University of Wisconsin-Madison

Institute for Research on Poverty

Discussion Papers

Jacques van der Gaag

ON MEASURING THE COST OF CHILDREN

DP #663-81

On Measuring the Cost of Children

Jacques van der Gaag

The research on which this paper is based is part of the Child Support Project of the Institute for Research on Poverty, University of Wisconsin, Madison. I thank Yves Balcer, Sheldon Danziger, Irv Garfinkel and Eugene Smolensky for various comments on an earlier draft and for numerous useful discussions on the topic. However, I have no reason to believe that any one of them would like to be held responsible for any mistakes or omissions; thus the usual disclaimer applies.

ABSTRACT

The primary objectives of this paper are to give a precise definition of the costs of a child and to compare various approaches in the literature to estimating these costs.

Two alternative approaches are discussed at length. The first goes back to the work of Engel (1895) on household consumption patterns. The second is based on methods to directly estimate Individual Welfare Functions of Income (van Praag, 1968). Both approaches produce well-defined measures of levels of well-being. These measures are adopted to define the cost of a child. We will show that both approaches can be viewed as stemming from the same general methodological framework. A short discussion of methodologies that do not fit this framework is included. It is concluded that there exists in the literature no consensus on the exact value of the cost of a child. The estimates for a first child range from 0 percent to 42 percent of a household's yearly income. However, the century-long development of this topic in the economic literature has produced a precise definition of the cost of a child that—once adopted—will enable researchers to improve their estimates and narrow down this disappointingly wide range.

On Measuring the Cost of Children

1. INTRODUCTION

Estimating the "true" cost of a child clearly is of more than just academic interest. Eligibility for social welfare programs is often defined on the basis of a household's income adjusted for family size. These adjustments implicitly reflect the cost of children. Virtually all tax systems in the western world allow tax deductions for families with children. In many countries families receive child allowances from the government to alleviate the financial burden of raising children. And for child support payments in divorce cases the cost of a child has to be estimated in one way or another. A recent estimate of the cost of raising a child from 0 to 18 years of age is some $85,000 (Espenshade, 1980). The discussions following the publication of such a number often suggest that it is accepted by the researchers dealing with this topic as a reasonable ball-park estimate. A closer look at the literature, however, shows a different picture. In order to produce an estimate of the cost of a child, many problems have to be solved, not the least of them the definition of this cost. Various solutions to these problems yield a large variety of cost estimates, showing a surprising lack of consensus about the "true" number.

This paper sketches the development of the topic in the economic literature. Consequently the paper has the character of survey, but I do not claim completeness. I focus on a number of papers that I consider to have contributed significantly to the development of a clear definition of the cost of a child. This definition is presented, and a selection of estimates based on it are discussed. I also briefly discuss some efforts that do not fit the general framework.

The economic literature on the cost of children is embodied in the literature on the demand for consumption goods. A clear historical development can be observed from the seminal work of Engel at the end of the nineteenth century via Sydenstricker and King (1921) and Barten (1964) to Muellbauer's (1974) work on household equivalence scales, on which most of the recent work in this area is based.

In the next section I shall discuss this development in some detail, and derive a formal definition of the cost of a child. The rest of the paper will focus on the consequences of this definition for the actual measurement of the cost of children.

Section 3 will review the various approaches to the measurement problems found in the literature on household equivalence scales. In section 4 I shall try to construct estimates of the cost of a child, based on the estimates found in the literature, my subjective evaluation of these estimates and our own empirical work in this area.

In section 5 I digress a little to discuss some basic methodological problems that have been pervasive in the literature throughout the century-long development of the topic.

Section 6 discusses some approaches that do not fit this framework.

In section 7 I shall summarize this paper by specifying a set of questions that should be addressed in order to obtain reliable, well-defined estimates of the cost of a child.

2. HOUSEHOLD EQUIVALENCE SCALES AND THE DEFINITION OF THE COST OF A CHILD

The literature that considers the determination of scales to adjust income (or consumption) levels of families of different composition in order to make them "equally well off," goes back to the work of E. Engel (1895).

Engel postulated an expenditure function of the following form as an appropriate method of incorporating household composition effects into the analyses of consumer demand:

$$\frac{q_i}{m} = q_i\left(\frac{C}{m}\right), \qquad (1)$$

where q_i is expenditures on good i by a given household,

C is income (equal to total consumption expenditures),

m is a measure of household equivalence.

In the simplest case, m equals family size, and equation (1) says that the per capita expenditures on good i by a given household are a function of per capita household income. As we will see later on, the measure of household equivalence, m, contains the information we need to calculate the cost of a child. In the per capita case, children are considered "equally expensive" as adults, and it is implicitly assumed that there are no economies of scale in household production. In this case, for example, food expenditures, housing expenditures, transportation expenditures, etc., are considered to be twice as high for a couple with two children as for a childless couple in order to reach the same welfare level.

It goes without saying that in order to obtain more realistic estimates of household equivalence, a more realistic representation of household consumption behavior is needed.

As early as 1921, Sydenstricker and King criticized the approach represented in equation (1) as too restrictive, because each good consumed is rescaled by the same amount m (and consequently income is rescaled by m). In 1955, Prais and Houthakker rediscovered this objection to the original Engel approach. They reformulated equation (1) as:

$$\frac{q_i}{m_i} = q_i\left(\frac{C}{m}\right) \tag{2}$$

where m_i is a commodity-specific weighting factor and m is a weighted sum of the m_i's. In this formulation, it is possible to allow for large economies of scale for, say, housing and hardly any for, say, clothing.

Barten (1964) showed how this approach can be incorporated in a utility framework, thus giving a formal base to the notion of "household equivalence." Two households are considered to be equally well off if both have the same level of utility (economic well-being). This level of utility can be inferred from differences in consumption patterns between households of different composition. Muellbauer (1974) showed how this approach can yield household equivalence scales, the formal definition of which will be the base of our definition of the cost of a child.

Following these developments, most of the literature on household equivalence scales now starts with the familiar assumption in economics, that households maximize a utility function under a budget constraint. More specifically, a utility function U is postulated with arguments q_1, q_2, ..., q_K, the quantities of goods 1, 2, ..., K.

Households are assumed to choose that bundle of consumption goods that maximizes their utility, given their total income.

Thus households face the following choice problem:

$$\text{maximize } U = U(q_1, q_2, \ldots, q_K), \tag{3}$$

$$\text{under } \sum_i^K p_i q_i = C, \text{ the budget constraint,} \tag{4}$$

with q_i, quantities consumed of good i,

p_i, price of good i,

C, "income" (= total consumption expenditures).

This constrained maximization results in a set of demand equations of the form:

$$q_i = q_i(p_1, \ldots, p_K, C) \quad i = 1, \ldots, K. \tag{5}$$

Thus the demand for good i is a function of all prices, and income.

Since it is likely that large households are worse off with a given bundle of consumption goods than small households, a more realistic representation of the utility function is:

$$U = U(q_1, \ldots, q_K; h), \tag{6}$$

i.e., the utility level reached with a given bundle of goods is conditional upon household composition h. The corresponding set of demand equations now reads:

$$q_i = q_i(p_1, \ldots p_K, C; h) \quad i = 1, \ldots, K, \tag{7}$$

saying that the demand for goods is a function of all prices and income, given household composition. Note that equations (1) and (2) from an earlier time, are, assuming that prices are constant across households, special cases of equation 7.

