
Do Husbands and Wives Pool Their Resources?

Evidence from the United Kingdom Child Benefit

Shelly J. Lundberg
Robert A. Pollak
Terence J. Wales

ABSTRACT

Common preference models of family behavior imply income pooling, a restriction on family demand functions such that only the sum of husband's income and wife's income affects the allocation of goods and time. Testing the pooling hypothesis is difficult because most family income sources are not exogenous to the allocations being analyzed. In this paper, we present an alternative test based on a "natural experiment"—a policy change in the United Kingdom that transferred a substantial child allowance to wives in the late 1970s. Using Family Expenditure Survey data, we find strong evidence that a shift toward greater expenditures on women's clothing and children's clothing relative to men's clothing coincided with this income redistribution.

I. Introduction

Traditional models of family behavior assume that family members act as if they are maximizing a single utility function. The common preference ordering may be the outcome of consensus among the family members (Samuelson 1956) or the preferences of a dominant family member (Becker 1981).

Shelly J. Lundberg is a professor of economics at the University of Washington; Robert A. Pollak is a professor of economics at Washington University in St. Louis; and Terence J. Wales is a professor of economics at the University of British Columbia. The authors acknowledge the research assistance of Debra Fogarty and Kelli Schonher and the comments of Jane Friesen, David Lam, Stephen Jenkins, an anonymous referee, and participants in workshops at Stanford University, the Institute for Research on Poverty at the University of Wisconsin-Madison, and the University of Washington. The data used in this article can be obtained beginning in November 1997 through October 2000 from Shelly Lundberg at the Department of Economics, Box 353330, University of Washington, Seattle WA 98195.

[Submitted April 1996; accepted December 1996]

Other models have challenged this "unitary" or "common preference" approach and have attempted to incorporate divergent and conflicting preferences of individual family members into economic analysis. The allocation mechanism in these individual utility models includes cooperative bargaining (Manser and Brown 1980; McElroy and Horney 1981; Lundberg and Pollak 1993), noncooperative bargaining (Kanbur and Haddad 1994; Lundberg and Pollak 1994; Bergstrom 1996), and a generic "collective" approach which avoids specifying a particular model of intrafamily allocation but assumes that family allocations obey a Pareto-efficient sharing rule satisfying certain regularity conditions (Chiappori 1988, 1992).¹

The common preference model implies pooling, a restriction on family demand functions that is both simple and of considerable practical importance. If all income is pooled and then allocated to maximize a single objective function, only total family income will affect family demand. Which family member receives or controls income is irrelevant to the allocation of family resources. Thus, the pooling hypothesis implies the ineffectiveness of targeted transfer policies: transfer policies that attempt to redistribute income to particular family members will be neutralized by the intrafamily allocation process—a form of Ricardian equivalence. In contrast, individual utility models of the household permit the income received or controlled by one family member (for example, the wife) to have a different effect on consumption and time allocation than income received by another (for example, the husband). This is easy to see in cooperative Nash bargaining models in which each spouse's income affects the threat point, and thus the equilibrium allocation.

Tests of the hypothesis that, with prices and wage rates held constant, only the sum of husband's income and wife's income affects the demand for goods and the allocation of time are conceptually simple. However, the implementation of such tests has been hampered by difficulty in finding data on income from sources that are exogenous to the demands and allocations being analyzed. In a cross-section, differences in earned or unearned income of husbands and wives are likely to be correlated with differences in prices (including wage rates) and differences in preferences. In this paper, we present an alternative test of the pooling hypothesis based on a "natural experiment"—a policy change in the United Kingdom that transferred a substantial child allowance to wives in the late 1970s. We find strong evidence that a substantial shift toward relatively greater expenditures on women's goods and children's goods followed this income redistribution.

