97 (4.1)	-5.17 (10.1)	-4.78 (8.8)	-5.86 (8.9)	-7.45 (10.8)
Fourier	-1.41 (2.6)	-.58 (1.2)	-1.23 (2.5)	-.96 (4.1)	-5.15 (10.0)	-4.79 (8.7)	-5.64 (8.8)	-7.54 (11.7)
Estes-Honoré	-2.11 (3.5)	-1.12 (2.1)	-1.08 (1.9)	-.81 (2.3)	-4.61 (5.8)	-5.48 (6.0)	-5.77 (7.6)	-8.28 (7.9)
IV	-1.04 (2.2)	-.32 (.6)	-.66 (1.2)	-1.54 (7.9)	-5.00 (10.7)	-5.54 (9.5)	-7.53 (6.6)	-11.07 (15.4)
Sample size	13,771	5,005	6,407	14,942	4,176	5,095	4,640	4,955

NOTE.—Absolute t -values are in parentheses. All regressions contain the logarithm of per capita total outlay, the logarithm of household size, the ratios to household size of numbers of people by age and sex category, and a set of other controls listed in the Appendix. The models differ in the treatment of $\ln(x/n)$; in the linear form it appears as is, and in the Fourier form through its level, its square, and as sin and cos of $\ln(x/n)$ for values of $j = 1, 2, 3$. In the Estes-Honoré model, the data are sorted by $\ln(x/n)$ and first-differenced prior to an OLS regression. For the linear and Fourier models, standard errors and the resulting t -values are Huberized to account for the cluster structure of the samples in the case of Pakistan, South Africa, Taiwan, and Thailand and for the fact that many households are repeatedly interviewed in the United States. Geographical clustering information is not available on public-use samples for the United States, Great Britain, and France. For all countries except Britain, the linear and Fourier models are estimated using weighted least squares, with weights inversely proportional to the sampling weight provided by the survey. In the instrumental variables model, $\ln(x/n)$ enters linearly but is instrumented using the logarithm of per capita income; the income measure excludes in-kind items that are used in the calculation of both income and expenditure. The standard errors for the Estes-Honoré estimates should be regarded as lower bounds.

The most striking feature of table 3 is the almost uniform decline in the coefficient on household size as we move from richer to poorer countries. (Recall that for South Africa, we use data for only African households.) For African households in South Africa—possibly the poorest group in the analysis—as well as for the households in Thailand and Pakistan, the negative effects of scale are large: doubling household size (and household total expenditure) causes a fall in the food share of between five and 10 percentage points. We can translate this into the effect on food *expenditure* per head by noting that the elasticity of food expenditure per capita with respect to household size at constant PCE is γ/w_f . The average food share for African households is .5, so that doubling both household size and household total expenditure gives a decline of per capita food expenditure of between 10 and 20 percent.

The ranking of the coefficients from rich to poor countries strengthens the puzzle since according to the theoretical analysis, the failure of the prediction—that food expenditures per capita should rise with scale at constant PCE—is more likely to occur among rich households in which the expenditure elasticity of food is low and the price elasticity relatively high. On the contrary, it is in the three poorest countries that the failure is most marked. One possible explanation is that “expenditure” on home-produced food is systematically understated, and home-produced food makes up a larger share of total food expenditure for larger households in the poorer countries. However, this explanation cannot account for our results. We reestimated the equations in table 3 for Thailand, Pakistan, and South Africa using only nonfarm households that had no family members who were self-employed in agriculture. Although this resulted in the loss of many observations—for example, the rural Thai sample fell from 5,604 to 1,913 observations—the estimates of the effects of household size on the food share were essentially unchanged.