For our present purposes it is useful to restate the maximization problem: instead of solving for the maximum utility level that can be reached with a given income, given household composition and prices, we can solve for the minimum expenditure level (income) needed to reach a given utility level, given household characteristics and prices.

More formally, the original problem:

$$\underset{q_1, \ldots, q_K}{\text{maximize}}\ U = U(q_1, \ldots q_K; h) \text{ subject to } \sum_{i=1}^{K} p_i q_i = C^\circ$$

with C° a prespecified income level,

can be stated as

$$\underset{q_1, \ldots, q_K}{\text{minimize}}\ C = \sum_{i=1}^{K} p_i q_i \text{ subject to } U(q_1, \ldots q_K; h) = U^\circ$$

with U° a prespecified utility level.

Both approaches are equivalent and yield a set of demand equations as given in (7). The solution of the second optimization problem, the minimum income, C_{min} needed to reach U°, is obviously a function of prices, the utility level chosen, and household characteristics, so

$$C_{min} = C(p_1, \ldots p_K, U^\circ, h). \tag{8}$$

The function $C(\cdot)$ is the cost function.

This cost function allows us to calculate "household equivalence scales." These scales tell us the factor by which a given income level of a family with characteristics h_2 should be multiplied to become "equivalent to" (i.e., to yield the same utility level, to give the same level of economic well-being as) the income of a family with characteristics h_1.

Let U° be a given level of utility for the reference family, h_1, and let $C_1 = C(p_1, \ldots, p_K, U^\circ, h_1)$ be the minimum cost to reach that level, and C_2 be similarly defined for a household with characteristics h_2.

Then,

$$\frac{C_2}{C_1} = \frac{C(p_1, \ldots, p_K, U^\circ, h_2)}{C(p_1, \ldots, p_K, U^\circ, h_1)} \tag{9}$$

is the ratio with which to calculate the income of the first household (h_1) to get the equivalent income for the second household (h_2). We will now adopt the concept of the cost function to define the cost of a child.

Let h_1 represent a couple without children, and h_2 a couple with one child. The cost of one child is defined as:

$$C_2 - C_1 = C(p_1, \ldots, p_K, U^\circ, h_2) - C(p_1, \ldots, p_K, U^\circ, h_1), \tag{10}$$

i.e., the cost of one child is equal to the difference in the incomes of a one-child household and a childless household ($C_2 - C_1$) that is needed to reach the same given level of economic well-being (U°).

Since the analysis of the cost of a child is usually based on cross-sectional data, for which it is assumed that all households face the same prices, equation (10) can be rewritten as:

$$C_2 - C_1 = C(U^\circ, h_2) - C(U^\circ, h_1). \tag{11}$$

This definition of the cost of a child, or, more generally, of the difference in cost[1] of reaching the level of economic well-being ($U°$) between a household with characteristics h_2 and one with characteristics h_1, has two important consequences.

First, from (11) it is clear that the cost of a child is a relative concept, i.e., it depends on prior choice of the utility level $U°$. Thus "the cost of one child" will generally be different for "the rich" (with a high income-utility level) than for "the poor;" this is intuitively plausible. It does imply, however, that we cannot specify the cost of a child if we do not first specify the utility level (income level) to which this cost refers.

Second, the cost of a child depends on "the difference" between h_2 and h_1. It is not known, a priori, how this difference should be measured. In our example, h_2 was a couple with one child, h_1 a childless couple. But in general we have to address the following questions: does the sex of the child matter, does the age of the child matter, is the cost of a second child the same as the cost of a first child, do age and sex differences between the first two children matter, and what about subsequent children? Are other household characteristics relevant, like the employment status of the spouse?

Before addressing these questions, we will answer the question of how to measure levels of economic well-being in the next section.

3. METHODS OF MEASURING ECONOMIC WELL-BEING AND THEIR RELATION TO HOUSEHOLD EQUIVALENCE SCALES

3.1 Indirect Measures of Economic Well-Being

As we have seen in the previous section, the theory starts with the

---

[1]In economic jargon: the compensating variation.

concept of a utility function to be maximized under a budget constraint. The result is a set of demand equations, explaining the consumption of goods and services as a function of prices,[1] income, and household characteristics.

In applied work, we work the other way around. We observe the consumption of different market bundles by households with different incomes and of different family composition. From this consumption behavior we infer differences in economic well-being (utility).

One of the best-known examples of this approach again goes back to the work of Engel. One of Engel's observations was that the proportion of income spent on food declines as income rises ("Engel's law"). A similar observation was made with respect to family size: large households spend a larger proportion of their income on food than small households.

This suggests that the food share can be used as a measure of well-being. It is often assumed that two households are equally well off if they spend the same proportion of their income on food. Once this measure of well-offness is accepted, the measurement of the cost of a child is straightforward, as in the following example:

Assume we observe two households. One is childless, has an income of $10,000 and spends 25% of that income on food. The other has one child, the same income, and spends 30% on food. According to our food-based definition of economic well-being, the childless couple is "richer." The question is: How much additional income is needed to make the second household equally well off? We can answer this question by observing

---

[1]Since in all that follows I assume prices to be constant across households, I shall ignore price differences from now on.

one-child families at different income levels. Suppose we find that the average one-child family spends 25% of its income on food at an income level of $12,000. We conclude that the cost of a child is $2,000. (Alternatively, say that $12,000 is the equivalent income for a one-child household, as compared to a childless couple with $10,000.) Equivalent scales based on this principle are widely used (BLS worker budgets, see U.S. Department of Health, Education and Welfare, 1977, Espenshade, 1973, Dubnoff, 1979b) and I shall discuss them later.

The advantages of using the food share as a measure of economic well-being are clear, and the measure is easy to calculate. The amount of information needed is limited. It is based on some intuitive notion of basic needs: large families "need" more food than small families. Finally, it is based on Engel's early observations of household consumption behavior, and "Engel's law," based on a small nineteenth-century survey among blue-collar workers in Belgium, have been repeatedly confirmed in later work.

The problems with this measure are equally clear. Food is an obvious necessity, and for poor households it is plausible to assume that they first spend part of their income on food, before deciding how to spend the rest on other commodities. But an equally plausible assumption can be made with respect to housing and, maybe to a somewhat lesser extent, clothing. Especially in a rich society, in which basic food needs can virtually always be met, the focus on food seems somewhat arbitrary, and is too restrictive. Furthermore, the observation that food shares decline as income rises, and rise if family size increases, does not imply that equal food shares represent equal welfare levels (Friedman, 1952).

As Watts (1977) has shown, the food share approach can easily be extended to include other commodities. The "iso-prop" index he developed is based on the assumption that households spending equal proportions of their income on "basic necessities" (food, housing, clothing, and transportation) are equally well off. Measuring the cost of a child based on this definition of economic well-being is accomplished in the same as when the food share is used (see Seneca and Taussig, 1971).

Though Watts's approach is an obvious improvement over the measure based on food alone, a number of problems remain. The choice of the goods to be called "basic necessities" is again somewhat arbitrary. Moreover, the intuitive appeal (households first have to spend part of their income on basic necessities; the more they have left thereafter, the better-off they are) becomes less convincing in a rich society[1] where the concept of "necessities" is less anchored in a notion of physical needs than in some notion of socially acceptable minimum living conditions (which might include such "unnecessaries" as a color TV, theater tickets and, say, one two-week vacation per year).

Pushing the idea that households first spend part of their income to satisfy some "basic needs" to their limit, it seems reasonable to assume that households first spend part of their income on some specific minimum level of a large number of goods, before they decide how to spend the remainder of their income. It turns out that, conceived of in this broader way, we do not have to specify in advance which goods or services belong in the category "basic necessities".