II. Tests of the Pooling Hypothesis

Empirical tests of the pooling hypothesis have used a variety of data sources, but the measure of "income" received by husbands and wives is either earned income (Phipps and Burton 1992; Bourguignon et al. 1993) or some

1. The implications of nonunitary models of the household are discussed more fully in Lundberg and Pollak (1996).

measure of unearned income. Unearned income can include income from pensions, social security and other public transfers, income from physical assets (rents), income from financial assets, and private transfers (including gifts and bequests). Some studies use total unearned income received by the husband or wife (Thomas 1990) and others use income from particular sources such as property income or transfers (Schultz 1990).

Data on the labor earnings of husbands and wives are readily available, reliably measured, and, most important, are the best measure of the relative control over resources of husbands and wives. However, earnings are clearly endogenous with respect to the household's allocation decisions. In a cross-section, differential effects of husband's earnings and wife's earnings on consumption patterns are consistent with the common preference framework, because households with different ratios of husband's earnings to wife's earnings are likely to face different prices and have different preferences, even with total household income held constant. If we think of each spouse's earnings as the product of hours worked and his or her fixed market wage rate, then the first factor, hours worked, is a standard choice variable in models of household behavior and is determined simultaneously with the consumption patterns the pooling test examines. The assumption that hours worked are exogenously fixed, though sometimes invoked, is not plausible. The second factor, the market wage, measures the price of time for the husband or wife, and enters commodity demands directly in the common preference model. Phipps and Burton find that expenditures on restaurant meals are more elastic with respect to the wife's earnings than the husband's earnings. The bargaining interpretation of this result is that, as the wife's earnings rise relative to the husband's, she gains more influence over the household's spending patterns and that increased expenditures on restaurant meals reflect her preferences. An alternative interpretation, entirely consistent with the common preference model and the pooling hypothesis, is that restaurant expenditures depend upon the cost of substitutes, and that the wife's wage is an important component of the cost of home-prepared meals.

Tests of pooling based on unearned income instead of earnings mitigate but do not completely avoid these problems. Unearned income, unlike earnings, is not contaminated by price effects, but most unearned income sources are not entirely exogenous with respect to household behavior, past or present; furthermore, variations in unearned income over a cross-section are likely to be correlated with other (possibly unobservable) determinants of consumption. The separate components of unearned income are also problematic. For example, property income reflects, to a considerable extent, accumulated savings and is therefore correlated with past labor supply and, through fixed individual effects, current labor supply. Public and private transfers may be responsive to household distress due to unemployment or bad health, and may be related to expenditures through the events that prompted them.² *Unexpected* transfers affect resources controlled by family members but not prices, but unexpected gifts or bequests are likely to be small and sporadic sources of income for most families and therefore, without a specially constructed data set, cannot be used to test the pooling hypothesis.

2. See Schultz (1990) for a discussion of this problem.

In general, tests based on cross-sectional variation in income received by husbands and wives have decisively rejected the pooling hypothesis. There is reason to be concerned, however, that existing tests are not reliable, since they exploit differences in income that are not exogenous with respect to prices, preferences, and other important determinants of consumption behavior. What is required is an experiment in which some husbands or wives are randomly selected to receive an income transfer that is withheld from others or, alternatively, in which there is an exogenous change over time in the income received by husbands or wives.

Lundberg and Pollak (1993) consider, as a thought experiment, a government policy that redistributes income from husbands to wives and investigate the conditions under which it could affect distribution within marriage. More specifically, they consider a lump-sum cash transfer, labeled a "child allowance," that is made to all families. Under one policy regime, the hypothetical transfer is made to the husband in each family; under the alternative regime, it is made to the wife. In the context of a specific cooperative bargaining model of intrafamily distribution, the separate spheres bargaining model, these child allowance policies imply different equilibrium distributions in existing marriages. Child allowances are a convenient exemplar of a cash transfer scheme because they are provided on a universal basis by many countries, and are unlikely to be related to prices or individual preferences.³ A real-world counterpart to this thought experiment would be a change in the child allowance policy of some country from a scheme in which cash transfers are paid to the husband to one in which transfers are paid to the wife. A unitary model of family behavior predicts that such a shift in the child allowance policy would have no effect on consumption patterns, while an individual preference model predicts that consumption may shift toward goods more highly valued, or private to, the spouse receiving the child allowance.⁴ Thus, one could test the pooling hypothesis by comparing expenditure patterns before and after the policy change.