Table 4 examines the second theoretical prediction, that with PCE and household size held constant, the substitution of a child for an adult should decrease the per capita demand for food and thus the food share. The coefficients in the table come from regressions with the same form as those in table 3. Since we wish to compare the effect of a child with the effect of an adult and since the full set of results differ by the sex of both the adults and children, we reestimate the equations suppressing the gender effects and with prime-age adults as the omitted category. The reported numbers are the regression coefficients on the ratio of the youngest child group (aged 0–4 in Britain and Thailand, and aged 0–5 in all other countries) to total household size, multiplied by 100. As predicted by the

TABLE 4

REGRESSION COEFFICIENTS ON THE RATIO OF YOUNG CHILDREN TO HOUSEHOLD SIZE ($\times 100$)

Model	United States	Great Britain	France	Taiwan	Urban Thailand	Rural Thailand	Pakistan	South Africa (Africans)
A. All Households								
Linear	-8.11 (9.7)	-12.06 (15.5)	-12.12 (14.2)	-6.46 (13.5)	-1.04 (.7)	-.09 (.1)	5.68 (3.3)	9.62 (5.8)
Fourier	-8.46 (10.2)	-12.78 (16.3)	-12.74 (15.0)	-6.52 (13.6)	-1.09 (.7)	-.23 (.1)	4.06 (2.4)	8.47 (5.3)
Estes-Honoré	-8.81 (9.2)	-11.77 (12.9)	-13.43 (13.6)	-6.26 (9.4)	-1.56 (.8)	1.66 (.8)	3.11 (1.5)	8.46 (4.1)
IV	-8.03 (12.5)	-11.15 (12.9)	-11.31 (11.2)	-6.97 (16.5)	-.40 (.3)	-2.03 (1.3)	3.48 (1.7)	6.55 (4.1)
B. Households with More than One Adult								
Linear	-9.17 (9.6)	-11.24 (12.5)	-11.29 (39.6)	-6.48 (13.3)	-.82 (.5)	.71 (.4)	5.67 (3.3)	9.35 (4.8)
Fourier	-8.87 (9.3)	-11.66 (12.9)	-11.60 (12.1)	-6.51 (13.3)	-.78 (.5)	.64 (.4)	4.08 (2.4)	8.39 (4.5)
Estes-Honoré	-8.71 (9.2)	-10.07 (10.3)	-11.83 (10.5)	-6.57 (10.0)	-1.41 (.6)	2.68 (1.0)	2.80 (1.4)	7.28 (2.8)
IV	-8.69 (12.0)	-10.55 (10.7)	-10.67 (10.2)	-7.04 (16.5)	.41 (.3)	-1.30 (.8)	3.21 (1.5)	5.83 (3.2)

NOTE.—Absolute *t*-values are in parentheses. These regressions have the same dependent and independent variables as those described in table 3, except there are no gender distinctions in the ratios to household size of people in various age categories. In this table we report the coefficient on the ratio of people in the youngest age group, aged 0-4 in Britain and Thailand and aged 0-5 in all other countries. The omitted age category is prime-age adults, aged 18-64 in the United States, Britain, France, and Taiwan; aged 16-59 in South Africa and Pakistan; and aged 16-29 in Thailand. The Thai results are similar if the omitted category is extended to include those between the ages of 16 and 54. See table 3 for sample sizes and other details.

Barten model, the coefficients are negative and significant for Great Britain, the United States, and Taiwan but contrary to the model are insignificantly different from zero for Thailand, zero or positive for Pakistan, and large and positive for South Africa. The most obvious feature of the cross-country pattern is that the substitution of a child for an adult reduces the food share in the more developed countries and leaves it unchanged or increases it in the poorer countries.

In contrast to the patterns between food and household size, the effects of children are more easily rationalized if there are fixed costs associated with the presence of adults and children. Among poor households, such as those in South Africa, the Barten model predicts that the substitution of a child for an adult will decrease per capita expenditures on food because children eat less than adults and because the offsetting income effects cannot reverse the sign when the income elasticity of food is less than one. But if there are fixed costs of both adults and children and if the former are larger than the latter, the replacement of an adult by a child will release resources that will increase the demand for food. Now consider a rich country and suppose that the fixed costs of children are relatively large because children are generally expensive in rich countries. We would then predict what we observe, that the negative effect associated with the replacement of an adult by a child is not offset and possibly even enhanced, so that food expenditure per capita falls.

When the food share falls, the shares of other goods must increase. To investigate which goods increase and whether there are systematic patterns across countries, we estimated the regressions (14) for a set of commodities including clothing and footwear, housing (which is broadly defined to include household furnishings and utilities), recreation and entertainment, and alcohol and tobacco (in Pakistan this contains only tobacco). This list does not cover all commodities, but together with food it covers more than half of the average budget in all the countries. In some countries the data permitted us to break food into food consumed at home and food eaten away from home, and we also estimated budget share equations for these two groups.