---

[1]Watts developed his measure explicitly to refer to a poverty line concept similar to the so-called Orshansky poverty line.

This is an approach that has a base in economic theory, that identifies these unnecessities, as I shall now show.

Probably the best known utility function is the Stone-Geary function, which has the following form:[1]

$$U = \sum_{i=1}^{K} \beta_i \ln(q_i - \gamma_i), \tag{12}$$

where β_i and γ_i are parameters, $\Sigma\beta_i = 1$ and $\gamma_i < q_i$, $i = 1, K$. Thus the utility level derived is a weighted sum of the logarithm of the goods consumed, insofar as the quantity of each good consumed exceeds some minimum level γ_i. The β_i's are the relative weights.

Maximization of this utility function under the budget constraint yields the following set of demand equations:

$$q_i = \gamma_i + \beta_i(C - \sum_{j=1}^{K}\gamma_j) \qquad i=1, \ldots, K. \tag{13}$$

That is, households first buy the quantities γ_i for each commodity, then they spend the rest of their income, $(C - \sum_{i=1}^{K}\gamma_i)$, in the proportions β_i.

Thus the proportion of income spent on the minimum levels γ_i for all K commodities can be adopted as a measure of economic well-being: the smaller this proportion, the "better off" you are.[2] If we can identify the

---

[1]Ignoring houshold characteristics for the moment.

[2]The minimum levels γ_i are sometimes referred to as "subsistence levels," again giving the impression that those levels have some base in physical needs. A more recent label is "committed consumption." However, this interpretation is not necessary, and, in fact, breaks down if one or more of the γ's appear to be negative. We nevertheless adopt this interpretation for expositional convenience.

minimum levels γ_i for households of different composition, we can again obtain a measure of the cost of a child in a straightforward way.[1]

As stated in Section 2, Barten (1964) has shown how household characteristics can be incorporated in a utility-maximizing framework. Following his approach, (10) is rewritten as

$$U = \sum_{i=1}^{K} \beta_i \ln \left(\frac{q_i}{m_i} - \gamma_i\right), \tag{14}$$

with m_i a commodity-specific weighting factor. This factor is a function of household characteristics

$$m_i = 1 + \delta_i' h \qquad i=1, \ldots, K, \tag{15}$$

with h a vector of household characteristics.

The demand equations resulting from the maximization of (14) are:

$$q_i = \lambda_i m_i + \beta_i \left(C - \sum_j \gamma_j m_j\right) \qquad i=1, \ldots, K. \tag{16}$$

Thus households first buy the quantity $\gamma_i m_i$ of each commodity i, and then spend the rest of their income $(C - \sum_{i=1}^{K} \gamma_i m_i)$ in proportions β_i.

Let us take as our base household a childless couple, setting the weights m_i for this couple equal to 1.0. If the h in equation (15) represents the number of children, the weights for a couple with one child are equal to $m_i = 1 + \delta_i$, $i = 1, \ldots, K$. Thus, if the childless couple spends $\Sigma\gamma_i$ on "committed consumption," the one-child household spends $\Sigma\gamma_i m_i = \Sigma\gamma_i + \Sigma\gamma_i\delta_i$. The difference, $\Sigma\gamma_i\delta_i$, equals the cost of one

---

[1]Goldberger (1967) has shown that the proportion of income spent on "committed consumption" is directly related to Frisch's formal measure of economic well-being, "money flexibility" (the income elasticity of the marginal utility of income).

child, as defined in section 2. Note that this cost is based on the income level needed just to buy the minimum levels γ_i. We cannot say anything yet about the cost of a child at higher income levels.[1]

In order to be able to do that, the cost-function as introduced in section 2 can be employed, as Muellbauer (1974) has shown.

The cost function corresponding to the Stone-Geary utility function reads:

$$C(U^\circ, h_1) = \Sigma\gamma_i m_i^1 + \exp[U^\circ - \Sigma\beta_i \log\beta_i + \Sigma\beta_i \log m_i^1], \qquad (17)$$

where m_i^1, $i = 1, \ldots, K$, is the commodity-specific weight for a household with characteristics h_1. So, again, if h_1 is the childless couple and h_2 is the couple with one child, the cost of one child, at utility level U°, equals $C(U^\circ, h_2) - C(U^\circ, h_1)$, where $C(\cdot)$ is specified as in (17).

Thus, we have observed a rather straightforward development from measuring economic well-being on the basis of the proportion of income spent on food, through Watts's iso-prop index, which is based on necessities, and finally to total expenditures for minimum consumption levels of all goods. All three measures (food share, necessity share, committed consumption on all goods) depend on the size of the household, which gives us the information needed to obtain an estimate of the cost of a child.

---

[1]In this framework we will never be able to say anything about households with incomes <u>below</u> $\Sigma\gamma_i$, since in that case $q_i < \gamma_i$ for at least one good, and the utility function is not defined. In fact the utility level is not defined for $C = \Sigma\gamma_i$ either, but the interpretation of the γ_i's as committed consumption does make the interpretation of $\Sigma\gamma_i\delta_i$ as the cost of a child at this income level plausible. Alternatively one might think of it as the approximation of the cost of a child for an income level slightly above $C = \Sigma\gamma_i$.

The third measure (committed consumption) gives us the tie with the economic theory of utility maximization discussed in the previous section, if we adopt a specific form of the utility function (the Stone-Geary function). And, as shown in section 2, this utility-maximizing framework gives us a formal definition of the cost of a child, by employing the cost function. In principle the parameters of this cost function can be estimated from the parameters in the demand equations. Thus, in all cases discussed so far, observed household consumption patterns provide the information needed to obtain estimates of the cost of a child.

A slight modification of this approach is due to Henderson (1950a, 1950b). So far I have implicitly assumed that the household is the decision-making unit, and have concentrated on the "utility level of the household." Henderson concentrates on the welfare of the parents only, and obtains the cost of a child by observing the consumption of adults in households of differing size. Since breaking up household consumption data between children's consumption and adult consumption is hard (if not impossible) for most consumer goods, Henderson concentrates on alcoholic beverages, tobacco, and adult clothing. Two pairs of adults are considered to be equally well off if they consume equal shares of their income on these "adult consumption" items. Thus, by observing adult consumption differences among households of different composition, the cost of children can be calculated along the lines previously sketched for the food proportion and the other indirect techniques.

In the next subsection I shall discuss an alternative approach to estimating the cost of a child. Instead of <u>indirectly</u> obtaining the parameters of the cost function from observed consumption differences between households of different composition, a method is proposed by which to measure the cost function directly.

3.2 Direct Measures of Economic Well-Being

From the definition of the cost of a child it is clear that all that is needed is a dollar amount that equates the welfare level of a couple with a child to a prespecified welfare level of a childless couple. Thus, in general, we try to measure how much it takes, under various circumstances, to reach a given welfare level.

One straightforward way of obtaining this measure is to conduct a survey in which people are asked to say what it takes.

The best-known example of this approach can be found in the Gallup polls: respondents are asked to specify the minimum amount of money required by a family of four to "get along." Clearly the answer to this question is the "cost" of reaching a prespecified welfare level, "to get along," for a household of given composition. If the same question were to be asked for different household sizes, the cost of children could be directly estimated by analyzing the systematic differences in the answers obtained. Rainwater (1974) presents such an analysis based on the Boston Social Standards Survey.