Changes in the U.K. child benefit scheme in the late 1970s provide an opportunity to test the pooling hypothesis in this manner. The universal child benefit, which had consisted primarily of a reduction in the amount withheld for taxes from the father's paycheck, was replaced by a cash payment to the mother. This represented a substantial redistribution of income; by 1980 child benefits amounted to £500 per year for a family with two children, or about 8 percent of average male earnings in the United Kingdom. Under the pooling hypothesis, however, this change in the nominal recipient of the transfer within the household should have no effect on expenditure patterns.

Prior to April 1977 the U.K. child benefit scheme consisted of two separate programs: a taxable Family Allowance payment to the mother; and a Child Tax Allowance, dependent on the age of the child, available to the household as a deduction from income for tax purposes. Over the 1977-79 period these two programs were replaced by a single Child Benefit program which made a nontax-

3. Endogenous fertility is, of course, a problem.

4. In some individual preference models, such a shift will have no effect; this is the case, for example, in a Cournot-Nash voluntary contribution model in which both spouses make strictly positive contributions to the provision of a public good.

Table 1
Family Allowance and Child Benefit Rates (£/week)

	First Child	Second Child	Third and Subsequent Child
1968-74	—	.90	1.00
8 April 1975	—	1.50	1.50
5 April 1977	1.00	1.50	1.50
3 April 1978	2.30	2.30	2.30
13 November 1978	3.00	3.00	3.00
2 April 1979	4.00	4.00	4.00
24 November 1980	4.75	4.75	4.75
23 November 1981	5.25	5.25	5.25

Source: U.K. Department of Health and Social Security (1991), Table G1.01, p. 253.

Note: Child Benefit replaced Family Allowance on 5 April 1977.

able weekly payment to the mother. The Family Allowance program was replaced by the Child Benefit on the 5th of April 1977. Table 1 contains Family Allowance and Child Benefit rates from 1968 through 1981. After 1981 there were periodic increases in benefits, but they were not large enough to maintain their real value; the 1990 rate was £7.25 per week per child.

In addition to replacing the Family Allowance by the Child Benefit, the reform gradually phased out the Child Tax Allowance, a process that was completed by April 1979. From our point of view the most important aspect of these changes was that the mother became the sole direct recipient of the payment. Prior to April 1977 a substantial part of the benefit was in the form of the Child Tax Allowance, which, because it was a deduction from income for tax calculation purposes, generally resulted in an increase in the father's take-home pay. This redistributive aspect of the child support policy change is discussed in Brown (1984, pp. 63-64):

The proposal to abolish Child Tax Allowances, on the other hand, would reduce the take home pay of married men with children and though this would, in family terms, be compensated for by increased Child Benefit, it involved a redistribution of family income from men to women, expected to be more popular with the latter than with the former. Indeed so convinced did some Ministers become that a transfer of income "from the wallet to the purse" at a time of wage restraint would be resented by male workers, that they decided at one point in 1977 to defer the whole child benefit scheme.

The following statement by a Member of Parliament voices a similar sentiment (*U.K. House of Commons Hansard*, 13 May 1975):

. . . far from a new deal for families, it will take money out of the husband's pocket on the Friday and put it into the wife's purse on the following Tues-

Table 2
*Child Tax Allowance, Family Allowance, and Child Benefit
 Combined as a Percentage of Average Male Annual Earnings*

	Two Children		
	Younger than 11 years	Between 11-16	Older than 16
1972-73	7.4	8.5	7.5
1973-74	6.3	7.3	8.2
1975-76	6.4	7.2	7.8
1977-78	6.4	7.0	7.6
1979-80	8.1	8.1	8.1

Source: *U.K. House of Commons Hansard*, 14 January 1980, pp. 641-2.

day. Far from being a child benefit scheme, it looks like being a father dis-benefit scheme.