Table 5 reports the coefficient on household size for the budget share equations. We report the results for all goods only for the sample of households containing at least two adults, though the results for all households are also shown for the case of food. There are several noteworthy patterns. First, it is not the case that the negative effect on food of household size can be attributed to a reduction in spending on restaurant meals, which have a large service component and are likely candidates for replacement when a larger family size

TABLE 5

REGRESSION COEFFICIENTS ON THE LOGARITHM OF HOUSEHOLD SIZE IN BUDGET SHARE REGRESSIONS ($\times 100$)

Share of Expenditure Allocated to	A. All Households							South Africa (Africans)
	United States	Great Britain	France	Taiwan	Urban Thailand	Rural Thailand	Pakistan	
Food	-.77 (2.8)	...	-.75 (2.3)	...	-5.37 (14.6)	-5.48 (12.7)	...	-5.98 (10.9)
Food at home	-.10 (.4)	...	-.57 (1.8)	...	-3.77 (10.3)	-6.05 (14.0)	...	-6.81 (12.0)
Food away from home	-.67 (4.2)	...	-1.31 (5.4)	...	-1.60 (4.1)	.56 (2.1)62 (5.1)
	B. Households with More than One Adult							
Food	-1.01 (1.9)	-.64 (1.3)	-1.18 (2.1)	-.97 (4.1)	-5.17 (10.1)	-4.78 (8.8)	-5.86 (8.9)	-7.45 (10.8)
Food at home	-1.44 (2.8)	...	-2.52 (4.8)	...	-6.48 (13.4)	-6.06 (11.4)	...	-7.96 (11.4)
Food away from home	.43 (1.6)	...	1.34 (4.4)	...	1.32 (2.9)	1.28 (4.6)48 (3.4)
Clothing	1.94 (6.5)	4.09 (7.4)	2.29 (6.3)	1.32 (4.9)	.48 (1.7)	.94 (2.8)	3.09 (7.2)	.65 (4.2)
Entertainment	.87 (3.0)67 (3.1)	4.82 (21.3)	-.41 (4.9)	.09 (1.0)14 (3.4)
Housing, including utilities and furnishings	-5.23 (6.6)	-13.27 (18.2)	-9.20 (16.6)	-7.37 (23.7)	-3.01 (5.9)	-4.48 (8.6)	.20 (.4)	.84 (.6)
Alcohol and tobacco	-1.46 (6.1)	2.81 (4.5)	-1.24 (6.2)	-1.35 (16.5)	-.75 (3.6)	-.19 (.9)	-.82 (4.4)	-.16 (.7)

NOTE.—Absolute *t*-values are in parentheses. See table 3 for the sample sizes and other control variables. For Pakistan, the category "alcohol and tobacco" contains only tobacco. In all countries, clothing includes clothing and footwear for adults and children. The entertainment category differs somewhat across countries but generally contains travel and recreation expenses. Housing includes utilities and home furnishings in addition to maintenance expenses and the rental value of the home.

makes it more attractive to incur the fixed costs of preparing food at home. For all the countries and when one-adult households are excluded, expenditure on food away from home *rises* with household size at constant PCE, so that if we confine ourselves to food eaten at home, the puzzle is intensified, not resolved. Households with only one adult spend more on meals away from home than other households, and in the rich countries and urban Thailand, this fact dominates the results for all households. In France, but nowhere else, the effect is strong enough that the effect of household size on food switches sign and becomes positive once food eaten out is excluded. But this rather satisfactory result comes entirely from the comparison of one-adult with other households and is not present in panel B of the table, where one-adult households are excluded.

Second, the coefficients on household size are generally positive for clothing and entertainment and have mixed signs for alcohol and tobacco. Clothing, alcohol, and tobacco would usually be classified as private goods, and entertainment has both public and private components, so that the results neither confirm nor contradict the theory. Third, and perhaps most important, are the effects of household size on housing, the most public of the goods in our list. If there is limited substitution toward housing in response to its fall in effective price with household size, the effect should be negative, which is what we observe for all but the two poorest countries, Pakistan and South Africa. The results for Pakistan and South Africa can be reconciled with the theory if housing is sufficiently price elastic, which would also help explain the large negative effect of household size on food in these countries. Even so, we are still left with the puzzle of why the demand for food has the highest substitution elasticity in those countries whose inhabitants are closest to subsistence.