One obvious shortcoming of this approach is that respondents are asked to judge the economic well-being of a hypothetical household. The respondents' own current situation may differ both in economic well-being and in household composition.

Goedhart et al. (1977) asked the following question: Living where you do now, and meeting expenses you consider necessary, what would be the very smallest income you (and your household) would need to make ends meet? This way of posing a question that has to do with welfare levels refers directly to the respondent's own circumstances. Not surprisingly, the answer one gets varies systematically with those circumstances.

More specifically, Goedhart et al. show that the answer to this question is a function of the income level of the respondent and his family size. Thus it contains all the information we need to calculate the cost of a child.[1]

Goedhart's analysis is part of a larger body of literature on the individual welfare function of income, developed by van Praag (1968, 1971). Instead of asking for a level of income that corresponds with only one welfare level ("get along," "make ends meet") respondents are asked to state the income level they associate with six or seven welfare levels, ranging from feeling "terrible" to feeling "delighted."

The answers to this income evaluation question are transformed into a so called individual welfare function of income (WFI), which gives a utility level (on a 0-1 scale) associated with each income level. Two typical WFI's are given in Figure 1.

---

[1]Goedhart et al. use this question to obtain a poverty line for households of different size.

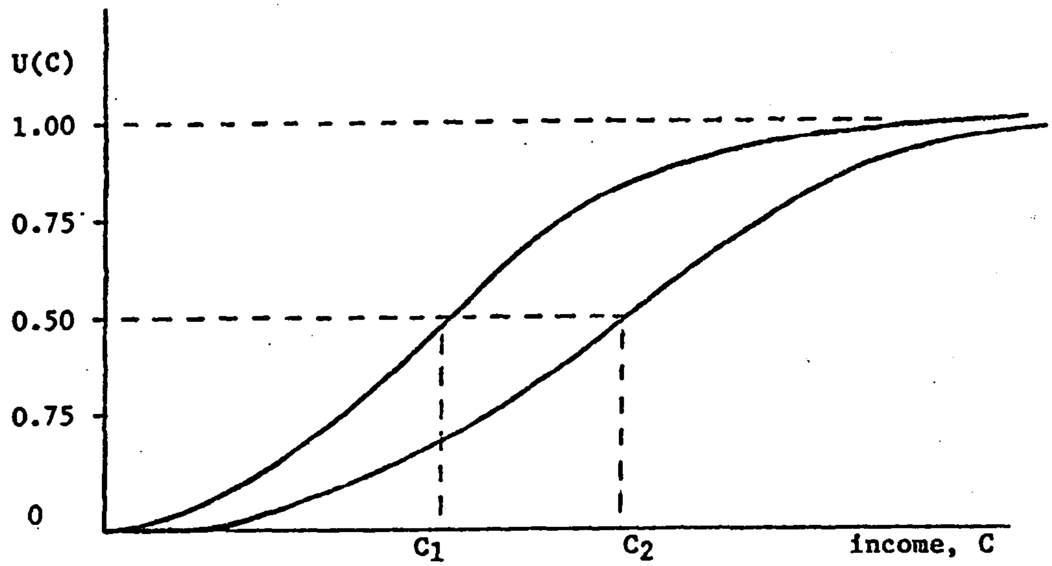

Figure 1. Two individual welfare functions of income.

Given the answers to the income evaluation question, a WFI can be estimated for each household. Again the answers show systematic variation with family size, thus enabling the researcher to calculate the cost of a child. For instance, if graph (1) in Figure 1 refers to a childless couple and graph 2 to a couple with one child, the cost of a child at household utility level .50 is calculated as $C_2 - C_1$.

Using this approach, Kapteyn and van Praag (1976) derive a full set of family equivalence scales. In the next section I shall discuss their results, together with results obtained by the various other methods of measuring welfare that we have discussed in this section.

4. ESTIMATES OF THE COST OF A CHILD

Virtually all studies from which I shall obtain estimates of the cost of a child deal primarily with the estimation of complete "family equivalence scales." To ease the exposition I shall start this section by discussing estimates of the cost of a first child. Thus the question is: How much income does a couple with one child need, to obtain the same (prespecified) level of economic well-being as a childless couple?

The second part of this section deals with the cost associated with subsequent children.

4.1 The Cost of a First Child

Table 1 presents estimates of the cost of a first child, obtained by first converting the equivalence scales found in the literature so that a childless couple is the reference household. Then the cost of a first child is obtained by multiplying the reference household's income level by this equivalence scale.

As the table readily reveals, there is not much consensus about the numbers. The percentage increase of income needed to compensate a couple for having a first child runs from 0% to 42%. There seems to be no systematic relation between the outcome and the technique used. Henderson, basing his estimates on "adult consumption," gets numbers between 17% and 22%, depending on income level. The finding that the percentage increases with income, however, is counter-intuitive, though Seneca and Taussig show a similar result for lower levels of income.

The "proportion spent on food" method yields an increase anywhere between 0% and 42%. The "necessity" method lies between 7% and 40%. The various direct approaches yield between 13% and 30%, while the constant utility approach results in 0% to 35% increases. Even if we disregard the outliers (the three zeros and the four numbers over 35%), we are left with multiples evenly distributed between 6% and 35%.

Table 1

Estimated Cost of a First Child

Author	Technique	Income[a] Level	Income Increase Needed (%)	Cost of[b] First Child	Remarks
Henderson (1950)	adult cons.	600 pence/ week, 1937	17%		
		1000	16		
		1600	22		
Espenshade (1973)	food prop.	$ 7,360	40	$2,944	These are
		11,657	32	3,730	averages over
		18,223	26	4,738	an 18-year
Dubnoff (1979a)	food prop.	*	30		period
Dubnoff (1979b)	direct	8,522	28	2,526	
Seneca, Taussig (1971)	food prop.	5,544	1	55	
		12,312	42	5,171	
		32,160	29	9,326	
	necessities	5,544	7	388	
		13,608	40	5,443	
		34,560	26	8,986	
Goedhart et al. (1977)	direct	5,220	13	691	
Kapteyn, v. Praag (1976)	direct	*	19		child age: 2
		*	14		22
Muellbauer (1977)	constant U	£20/week, 1975	16		child age: <5
	40		9		<5
	100		0		<5
	20		30		>5
	40		22		>5
	100		13		>5
McClements (1977)	constant U	£27.50/ week,1972	8		child age: 0-1
			22		8-10
			35		16-18
van der Gaag, Smolensky (1981)	constant U	$11,239	0	0	child age: <6
		8		899	6-11
		12		1,349	12-17
		34		3,821	18+
BLS (U.S. DHEW, 1977)	food prop	*	37		

*not dependent on income level.

[a]These income levels refer to childless couples. All amounts are 1979 dollars, unless otherwise stated.

[b]The cost of a child is defined as the additional income needed if one child is added to a childless couple. The additional income will keep the household at the same level of economic well-being as it was before the addition of a child.

Some of this variation is explainable. With the two exceptions mentioned above, the percentage increase needed declines with income. Looking at the income levels around $12,000, we find 32% (Espenshade) and 0%-35% (van der Gaag, Smolensky).[1] The latter result depends on the age of the child, another major source of the variation in the results. All estimates show that the cost goes up with the age of the child, except Kapteyn and van Praag.

Espenshade's result (32%) is the average cost of a child over an 18-year period. The van der Gaag-Smolensky result is consistent with the assumption that the cost increases with approximately two percentage points each year, yielding an average cost of 18%. Note finally that Muellbauer's results are consistent (for the midrange of income) with a 2% to 3% increase per year of age, implying an average cost of 18%-27%, and McClements's results also fall in this range, showing 22% for the "average nine-year-old."