For many families, the new Child Benefit was a substantial increase over the old Family Allowance and Child Tax Allowance, and was a significant component of total income. Table 2 shows the magnitude of the benefits over the transition period expressed as a fraction of average male annual earnings. The table appears to show that benefit levels did not change much over the period, but the appearance is misleading. First, some of the benefit prior to 1979 took the form of the Child Tax Allowance, and thus was of value only to those with incomes high enough to be taxable. Indeed Family Allowance benefits alone were only about 2.3 percent of average male annual earnings for families with two children in 1972-73 and 2.0 percent in 1973-74. Since low-income families did not stand to gain from the Child Tax Allowance, their benefits as a fraction of average male earnings actually increased from about 2 percent in the early 1970s to about 8 percent by the end of the decade.

Second, even though the level of aggregate benefits was fairly stable over the period for families with incomes high enough to benefit from the Child Tax Allowance, there was still a significant redistribution of these benefits from men to women. If we interpret Child Tax Allowance as primarily paid to the husband and the Child Benefit (and the earlier Family Allowance) to the wife, then from 1972-73 to 1979-80 payments to the husband decreased by about 6 percent of average male earnings while payments to the wife increased by about this same amount in households with two children. To the extent that expenditure patterns depend on who receives the family income, a change of this magnitude may indeed have an effect.

To determine whether a shift in spending patterns did occur we wish to focus on an expenditure category that is likely to be a prime candidate for change. An obvious choice is clothing, because clothing expenditures are "assignable" to

individual household members.⁵ Most consumer expenditure surveys report expenditures on men's, women's, and children's clothing separately. The assignment of men's clothing to the husband and women's clothing to the wife is relatively straightforward for households that have only one adult male and one adult female. Furthermore, given traditional gender roles in the household, it also seems plausible to expect the wife to have a greater interest than the husband in children's clothing. We conjecture, therefore, that expenditures on children's clothing and women's clothing increased relative to expenditures on men's clothing as a result of the policy change that shifted control of child benefits from the husband to the wife.⁶ We now turn to the data and the model we use to test the pooling hypothesis.

III. Data and Model Formulation

In order to test the pooling hypothesis we use U.K. Family Expenditure Survey (FES) data for the period 1973–90. The FES reports separately expenditure patterns for families consisting of one man, one woman, and one child; two children; and three children.⁷ The FES cross-classifies households by income and number of children, and for each cell reports mean expenditures on various consumption categories and the age distribution of the children. These cells are our observations, and we treat cell means (appropriately weighted) as though they were the consumption patterns of representative households. Table 3 presents descriptive statistics for the important variables in our model.

Since the changes in the child benefit were phased in between April 1977 and April 1979, we break the sample into three periods. We use the period 1973–76 to represent the consumption regime before the policy change, drop the intermediate years 1977–79, and use the period 1980–90 to represent the regime after the policy change. The pooling hypothesis implies no significant difference in expenditure patterns before and after the policy change.

The FES provides data on expenditures on women's, men's, and children's clothing. Because there are not separate price indexes for these categories, we cannot estimate a demand system and instead employ single-equation analysis. We do this by selecting an indicator of consumption patterns and relating it to a measure of income or expenditure, family size and age composition, and a set of dummies that distinguishes years before and after the policy change.

We use two indicators of consumption patterns in our tests of the pooling hypothesis. Our first indicator is the ratio of current expenditures on children's

5. Browning et al. (1994) describe a private consumption good as assignable if we can observe individual consumptions. The benefits of clothing consumption, of course, need not be exclusively private.

6. Another category that might well be affected by the policy change is alcohol, but in preliminary estimates we found no evidence of a substantial shift in consumption.

7. Childless couples are excluded because the wide age range of families in these cells (which include the elderly) is likely to result in expenditure patterns noncomparable to those of households containing children. We do not use years before 1973 because the definition of children was different. After 1972 persons were classified as children if they were younger than 18, while in earlier years they were so classified if they were younger than 16.