D. *Explanations*

There are a number of possible explanations for our finding that, given PCE, larger households spend less per head on food. None holds out the promise of resolving the puzzle, though several are clearly worth further research.

i) *Direct economies of scale in food consumption.*—Larger households may benefit from buying in bulk and thus paying less per unit. As a result, expenditures may fall even when quantities are rising. Robin (1985) provides evidence that is consistent with this from a food survey in France. For some other countries, for example, India and Pakistan, there exist data on unit values, and they show the opposite effect. At constant PCE, unit values increase with household size. Indeed, this is what Prais and Houthakker (1955) originally found

in their 1938 British data and what they attributed to the welfare-enhancing effects of economies of scale: larger households with the same PCE are better off and therefore buy higher-quality items that cost more per unit. There is similar evidence from earlier work by one of us using data from Côte d'Ivoire, Indonesia, Pakistan, and India (see Deaton 1997, chap. 5).

ii) *Economies of scale in food preparation.*—These almost certainly exist, but they do not help to explain the phenomenon. Our data pertain to food purchases, not to the combined expenditure on food and time spent by household members in preparing food. Economies of scale in preparation cheapen an input—time—that is complementary with food consumption. They should therefore lead to increased (at home) consumption of food at a lower overall cost (including the time cost), which only deepens the puzzle. Another issue, related to point i above, is that preparation time can be reduced by purchasing prepared foods, restaurant meals being the extreme example. As household size rises, households may switch to more home preparation and away from already-prepared foods, causing a reduction in expenditure per person. However, the evidence in table 5 on restaurant meals and meals eaten at home indicates that this is unlikely to resolve the puzzle.

iii) *Wastage.*—Larger households are better at eliminating waste through better management of storage, refrigerators, and leftovers. This is hard to rule out, though we find it hard to imagine that there is a great deal of avoidable waste among poor households in South Africa, Pakistan, or Thailand.

iv) *Collective models.*—We have worked with a standard unitary model of household preferences. While it would be possible to consider more general models, the basic argument about economies of scale, housing, and food does not seem to depend on the unitary assumptions. It is possible that larger households have systematically different compositions of people, people who might have different tastes for food. Indeed, the regressions that are only summarized in tables 3 and 5 show that the sex composition of households has large effects on the composition of the budget. But since these regressions control for the sex and age composition of the household, there is no obvious explanation along this route.

v) *Price elasticity of food.*—If we maintain the analytical framework of Section II and reject the supposition that there are economies of scale in food, we are left with the implication that the compensated own-price elasticity of food is large, perhaps (numerically) in excess of a half. This is most plausible in the United States, France, or Britain, where (e.g.) new parents may switch from restaurant meals to cheaper meals at home. But as we have seen, the behavior of food

eaten away from home does not conform to this explanation. Nor is the elasticity story supported by evidence of a switch into housing: housing expenditures per head fall along with food expenditures per head, except in the two poorest countries, where a high price elasticity of food is the hardest to believe (although this does not mean that it does not exist).

vi) *Measurement error*.—Food expenditures may be systematically more understated the larger the household. Surveys typically interview a single respondent, who provides detail on expenditures for all household members. The respondent may not know everything about all household members, and the understatement could be larger the more there are to know about. But for such biased reporting to account for our results, the respondent must be worse informed about others' food consumption than about others' non-food consumption. There seems no reason why this should be the case, especially in households in which everyone eats from a common pot so that food expenditures are the least likely to be understated.

vii) *Calorie overheads*.—Households in poor countries may face fixed costs in terms of food or calories. For example, one person in the household may have to collect firewood for fuel and may need to be given more food in order to undertake the task. If so and if firewood needs rise less rapidly than household size, larger households will need fewer calories per head, and the demand for food may fall.

viii) *Intrahousehold inequality*.—If the food Engel curve is concave, the average demand for food in the household will be reduced the more unequally income is distributed within it. If larger households are systematically more unequal, per capita food consumption could fall with household size. It would also have to be the case that this relationship between inequality and household size is more pronounced in the poorer economies. The common pot is also a problem here: when all household members eat from the same pot, it is implausible that the food allocation is unequal.