Thus, this effort to reduce the range of costs yields the result that between 18% and 32% additional income is needed for a couple with about $12,000 income, but if I were obliged to give an estimate on the basis of the information given above, I would say that the "true value" of the cost of a first child is between 20% and 30% of a childless couple's income. An obvious point estimate would be 25%.

Thus a couple with a yearly income of $12,000 needs, on average, $3,000 more per year to enjoy the same level of economic well-being with one child. But we would like to emphasize the large variance in the

---

[1]Ignoring here the British pence and pounds contributions. Their "mid-range" estimates run from 9%-22%. We also ignore the outliers.

estimates. Other observers might easily reach a different point estimate.

A final word on the effect of the income level. It can be shown that for the constant utility approach (Barten, 1964, Muellbauer, 1977) the percentage of compensating income decreases if the income (utility) level increases. Muellbauer's results are in accordance with that, and so are the results of van der Gaag and Smolensky. However, the latter show that the equivalence scale is virtually constant over a large income range. Only at very high incomes does the scale become flatter. Adopting this last result for the next subsection, I shall proceed under the assumption that the equivalence scales are approximately constant over the relevant income range.

4.2 The Cost of Second and Subsequent Children

Though I did derive at a point (gu) estimate for the cost of a first child in the previous subsection, I was able to do so only after extensive manipulation of the data. Unfortunately, the consensus about the cost of subsequent children is even weaker than that for the first child. Table 2 shows the increase, in percentages, in income needed to compensate a household if one or more children are added.

Espenshade estimates the cost of the second child to be about half that of the first child. Some of the estimates of Henderson (low income), Seneca and Taussig (high income), Kapteyn and van Praag, and Muellbauer (young children) are in agreement with this result. However, Dubnoff's results imply approximately constant cost per child, as do the BLS scales. Some of the results of Seneca and Taussig (for low income

Table 2

Estimated Income Increases, in Percentages, Needed to Compensate for Increasing Family Size (the reference household is a childless couple)

Author	Income Level[a]	1 Child	2	3	4	5
Henderson (1950)	600 pence/week, 1937	17	8			
	1000	16	13			
	1600	22	16			
Espenshade (1973)	$7,360	40	18	17		
	11,657	32	15	14		
	18,223	26	13	12		
Dubnoff (1979a)	*	30	29	25	26	22
Dubnoff (1979b)	8,522	28	26	22	21	
Seneca Taussig (1971)	5,544	1	29	35	42	53
	12,312	42	34	27	31	26
	32,160	29	19	11	4	8
	5,544	7	34	21	22	27
	13,608	40	15	16	18	20
	34,560	26	13	10	11	11
Goedhart et al. (1977)	5,220	13	11	8	8	7
Kapteyn, van Praag (1976)	*, young children	19	9	8	3	
	*, older children	14	7	6	3	
Muellbauer (1977)	£20/week, young children	16	9			
	40/week, young children	9	4			
	100/week, young children	0	−4			
	£20/week, older children	30	27			
	40/week, older children	22	20			
	100/week, older children	13	11			
van der Gaag, Smolensky (1981)	$11,239 children < 6	0	9	7	5	4
	" 6–11	8	10	10	5	4
	" 12–17	12	9	7	5	4
	" 18+	34	10	8	5	4
BLS (U.S. DHEW, 1977)	*	37	31	32	34	

*not dependent on income level.

Note: The reference household is a childless couple.

[a]These income levels refer to childless couples. All amounts are 1979 dollars, unless otherwise stated.

levels) and van der Gaag and Smolensky (young children) find the second child to be more expensive than the first.

The results for the third child are quite similar across studies: the third child is approximately as costly as the second. The cost for subsequent children decreases fast, according to Goedhart, Kapteyn and van Praag, and van der Gaag and Smolensky, but does not change much according to Dubnoff or the BLS. The Seneca and Taussig estimates are erratic in this respect.

It should be noted that where the age of the children is taken into account, we generally find the second child to be roughly half as expensive as the first. (This includes Espenshade's results, giving the average cost over an 18-year period.) The cost of the third child is roughly equal to the cost of the second. And, in addition, we find the cost decreases rapidly after the third child.

Since, as we saw above, the age of the child is an important factor in determining its cost, it is likely that where the age of the child is ignored, the effect of the number of children is contaminated by the age effect. This could explain to some extent the deviant results of Dubnoff and Seneca and Taussig. It does not explain the results of van der Gaag and Smolensky for young children, however.

Thus, if any general result can be derived from Table 2, one could argue that the second child costs about half as much as the first, the third costs the same as the second, and the subsequent children are about half as expensive as the second and third. If we tie this to our previous (gu) estimate of 25% for the first child, we obtain Table 3. Again I should emphasize that, because of the large variance in the estimates, Table 3 could only be obtained after an unsatisfactory data manipulation.

Table 3

Average Cost of Children

Number of Children (1)	Cost of Subsequent Child (%) (2)	Equivalence Scale (3)	Income (4)	Cost of Subsequent Child ($) (5)	Cost of All Children ($) (6)	% income "shared with children" (7)
0	-	100	$12,000	-	-	-
1	25%	125	15,000	$3,000	$3,000	20%
2	12.5	137.5	16,500	1,500	4,500	27
3	12.5	150	18,000	1,500	6,000	33
4	6.25	156.25	18,750	750	6,750	36
5	6.25	162	19,440	750	7,500	39

Note: The reference household, a childless couple = 100; the reference income is $12,000.

Column 2 of Table 3 shows the percentage of compensating income needed for the additional child, to keep the household "as well off as a childless couple with $12,000." Column (4) shows equivalent income levels from which the dollar cost of a child can be obtained (column 5). Column 7 shows this dollar cost as a percentage of the equivalent income. Thus this column can be interpreted as the percentage of income that the parents use for their children.

Up to now, some readers might have the impression that, in spite of the variety in the techniques used, the theory is well established and some consensus can be reached from the empirical results. In the next section we will partially discomfort those readers.

5. SOME PROBLEMS RELATED TO MEASURING THE COST OF CHILDREN

As we have seen in the previous sections, the concept of the cost of a child can be considerably clarified if we start with the assumption that households maximize a utility function given their resources. The welfare comparisons based on the utility levels reached provide the information for measuring the cost of a child. Up to now, I have assumed that the only arguments in the utility function are consumption goods. Consequently, the appropriate budget constraint refers to total consumption expenditures. In what follows I shall discuss a number of other arguments that should enter the utility function in order to make meaningful welfare comparisons.

5.1 How Many Children?

The most controversial additional source of utility is the children themselves. Why do couples decide to have children, if this results in a drop in their welfare level? Or, to paraphrase Deaton and Muellbauer (1980a), where do children come from, from the storks? Clearly, if children themselves are a source of utility, the cost of a child cannot be obtained from analyzing constrained consumption behavior alone.

Pollak and Wales (1979) show that in order to make unconditional welfare comparisons (i.e., comparisons in which children are not treated as given, but are treated as arguments in a household utility function) we need information of the following kind: what would you prefer, a household with two children and an income of $12,000, or a household with three children and a $15,000 income. But of course, many households prefer three children over two, even without the income adjustment.