Table 3
Family Expenditure Survey Data: Descriptive Statistics for Cell Means
(weighted by number of households in cell)

	Mean	Standard Deviation	Minimum	Maximum
One-child families	0.33			
Two-child families	0.51			
Three-child families	0.15			
Number of children younger than age 2	0.23	0.12	0.04	0.68
Number of children age 2 to 4	0.35	0.20	0.05	1.29
Expenditures in £ per week ^a				
Children's clothing	1.52	1.06	0.15	5.74
Women's clothing	1.84	1.32	0.11	7.16
Men's clothing	1.18	0.80	0.08	4.11
Total clothing and footwear	6.22	4.03	0.75	21.86
Total current consumption	53.71	27.97	11.17	144.69
Expenditure ratios, children's clothing/men's clothing			1973-76	1980-90
One-child families			0.97	0.97
Two-child families			1.26	1.63
Three-child families			1.72	2.20
Women's clothing/men's clothing				
One-child families			1.68	1.70
Two-child families			1.33	1.60
Three-child families			1.27	1.77

a. Deflated by retail price index with Jan. 1974 = 100

clothing to current expenditures on men's clothing. We conjecture that expenditures on children's clothing is of more interest to mothers than fathers, and therefore will increase at the expense of the latter as a result of the policy change. Our second indicator of consumption patterns is the ratio of current expenditures on women's clothing to current expenditures on men's clothing, with a change in consumption patterns interpreted as before.⁸ Table 3 shows that the means of both these ratios are higher in the period after the policy change than in the period before the policy change, at least for families with more than one child.

8. Browning et al. also examine consumption of men's and women's clothing and find, using Canadian data, that expenditures on women's clothing rise as the share of total household income earned by the wife rises. Their test of the pooling hypothesis using the total incomes of husbands and wives is, of course, subject to the same criticisms regarding endogeneity of the income shares as the studies discussed above.

If the price indexes for children's (women's) clothing and men's clothing are highly correlated, the ratio of current expenditures will be approximately equal to the ratio of real expenditures. This assumption enables us to proceed without separate price indexes for these categories. One possible source of variation in relative prices was the increase in the value-added tax (VAT) rate from 8 percent to 15 percent in 1979. Since the VAT was levied on adult clothing but not children's clothing, this change in the rate reduced the relative price of children's clothing at about the same time that the child benefit was introduced. This could have led to some substitution toward consumption of children's clothing, but does not present any problems for our test using women's and men's clothing expenditures.⁹

We use two measures of income or expenditure. First, we use a narrow measure of real income, namely the FES aggregate category "clothing and footwear," deflated by the retail price index for this category, available from the *Annual Abstract of Statistics* (U.K. Central Statistical Office). Second, we use a broader measure of real income, the FES aggregate category consisting of all current consumption expenditures, deflated by the corresponding retail price index. Because it is not clear which of these measures is more appropriate, we report both.

Clothing expenditures are likely to be affected by family demographics. Thus, we include in our basic model both the number and age distribution of children. To account for the number of children, we define three 0-1 dummies to reflect families with one, two, and three children, and include the latter two in our regressions; thus, the omitted dummy corresponds to families with one child. Our data also include for each cell the average number of children per household in each of three age categories: children younger than 2 years of age, children at least 2 years of age but younger than 5, and children at least 5 but younger than 18 years. We include the first two measures in the regressions; the third is redundant because the sum of the three age categories is equal to the number of children in the household.

To test for a change in consumption behavior before and after the policy change, we construct a set of dummy variables as follows. We define a variable D to be 0 in the early years and unity in the later years, and multiply D by each of our three dummy variables representing family size. Inclusion of the resulting three dummies, denoted as D_1 , D_2 , and D_3 , allows responses in consumption patterns to vary by family size in a general way. A test of the significance of the null hypothesis of no change in consumption patterns requires an F -test of the joint hypothesis that the coefficients of these three dummies are equal to 0.