Appendix

A. *Bootstrapping Procedure*

All the surveys we work with have a clustered design, and ignoring geographical clustering when bootstrapping will usually understate sampling variability (see, e.g., Deaton 1997, p. 60). In four of the countries—Taiwan, Thailand, Pakistan, and South Africa—we can identify which households are in the same geographical cluster, but in the others no cluster identifiers are provided. The U.S. data have the additional feature that each household

may be surveyed up to four times during the year and so can contribute up to four observations. The standard errors in tables 1 and 2 and the confidence bands in figure 2 are based on a bootstrapping procedure that takes into account clustering when information is available. For Great Britain and France, for which we have no information, we bootstrap by drawing 500 random samples with replacement and estimating the Engel curves, densities of PCE, and averages in tables 1 and 2 for each sample drawn. For Taiwan, Thailand, Pakistan, and South Africa, for which we have cluster identifiers, we bootstrap by drawing random samples of *clusters* with replacement and using all households in each cluster selected. In the United States, for which we do not have information on geographical clustering but we do know whether observations are drawn from the same household, we bootstrap by drawing random samples of *households* and then using all observations on the selected households. Our standard errors are likely understated for the countries in which clustering information is not available.

B. *The United States*

The U.S. data are drawn from the 1990 Consumer Expenditure Survey. The CEX interviews approximately 5,000 households per quarter, with a total of 20,517 household interviews in 1990. We excluded 13 households that either had no members aged 18 or over or had negative expenditure (which is possible in the CEX because of the treatment of refunds from health insurance companies). Unlike the other surveys we use, the CEX has a panel element: households may be surveyed in five consecutive quarters, and complete expenditure information is collected in the last four surveys. In our analysis we do not aggregate each household's information across quarters, but instead treat the information from each quarterly survey as a separate observation. Constructing annual expenditure items for each household is complicated by the fact that many households do not complete all the surveys.

The major variables are defined as follows. To calculate expenditure items, we use the summary expenditure measures constructed by the Bureau of Labor Statistics (BLS). Total expenditure is the sum of all expenditure items except the purchase of vehicles, payments for personal insurance, and contributions to pension and retirement plans. Food expenditure includes food eaten at home and away from home. The income measure used for the instrumental variables estimates of the food share equation is total after-tax income. Adults are defined as all household members aged 18 or older, and children are those aged 17 or less. We use the sample weight called "finlwt21" in the BLS documentation. This weight is a measure of the numbers of U.S. households the surveyed household represents in the quarter the survey was conducted.

The food share regressions shown in tables 3, 4, and 5 include controls for a number of variables. For tables 3 and 5 they include the ratios of numbers of members in various age and sex categories to total household size. The age and sex categories consist of males (females) in age groups 0-5, 6-11, 12-17, 18-64, and 65 and over. The regressions shown in table

4 include variables for the ratios of all members (i.e., both males and females) in each age category to total household size. The regressions in all tables include dummy variables for the quarter in which the survey was conducted, the reported race of the reference person in the household, whether the household lives in an urban area, whether the household receives food stamps, and whether the household lives in public housing. We also include a variable that equals the fraction of adults (aged 18 and older) who are "earners." The BLS codes anyone who earns income as an earner, regardless of whether the work was part-time or full-time. The instrumental variables results reported in tables 3 and 4 instrument the logarithm of PCE with the logarithm of per capita after-tax income. There are 1,666 households for which after-tax income is zero or negative. These households are excluded from the instrumental variables regressions. We estimated the linear OLS regressions (shown in row 1 of tables 3 and 4) on a sample that excluded these 1,666 households, and this had very little effect on the OLS results. Therefore, differences between the instrumental variables and OLS results are not due to changes in the sample.