Thus, Pollak and Wales argue, the household's preferences should not only be defined over consumption goods but should include the number of children. More generally, it should include all household characteristics that can reasonably be assumed to be an object of choice. Living alone, in conventional households or in extended households, for example, can be arguments in the utility function. Welfare comparisons should be based on these unconditional utility functions.

If one accepts this argument (and it is especially appealing in a world where having children is more and more the result of a conscious choice rather than of an unpredictable stork) one may wonder whether the approach to measuring the cost of a child, as sketched in the previous section, is valid.

I argue that it is. The conditional welfare comparison does yield a compensating amount of income, ignoring the utility derived from children. As such it provides us with a clear measure of the cost of a child (the gross cost, if one wants, since the benefits of having the child are ignored).

However, as the argument of Pollak and Wales makes clear, it is not obvious that a household should be financially compensated for this cost. While conditional welfare comparisons can form the base for estimating the cost of a child, it is questionable whether the cost estimates thus obtained should be used to correct household income in defining "equals." For many policy purposes[1] it might be reasonable just to accept that some people prefer large households over small ones, and consequently decide that they are better off with, say, two children than with one, even without any compensation in income.[2]

5.2 The Time Cost of Raising Children

One of the main aspects of the cost of children is the time input of the parents. Ignoring the cost of time will result in highly questionable estimates of the cost of a child, but in all the literature reviewed above, the parent's time input is not expressly included.

Conceptually the time cost can easily be incorporated in the present framework. We make the following simplifying assumptions: (1) only the

---

[1]Tax schedules, eligibility for social programs, transfer payments, etc.

[2]Note that in this entire discussion, we use the notion of "household welfare," and not welfare of the parents, welfare of the first child, etc. In making transfer payments, for instance, one might argue that the family size should be taken into account, in order to raise the welfare level of the children, who were not involved in deciding the family size, to an acceptable level.

wife's time is relevant;[1] (2) time can be used for two purposes only: working in the market place and leisure. "Leisure" includes all activities outside the market place, such as housework, child-raising and real leisure itself.

The standard model of utility-maximizing households can now be reformulated as:

$$\underset{q_1,\ldots,q_K,L}{\text{maximize}} \quad U = U(q_1, \ldots, q_K, L; h), \tag{18}$$

$$\text{subject to} \quad \sum_{i=1}^{K} p_i q_i + wL = y_o + wT$$

where: L is the leisure of the spouse,

w is the wage rate of the spouse,

y_o is household income not earned by the spouse (i.e. nonlabor income plus the husband's earnings),

T is total time available to the spouse.

Thus households maximize a utility function with, as arguments, the consumption goods $q_1, \ldots, q_K$ and the leisure of the spouse, L, measured in, say, hours per year. As before, this utility function is conditional upon household characteristics, h. The budget constraint says that expenses on goods, $\Sigma p_i q_i$, plus "expenses on leisure," wL, cannot exceed "full income."

The cost of leisure is equal to the opportunity cost that the spouse accepts for not working L hours in the market place. Thus the cost of an

---

[1]These assumptions are made to simplify the exposition. No sexist bias regarding parental duties is intended.

hour of leisure is equal to the spouse's hourly wage rate. Full income is defined as total household income in the event that the spouse works a total of T hours.

The maximization again yields a set of demand equations for goods, plus one for leisure. The demand-for-leisure equation can, of course, be transformed in a labor supply function by using H = T - L, where H is hours of work in the market place. This labor supply function specifies the number of hours the wife spends in the market place, given prices (including her wage rate, w), nonearned income y_o, and the characteristics of the household, h.

For example, we can specify the utility function in equation (18) as:

$$U = \Sigma\beta_i \ln(\frac{q_i}{m_i} - \gamma_i) + \beta_L \ln(\frac{L}{m_L} - \gamma_L) \tag{19}$$

Thus, equation (19) is an augmented form of equation (14) presented in Section 3. The augmentation specifies the contribution of the spouse's leisure to the household's utility. The weighting factor m_L is again a function of household characteristics, h, indicating that leisure in a household with children is different from leisure in a household without children. If h represents the number of children, and m_i, i=1, ..., K, L, is specified as in equation (15), "committed consumption plus leisure" for a couple with one child equals $\Sigma\gamma_i m_i + \gamma_L m_L$, and for a childless couple $\Sigma\gamma_i + \gamma_L$.

The total cost (money and time) of one child is thus:

$$(\Sigma\gamma_i m_i + \gamma_L m_L) - (\Sigma\gamma_i + \gamma_L) =$$

$$(\Sigma\gamma_i(1 + \delta_i) + \gamma_L(1 + \delta_L)) - (\Sigma\gamma_i + \gamma_L) = \Sigma\gamma_i\delta_i + \gamma_L\delta_L,$$

the additional cost of goods plus the additional cost of leisure.

As before, all parameters can in principle be estimated from consumption equations, now augmented by a labor supply equation for the spouse. From the literature on female labor supply, we know that the presence of children has a large impact on female labor participation. This suggests that the time cost of children is indeed considerable.

The model can easily be extended to include the husband's time input. The same model can be used to analyze the total cost of a child in single-parent households. It is likely that in these households the time cost is especially large.

It is important to emphasize two restrictions of this model. First it is assumed that the wife can choose the optimum number of hours she wants to spend in the market place. In practice there might be many restrictions in the labor market that are not accounted for in the model. Only full-time jobs might be available, or no jobs at all. Incorporation of these restrictions into the model would severely complicate the analyses.

Secondly, we distinguished between time in the market place and time at home ("leisure") only. Much would be gained if we could split the time at home into time related to child-raising and other time. Obviously this would produce serious measurement problems (which part of "time for cooking" is related to the child?).

Turchi (1975, Chapter 3) analyzed the hours a wife spends on "housework" (as distinguished from market work and leisure). He finds, for instance, that the first child adds about 835 hours per year[1] to the

[1]This is an average over 22 years.

time spent on housework, or more than 16 hours a week. At a wage rate of, say, $5.00 per hour, this means that the wife invests per year more than $4000 worth of her time in raising the child. A considerable amount, indeed, as compared to the money cost figures presented in Section 4.

This shows, as stated before, that measures of the cost of a child that ignore the time inputs of the parents will be seriously biased downwards. Though we did find it reasonable, for some purposes, to ignore children as arguments in the household utility function, it is much harder to find examples where ignoring leisure as a factor relevant for a household's welfare can be theoretically justified.[1] Consequently, in estimating the cost of a child, the parent's time input in raising the child should be included.

The next subsection will deal with yet another factor that is generally ignored in measuring the cost of a child.

5.3 Savings and Other Problems Related to Estimating the Cost of a Child

Up to now we have assumed that in any given period, say one year, households spend their entire income on goods and services. Thus we implicitly assumed that income equals total consumption expenditures. Savings or dissavings were ignored. However, it is quite likely that the presence of children will have an impact on a household's savings behavior, and this change should be taken into account when estimating the cost of a child.

---

[1]Household equivalence scales, of course, are generally used to make welfare comparisons based on household income only. But income is mainly chosen as the appropriate welfare indicator for practical reasons, not because it is theoretically the best measure.

The simplest way to do so is to treat savings as just another good, and proceed as sketched in Section 3. But this approach ignores the importance of the length of the period over which the households are observed. Households may anticipate having children, and save in advance to meet the higher cost of obtaining a specific level of living. Once the children are there, we will observe dissavings. However, this might change if the parents start to save to pay for, say, a future college education for the children.