In summary, we estimate linear regression equations in which the dependent variable is either the ratio of children's clothing expenditures to men's clothing expenditures, or the ratio of women's clothing expenditures to men's clothing expenditures. The independent variables are the average number of children younger than 2, the average number of children at least 2 but younger than 5, dummy variables for two- and three-child families, a broad or narrow income

9. Except to the extent that small women purchase children's clothing for their own use, record these purchases in the FES diaries at "women's clothing," and increase these purchases in response to the change in the VAT.

Table 4
Ratio of Children's to Men's Clothing Expenditures: Broad Income Measure

Variable	Sample Time Period		
	1973-76 1980-90	1973-76 1980-83	1980-83 1987-90
Children younger than 2	0.76 (1.2)	0.10 (0.1)	0.39 (0.3)
Children 2 to 4	-0.17 (0.4)	-0.47 (0.7)	-0.41 (0.6)
Two-child families	0.40 (2.0)	0.47 (2.0)	0.75 (3.1)
Three-child families	0.96 (4.2)	1.03 (3.7)	1.42 (5.4)
Late period X one-child (D_1)	0.11 (1.0)	0.11 (0.7)	0.12 (0.7)
Late period X two-child (D_2)	0.48 (4.8)	0.39 (3.0)	0.24 (1.7)
Late period X three-child (D_3)	0.55 (3.2)	0.52 (2.5)	0.44 (1.6)
Income/10	-0.12 (4.2)	-0.20 (3.3)	-0.17 (3.5)
Intercept	1.32 (4.8)	1.94 (4.0)	1.83 (3.8)
R^2	0.59	0.52	0.61
H_0	0.0005	0.0024	0.23
Observations	181	118	82

Notes: Figures in parentheses are t values. H_0 denotes the level at which the joint hypothesis $D_1 = D_2 = D_3 = 0$ is significant. "Observations" is the number of cell means used in the regression.

measure, and three dummy variables representing the period after the policy change for different family sizes.¹⁰ We assume that the disturbances are independently normally distributed with constant variances, and estimate the equations by weighted least squares using the square roots of the sample sizes of the cells as weights.¹¹

IV. Results

The results for our basic model of the ratio of children's clothing expenditures to men's clothing expenditures appear in the first columns of Tables 4 and 5. Table 4 uses the broad definition and Table 5 the narrow measure of

10. We have also estimated the model using logarithms of the (nondummy) variables, but do not present the results here since they yield similar conclusions.

11. Beginning in 1979 the data contained an indicator for each cell mean for which the sampling error was 50 percent or more. We excluded these 17 observations from the estimation. Prior to 1979 the data contained an indicator for each cell mean that had "a relatively high sampling error." Because the magnitude of this error is not reported, it is not clear whether these observations should also be excluded. Consequently we have estimated the model both with and without these observations, and found very similar results. The results we report below are based on estimates in which these observations are included.

Table 5
Ratio of Children's to Men's Clothing Expenditures: Narrow Income Measure

Variable	Sample Time Period		
	1973-76 1980-90	1973-76 1980-83	1980-83 1987-90
Children younger than 2	1.17 (1.9)	0.26 (0.3)	1.36 (1.2)
Children 2 to 4	-0.04 (0.1)	-0.29 (0.5)	-0.18 (0.2)
Two-child families	0.35 (1.7)	0.40 (1.7)	0.76 (3.0)
Three-child families	0.92 (3.9)	1.01 (0.37)	1.36 (4.9)
Late period X one-child (D_1)	0.16 (1.3)	0.22 (1.5)	0.08 (0.4)
Late period X two-child (D_2)	0.57 (5.3)	0.55 (4.3)	0.21 (1.3)
Late period X three-child (D_3)	0.57 (3.2)	0.66 (3.0)	0.32 (1.1)
Income/10	-0.56 (3.2)	-1.43 (3.5)	-0.74 (2.3)
Intercept	0.87 (4.0)	1.56 (4.3)	1.10 (2.6)
R^2	0.58	0.52	0.57
H_0	0.0005	0.00003	0.52
Observations	181	118	82

Notes: See Table 4.

income. The number of children in the household has a significant effect on the ratio of children's to men's clothing expenditures. As family size increases so do relative expenditures on children's clothing, with the effect for the three-child households being particularly significant. Real income has a significant negative effect on the clothing expenditure ratio, which implies that the income elasticity of demand for men's clothing is greater than the income elasticity of demand for children's clothing. We also estimated the model with the square of the income variable as well as income itself, but the square term was generally not significantly different from 0.