C. Great Britain

The British data are drawn from the 1992 Family Expenditure Survey. The basic data consist of observations on 7,418 households. We exclude three households that have no members aged 18 or more. The survey has a quarterly structure, in that surveyed households are asked about expenditures and income in the quarter before the survey, and roughly one-fourth of the households are surveyed in each quarter. However, there is no panel element: each household is surveyed only once.

The major variables are defined as follows. Total expenditure is measured as expenditure on all items, including consumer durables. Food expenditure includes food consumed at and away from home. The income measure is "normal gross income," which excludes imputed income from "self-supply" and other imputed items that could appear in expenditure. As for the United States, family size is the number of household members, and those aged 18 and older are treated as adults. The Family Expenditure Survey is designed to be self-weighting, so no sample weights are used. No information on geographic clustering is provided.

The regressions include controls for the ratio to household size of the numbers of males and females in age groups 0-4, 5-17, 18-64, and 65 and over. The gender differences in these groups are suppressed in the regressions shown in table 4. The other controls are dummies for the quarter in which the interview was conducted and the fraction of adults who are in the labor force. An adult is coded as a member of the labor force if he or she was an employee, was self-employed, or was looking for work at the time of the survey. The instrumental variables results use the logarithm of per capita income to instrument the logarithm of PCE. Only eight households have nonpositive income; excluding these eight households has virtually no effect on the OLS results.

D. France

The French data are taken from the 1989 Family Budget Survey (*enquête budget de famille*), which surveyed 9,038 households. We exclude 551 observations on households that could not provide complete income information, so our sample contains 8,487 observations. The survey asks questions about income in the year before the survey. The reference period for expenditure information varies by type of expenditure. For example, the survey collects information on expenditure on major durables over the past year, information on medical expenditure over the last six months, and expenditure on utilities (telephone, heat, etc.) over a reference period determined by the respondent. All items are grossed up to represent annual figures.

The variables are defined as follows. Total expenditure is expenditure on all goods, including durables. Food expenditure includes food at home and away from home. Income is after-tax money receipts. Adults are those aged 18 and older. The survey provides sample weights, which we use, but does not provide information on geographical clustering.

The regressions include controls for the ratio to household size of the numbers of males and females in age groups 0–5, 6–11, 12–17, 18–64, and 65 and over. As for the other countries, the gender differences are suppressed for the regressions reported in table 4. We included dummy variables that indicate whether the household is from a rural area, small town, medium town, large town (e.g., city), or Paris. An additional control variable is the fraction of adults who are in the labor force. There were 41 households with nonpositive income, and they are excluded from the instrumental variables regressions. Estimates of the linear OLS regression change very little when these 41 households are excluded.

E. Taiwan

The Taiwanese data are taken from the 1990 Survey of Personal Income Distribution. We use information on all the 16,434 households that were surveyed. The surveys contain information on income and expenditure in the year before the survey, although some expenditure items had a shorter reference frame.

The major variables are defined as follows. Expenditure is total household expenditure on all items, including durables. Food is food consumed at home and away from home. The income measure is total household after-tax income. The survey documentation provides little information on how the items that go into total income were constructed, and we cannot determine whether the values of home-produced items are included as income. Thus it is possible that measurement errors in income and expenditure will be correlated because common items are included in both, in which case income is not a suitable instrument for expenditure. Adults are defined to be those aged 18 and older. Both sample weights and geographical clustering information are provided. The 16,434 households are located in 1,493 clusters, with the numbers of households per cluster ranging from two to 70.

The control variables included in the regressions include a set of variables that measure the ratio to household size of the numbers of males and females in age groups 0–5, 6–11, 12–17, 18–64, and 65 and older. As for the other countries, the gender differences are suppressed for the results shown in table 4. Other controls include dummies for whether the household runs a farm and for whether the household lives in a rural area, small town, or city. We also include measures of the fraction of adults who are self-employed in agriculture, the fraction who are agricultural employees, and the fraction who do nonagricultural work (either self-employed or as employees). There were no nonpositive income numbers, so the samples for the instrumental variables and OLS estimates are the same.