A policy of treating savings as just another good and basing our estimates on observations during one year only cannot take these complications into account. For instance, using the spending and saving behavior of a childless couple as the reference point will yield biased results if some of these childless couples have modified their behavior in anticipation of having children. A lifetime welfare comparison seems to be in order, but for all practical purposes, comparing welfare levels of households with and without children over more than just a few years seems infeasible. Nevertheless, in making these short-term comparisons, we should be aware of the possible bias in the results arising from the short length of time the households are observed.

I end this section with three technical notes on the estimation of the cost of a child.

First, in the absence of all information on prices, it is not possible to obtain constant-utility household equivalence scales from demand equations without using additional information (Muellbauer, 1975, Cramer, 1969).

The cause of this fundamental identification problem is relatively simple: since households supposedly spend their entire budget, information on how much they spend on good K is redundant once we know how much they spend on the first K-1 goods. So we only have K-1 independent pieces of information, when we estimate K demand equations. This is not enough to derive the K good-specific weighting factors m_i that appear in the cost function.

It can be shown that this problem does not occur if we have observations on households that vary in family size and face different prices. However, most data sets that are rich in household composition data are poor in price variation and vice versa. Solutions to the problem include the use of additional information (e.g., nutritional requirements for households of different size) or the adoption of additional assumptions to the consumer demand theory (compare Kakwani, 1977, and van der Gaag and Smolensky, 1981). The results one obtains depend, of course, on the particular solution chosen.

A second problem relates to the particular form of the utility function (and consequently the cost function) chosen. The Stone-Geary function leads to the familiar Linear Expenditure System as in equation (13). However, the implications of this system are quite restrictive. We chose the system for expositional convenience and because it is one of the most widely used systems in empirical work.

The choice of the system is merely an empirical question, thus it is preferable to start with as general a specification as possible. Recent work on the Almost Ideal Demand System (Deaton and Muellbauer, 1980) might turn out to be important in this respect.

The final problem is somewhat related: How do we incorporate household composition variables in a demand system? Barten chose a particular form, known as scaling--compare equation (14). Recently Pollak and

Wales (1980) compared scaling with various other approaches giving insight into the sensitivity of the results to the method chosen. The scaling method performed quite satisfactorily, but it is not unlikely that future research in this area, using more general and more flexible forms of demand equations to represent household consumption behavior will improve the estimates of the cost of a child.

As we have shown, a large body of the literature on household composition and consumer behavior can be embodied in one general framework. This framework--utility maximization--provides a convenient way to define and estimate the cost of a child. One of the most widely used estimates of the cost of children, however, does not fit in this framework: the Orshansky Poverty Line Equivalence Scale. We will discuss the base of this scale and some related approaches in the next section.

6. ALTERNATIVE APPROACHES TO MEASURING THE COST OF A CHILD

The current official U.S. poverty measure consists of income cutoffs for 124 different family sizes and types. The cutoffs vary by the age of the household head, age of the children, sex of the household head, and total family size. These cutoffs are obtained as follows.

Food costs for families of different age-sex composition (family types) were derived by "costing out" food needs based on nutritional requirements (for men, women, and children of different ages) suggested by the National Research Council; this allows consideration of age and sex differences in measuring need. A multiplier was then applied to the food requirements to reflect nonfood needs (U.S. DHEW, 1976, p. 78).

Thus, where the equivalence scales discussed in the previous sections were all based on observed consumer behavior, the equivalence scales implicit in the U.S. poverty line (the "Orshansky scale") are primarily based on differences in nutritional requirement.

The U.S. Department of Agriculture thrifty-food plan was adopted to derive these differences across families of different composition, after which the total cost of the corresponding market basket was obtained. Thus, if the diet of a couple with a child is 13% more costly than that of a childless couple, the corresponding income equivalence scale is 113. Thus this scale implicitly assumes that goods-specific equivalence scales are the same for all goods. It ignores the possibility of differences in economies of scale between food needs and, say, housing needs. As such it is equivalent to the original Engel approach (see equation [1]), with the equivalence measure derived from food needs. Consequently, the criticisms of the Engel approach, starting with Sydenstricker and King (1921), hold for the U.S. poverty line equivalence scale.[1]

Yet another approach to estimating the cost of a child is due to Turchi (1975). For a given income group he estimates equations of the form:[2]

$$Exp_i = \alpha_i + \beta_i \text{ Child} \tag{20}$$

where Exp_i = expenditures on good i, and

Child = 1 if there is one child in the household,

= 0 otherwise.

---

[1]The multiplier used to transform the food cost into an income level stems from the 1955 Food Consumption Survey in which the average food expenditure-income ratio was found to be 1:3. However, for some family types a slightly different ratio was used, thus making the Orshansky scale a combination of a nutrition need and food proportion scale.

[2]Turchi uses a more general form; equation (20) is stated so as to compare expenditures on a given good between a childless couple and a couple with one child.

Using Turchi's interpretation, α_i equals the expenditures on good i for a childless couple, β_i the additional expenditures for a child. But it is misleading to interpret the sum of the β_i's over all goods as the cost of a child. Households can spend a given income only once, so if, because of a child, they spend more on some goods, they have to spend less on some other goods. Consequently, if all goods are taken into account, the β_i's will sum to zero.

Finally, there is the very extensive analysis of the cost of a child by Lindert (1978, 1980). Lindert defines the relative cost of a child as the ratio of all inputs into the child (goods, services and other family members' time) relative to the inputs in all activities that would have been enjoyed in the absence of the extra child. This relative cost-notion is defined over the entire planning horizon of the parents. In theory this approach could be fitted into a lifetime utility framework, in which households maximize their well-being by deciding how to distribute their available lifetime resources among raising children and other "enjoyable activities." In applied work, however, this lifetime approach has numerous practical problems, among them measuring the input of goods and time into raising children and defining and measuring the counterfactual: "inputs in all other activities that would have been enjoyed in absence of the extra child." Lindert's results are based on many ad hoc assumptions regarding these inputs, and therefore lack the theoretical base of the approaches discussed in Section 4.[1]

---

[1]The short discussion of the approaches of Lindert and Turchi cannot do sufficient justice to their work, especially since both authors are among the very few that explicitly tried to estimate the money and time cost of a child. However since the approaches of both authors are less theoretically justifiable, and based on a less precise definition of the cost of a child than the one presented in section 4, I only briefly mention them. The interested reader, however, is referred to Turchi (1975), Lindert (1978) and various chapters of Easterlin (1980).

7. SUMMARY AND CONCLUSION

In this paper we sketched the development of that part of the economic literature that implicitly or explicitly deals with the estimation of the cost of children. As we have shown, the development of the estimation of the cost of a child parallels the development of the analyses of consumer demand. In all cases which we referred to as the indirect techniques, differences in consumption patterns formed the basis for measuring differences in levels of economic well-being, and measured differences in economic well-being enabled us to estimate the cost of a child.

A relatively new and particularly attractive technique does not rely on observed consumer behavior, but directly obtains the necessary information through survey techniques.

Both the direct and the indirect approach fit into the same theoretical framework, since they try to answer questions of the following kind:

> How much does it take for a household of given composition to reach a given, prespecified level of well-being? How much more or less does it take for a household of different composition to reach the same level?

As we have seen, defining the cost of a child by using the answers to these questions highlights the problems that we must solve in order to obtain an estimate of the cost of one child.

First, we have to specify the basis of the "level of well-being;" in other words, we have to decide which arguments should enter the utility function. Generally, this will depend on the purpose of the analysis. If our goal is to make unconditional welfare comparisons, i.e., if we consider household composition as one of the results of the household's choices in maximizing its own welfare, the number of children should be one of the arguments in the utility function.