F. Thailand

The Thai data are taken from the 1992 Socioeconomic Survey. This is a survey of 13,456 households. We work with the 10,817 households that live in either urban or rural areas, excluding those that live in the peri-urban “sanitary districts.” We also exclude households with no adult members, and our final samples include 5,200 urban and 5,604 rural households. We analyze urban and rural households separately. The Socioeconomic Survey asks households about expenditure on food and the value of home-produced food consumed in the week before the survey. Information on other consumption items (including those that are home-produced) is also collected, with reference periods that vary according to the good. All figures are then grossed up to represent annual figures. Information on annual income is collected. The survey asks households about both cash and in-kind income.

The major variables are defined as follows. Total expenditure refers to all goods and services, including durables. It includes both purchased and home-produced items. Food expenditure includes food (purchased and home-produced) consumed at home and away from home. The income measure used for the instrumental variables estimates is cash income. Adults are those aged 16 and older. Both sample weights and geographic clustering information are available. On average, there are 6.4 rural households per cluster and 11.2 urban households per cluster.

The control variables included in the regressions include the ratios to household size of males and females in age groups 0–4, 5–9, 10–14, 15–29, 30–54, and 55 and older. The results in table 4 suppress the gender differences in the categories. Other controls are region dummies, dummies for the calendar quarter in which the survey was conducted, a dummy for whether the household is a farm household, the fraction of household members who are servants, the fraction of adults whose main occupation is agricultural self-employment, the fraction whose main occupation is agricultural employment, and the fraction who do nonagricultural work (the omitted category is “out of the labor force”). There were two urban households and 29 rural households with nonpositive income, and these households were excluded for the instrumental variables regressions; the results

of OLS regressions excluding these households were very similar to those reported in table 3.

G. Pakistan

The Pakistani data are taken from the 1991 Living Standards Survey collected by the World Bank. There were 4,971 households surveyed. Of these, 251 were excluded because of missing information on variables used in our analysis. The survey asks households about expenditure on purchased items and consumption of home-produced goods. The income side of the survey is quite complex and asks households detailed questions about farming activities, the operation of nonagricultural enterprises, and the participation of household members in the labor market.

The major variables are total expenditure on all goods and services and food expenditure. They include both cash purchases and the values of home-produced items. The measurement of food expenditure is complicated by the fact that many households purchase food in some parts of the year and consume home-produced items in others. For each purchased item, households were asked to specify the number of months of the year in which they usually purchased the item and the average amount spent per month in months in which purchases were made. These figures were used to calculate a figure for total food purchases. The value of home-produced food was computed in a similar way. The cash income figure used in the instrumental variables estimates is equal to total cash revenues from family-operated nonfarm enterprises, plus labor earnings, pensions, and transfers. Because of the complexity of the farm section of the survey, we did not attempt to measure cash income from farming. Although our measure of cash income is admittedly imprecise, it serves our purpose: because it is constructed solely from items that are not used to construct expenditure and is correlated with total expenditure, it serves as a valid instrument for expenditure. Adults are defined to be those aged 16 and older. The survey provides weights and geographical clustering information. There are 300 clusters, with an average of 16.7 households per cluster.

The regressions included controls for the ratio to household size of the numbers of males and females in age groups 0–5, 6–15, 16–59, and 60 and older. The gender differences in these categories were suppressed for the regressions in table 4. Other controls include dummies for the province of residence, urban/rural status, and whether the household is a farm household. Also included were variables that measured the fraction of adults who worked as agricultural employees, nonagricultural employees, agricultural workers on family farms, and workers in nonagricultural family enterprises. In contrast to the other surveys we use, household members who work do not report a “major” activity. A single household member who does more than one type of activity (i.e., works as an agricultural employee and on a family enterprise) will be counted more than once when the fractions described above are computed. The cash income measure is positive and non-missing for only 3,402 households, and the instrumental variables regres-

sions are estimated for this sample. However, OLS estimates based on this sample are similar to those reported in table 3.

H. South Africa

The South African data are taken from the Living Standards Survey collected between August and December 1993 by the World Bank and the South African Labour and Development Research Unit (SALDRU) of the University of Cape Town. The survey covers 8,850 households. We use observations only on households classified as "African" and exclude those identified as "white," "colored," and "Indian." Our final sample is 6,476, after exclusion of 37 households with no adult members or with missing expenditure information. Like the other surveys we use, the Living Standards Survey asks households detailed questions about expenditure on food and other goods, consumption of home-produced items, and income obtained from a variety of sources.

The major variables are defined as follows. Expenditure is expenditure on all goods and services and was constructed by the World Bank/SALDRU team. Food expenditure covers both purchased food and home-produced food and was originally reported on either a quantity or expenditure basis, with the former converted to the latter using prices collected in an associated community questionnaire. There is fairly elaborate allowance for imputations, including employer subsidies of food and transportation. Reporting periods vary from two to 52 weeks, depending on the commodity. Cash income estimates are calculated from the total income figures estimated by World Bank/SALDRU by deleting all imputations, including much of farm income for subsistence farmers. Adults are those aged 16 or older. The survey provides both sampling weights and geographic clustering information. There are 301 clusters with an average of 21.5 households per cluster.

The regressions included controls for the ratio to household size of the numbers of males and females in the age categories 0–5, 6–15, 16–59, and 60 and older. The gender differences in these categories were suppressed for the regressions in table 4. Other controls include dummies for the province of residence and rural/urban/metropolitan status. We also included variables that measure the fraction of adults who were engaged as "regular" wage workers and casual wage workers and were self-employed in agriculture and in nonagricultural activities. The cash income measure was positive for only 6,162 observations, and only these observations were used for the instrumental variables regressions. However, the OLS estimates that use this smaller sample were very similar to those reported in table 3.

References

- Barten, Anton P. "Family Composition, Prices and Expenditure Patterns." In *Econometric Analysis for National Economic Planning*, edited by Peter E. Hart, Gordon Mills, and John K. Whitaker. London: Butterworths, 1964.

- Deaton, Angus S. *The Analysis of Household Surveys: A Microeconomic Approach to Development Policy*. Baltimore: Johns Hopkins Univ. Press, 1997.
- Deaton, Angus S., and Muellbauer, John. *Economics and Consumer Behavior*. New York: Cambridge Univ. Press, 1980.
- Engel, Ernst. "Die Lebenskosten belgischer Arbeiter-Familien früher und jetzt." *Internat. Statist. Inst. Bull.* 9, no. 1 (1895): 1-124.
- Engle, Robert F.; Granger, C. W. J.; Rice, John; and Weiss, Andrew. "Semiparametric Estimates of the Relation between Weather and Electricity Sales." *J. American Statist. Assoc.* 81 (June 1986): 310-20.
- Estes, Eugena, and Honoré, Bo. "Partial Regression Using One Nearest Neighbor." Working paper. Princeton, N.J.: Princeton Univ., 1995.
- Eubank, R. L., and Speckman, Paul. "Curve Fitting by Polynomial-Trigonometric Regression." *Biometrika* 77 (March 1990): 1-9.
- Fan, Jianqing. "Design-Adaptive Nonparametric Regression." *J. American Statist. Assoc.* 87 (December 1992): 998-1004.
- Fraker, Thomas M.; Martini, Alberto P.; and Ohls, James C. "The Effect of Food Stamp Cashout on Food Expenditures: An Assessment of the Findings from Four Demonstrations." *J. Human Resources* 30 (Fall 1995): 633-49.
- Lanjouw, Peter, and Ravallion, Martin. "Poverty and Household Size." *Econ. J.* 105 (November 1995): 1415-34.
- Moffitt, Robert A. "Estimating the Value of an In-Kind Transfer: The Case of Food Stamps." *Econometrica* 57 (March 1989): 385-409.
- Prais, Sigbert J., and Houthakker, Hendrik S. *The Analysis of Family Budgets*. Cambridge: Cambridge Univ. Press, 1955.
- Robin, Jean-Marc. "Price of Goods and Consumer Behavior: The Case of Food." Manuscript. Paris: Inst. National de la Recherche Agronomique, 1985.
- Robinson, Peter M. "Root-N-Consistent Semiparametric Regression." *Econometrica* 56 (July 1988): 931-54.
- Rothbarth, Erwin. "Note on a Method of Determining Equivalent Income for Families of Different Composition." App. 4 in *War-Time Pattern of Saving and Spending*, by Charles Madge. Occasional Paper no. 4. Cambridge: Cambridge Univ. Press (for Nat. Inst. Econ. and Soc. Res.), 1943.