If, however, the purpose of our analysis allows us to treat household composition as given, we can make conditional welfare comparisons. The utility function is then conditional upon the number of children, but the number of children is not a choice variable. The utility function is defined over a set of goods and services only. "Leisure" should be included if one intends to estimate the "full cost" of a child in both money and time. Savings should be included too, but, as we have seen, this causes severe problems related to the period over which information is collected.

Various ad hoc appproaches have been used to define equal levels of well-being. Among them are equal proportions of income spent on food and on necessities respectively. Sometimes the cost of a child is estimated without much reference to the level of well-being specified. However, in principle the cost of a child in a rich household will differ from the cost of a child in a poor household.

The various approaches employed seem to yield different results, but no <u>systematic</u> relationship between the techniques used and the results obtained could be detected. It would be worthwhile to apply the various techniques, including the direct approach, to the same data set, in order to assess their relative merit.

Second, in the estimation of the cost of a child, we have to decide which household characteristics to include in the analysis, i.e., what factors constitute "different household composition." Obviously, the age of the child is an important factor. But so is the "age of the household," since households in various stages of the life cycle show different consumption and savings behavior. We have, therefore, to address several questions: Are we going to ask whether the cost of a child depends on the parents' ages, or on their employment status, or on the number of adults in the household? Or on other household characteristics?

If more than one child is involved we have to decide whether sex, birth order, and differences between ages are important. Obviously there is no theoretical answer to these questions. Again it is likely that the purpose of the analysis will suggest the factors that are relevant. If our aim is to set standards for a minimum welfare level of a child our choice might be different from what it might be if we try to define "equals" for public policy purposes. In addition, the data available will usually automatically reduce our set of choices.

Once we know how to measure "levels of well-being" and "differences 'n household composition" we have all the ingredients needed to measure the cost of a child (of given age, sex, with both parents present, etc.) at various levels of well-being.

As we have seen, the estimates currently available in the literature have a large variance. I derived a point estimate, ignoring time cost, of $3000 for the first child in a family with an income of $12,000. But I emphasize again that I could only obtain this estimate after what I consider to be excessive manipulation of the data. In spite of the century-long development of literature on the topic, little consensus as to the "correct answer" has been obtained. But at least there is consensus about the correct way of posing the question.

The approaches that do not fit within the utility-maximizing framework are all based on questionable or imprecise definitions of the cost of a child. It seems unlikely that further work in these directions will lead to better estimates of a child's cost. It is more likely that improvements of the cost estimates will come from further developments in consumer demand analysis and in the direct measurement of individual welfare functions. At the moment, however, we should be aware that the estimates of the cost of a child presented in the literature are based on a large number of varying assumptions. Therefore, in evaluating these estimates of the cost of a child, it is important to get precise answers to at least the following questions: what cost? which child? and whose?

REFERENCES

Barten, A. P. 1964. Family composition, prices, and expenditure patterns. In P. E. Hart, G. Mills and J. K. Whitaker (eds.), Econometric analysis for national economic planning. London: Butterworth.

Cramer, J. S. 1969. Empirical econometrics. Amsterdam: North Holland Publishing Co.

Deaton, A. S. and Muellbauer, J. 1980a. An almost ideal demand system. American Economic Review, 70.

________. 1980b. Economic and consumer demand. Cambridge: Cambridge University Press.

Dubnoff, S. 1979a. A method for estimating economic welfare of American families of any composition, 1860-1909. Institute for Social Research, University of Michigan. Mimeo.

________. 1979b. Experiments in the use of survey data for the measurement of income minima. Center for Survey Research, University of Massachusetts, Boston. Mimeo.

Easterlin, R. A. 1980. Population and economic change in developing countries. Chicago: The University of Chicago Press.

Engel, E. 1895. Die Lebenskosten Belgischer Arbeiter-Familien früher und jetzt. International Statistical Institute Bulletin, 9, 1-74.

Espenshade, T. J. 1973. The cost of children in urban United States. Population Monograph Series, No. 14. Berkeley: University of California.

________. 1980. Population Reference Bureau, Inc. News release, October 3.

Friedman, M. 1952. A method of comparing incomes of families differening in composition. Studies in Income and Wealth, 15. New York: National Bureau of Economic Research. Pp. 9-24.

Goedhart, T., V. Halberstadt, A. Kapteyn, and B. M. S. van Praag. 1977. The poverty line; concept and measurements. Journal of Human Resources, 12, 503-520.

Goldberger, A. S. 1967. Functional form and utility: a review of consumer demand theory. Social Systems Research Institute, University of Wisconsin, Madison. Mimeo.

Henderson, A. M. 1950a. The cost of children, parts I-III. Population Studies, 3, 130-50; 4, 267-98.

________. 1950b. The cost of a family. Review of Economic Studies, 17, 127-148.

Kakwani, N. C. 1977. On the estimation of consumer units scales. Review of Economics and Statistics, 59, 507-510.

Kapteyn, A. and B. M. S. van Praag. 1976. A new approach to the construction of family equivalence scales. European Economic Review, 7, 313-335.

Lindert, P. H. 1978. Fertility and Scarcity in America. Princeton: Princeton University Press.

________. Child costs and economic development. In Easterlin, 1980.

McClements, L. D. 1977. Equivalence scales for children. Journal of Public Economics, 8, 191-210.

Muellbauer, J. 1974. Household composition, Engel curves and welfare comparisons between households: a duality approach. European Economic Review, pp. 103-22.

________. 1975. Identification and consumer unit scales. Econometrica, 43, 807-809.

________. 1977. Testing the Barten model of household composition effects and the cost of children. The Economic Journal, 87, 460-487.

Pollak, R. A. and T. J. Wales. 1979. Welfare comparisons and equivalence scales. American Economic Review, 69, 216-21.

________. 1980. Demographic variables in demand analysis. Department of Economics, University of British Columbia, DP 78-48.

Prais, S. J. and H. S. Houthakker. 1955. The analysis of family budgets. 2nd ed. Cambridge: Cambridge University Press, 1971.

Rainwater, L. 1974. What money buys: inequality and the social meanings of income. New York: Basic Books.

Seneca, J. J. and M. K. Taussig. 1971. Family equivalence scales and personal income tax exemptions for children. Review of Economics and Statistics, 53, 253-262.

Sydenstricker, E. and W. I. King. 1921. The measurement of the relative economic status of families. Quarterly Publication of the American Statistical Association, 17, 842-57.

Turchi, B. A. 1975. The demand for children: The economics of fertility in the United States. Cambridge, Mass.: Ballinger Publishing Company.

U.S. Department of Health, Education and Welfare. 1977. The Measure of Poverty. Technical Paper IV, Bureau of Labor Statistics Family Budgets Program.

van der Gaag, J. and E. Smolensky. 1981. True household equivalence scales and characteristics of the poor in the United States. Institute for Research on Poverty, University of Wisconsin-Madison. Mimeo.

van Praag, B. M. S. 1968. Individual welfare functions and consumer behavior. Amsterdam: North Holland Publishing Company.

________. 1971. The welfare function of income in Belgium: an empirical investigation. European Economic Review, 2, 337-369.

Walker, K. E. and M. E. Woods. 1976. Time use: a measure of household production of family goods and services. Center for the Family of the American Home Economics Association.

Watts, H. W. 1977. The iso-prop index. In M. Moon and E. Smolensky, Improving the measure of economic well-being. New York: Academic Press, 1977.