To test the pooling hypothesis, we focus on the coefficients of the intertemporal dummy variables, the D 's, that are designed to measure a shift in consumption patterns following the policy change in the late 1970s. All three of these coefficients have the expected sign, and those for families with two and three children are individually highly significant. The insignificance of D_1 is surprising; the old Family Allowance program did not make payments to one-child families, so the Child Benefit program increased the direct payments to such households substantially. However, even though D_1 is not itself significantly different from 0, the joint hypothesis that $D_1 = D_2 = D_3 = 0$ is soundly rejected using the usual F -test. The calculated F values are 10.9 and 11.3 for the broad and narrow income measures respectively, both of which are significant at approximately the 0.0005 level. Although the parameter estimates differ slightly depending on which income measure is used, the basic conclusion is the same: the increase in the ratio of

children's to men's clothing expenditures following the policy change is highly significant.

It is comforting to find that the intertemporal dummies are highly significant and have the expected sign. Although this is consistent with the proposition that the policy change caused the change in consumption patterns, it does not prove causation. Indeed there is no way to prove, using these data, that the increased expenditures on children's clothing relative to those on men's clothing was caused by the policy change. Alternative explanations for the apparent regime shift include a change in tastes or fashion that increased the relative demand for children's clothing (for example, the evolution of brand-name consciousness among the school-age set), and a change in relative prices. Fashion change is less plausible when we observe similar changes in relative expenditures on both children's and women's clothing (see below). Furthermore, additional analysis suggests that the shift in consumption patterns is closely associated with the shift in child benefit policy.

To ensure that our results do not depend critically on the data from the late 1980s, well after the policy change took effect, we reestimate the same two equations using data from the periods 1973–76 and 1980–83. That is, we use the same four years of data before the policy change, but instead of using 11 years after the transition, we use only four. It is clear from the results in the second columns of Tables 4 and 5 that our basic conclusions remain valid. In particular, the coefficients of the intertemporal dummies have the expected signs and are highly significant. The joint test that $D_1 = D_2 = D_3 = 0$ is still easily rejected at the 1 percent level for both income measures. Further, the fact that the magnitudes of the coefficients of the dummies do not differ much from those in the first columns of the tables suggests that the change in consumption patterns had pretty well stabilized by 1983.

We next investigate the possibility that our results reflect an underlying trend in consumption patterns unrelated to the policy change. Figure 1a shows the time series of the expenditure ratio, with each observation a weighted average of the cell means in that year. A comparison of the pattern in the raw data with the pre- and post-change means (also plotted on the figure) is strongly suggestive of a one-time shift in consumption behavior, rather than a trend. We confirm this impression by estimating a model that treats the period 1980–90 in the same way we had treated the period 1973–83. That is, we use the observations for 1980–83 as “early” years, omit the observations for 1984–86, and use the observations for 1987–90 as “later” years. Our intertemporal family size dummies, D_1 , D_2 , and D_3 , are now defined to be 0 in the early years and unity in the later years for households with one, two, and three children respectively. All other variables are defined as before.

The results of these models appear in the third columns of Tables 4 and 5. Once again the number of children in the household has a significant positive effect on the ratio of children's to men's clothing expenditures, and the income measures have a significant negative effect. In this case, however, none of the coefficients of the dummy variables intended to reflect a possible change in consumption patterns is significantly different from 0. Indeed the test of the joint hypothesis that all three D 's are equal to 0 can only be rejected at the 23 percent

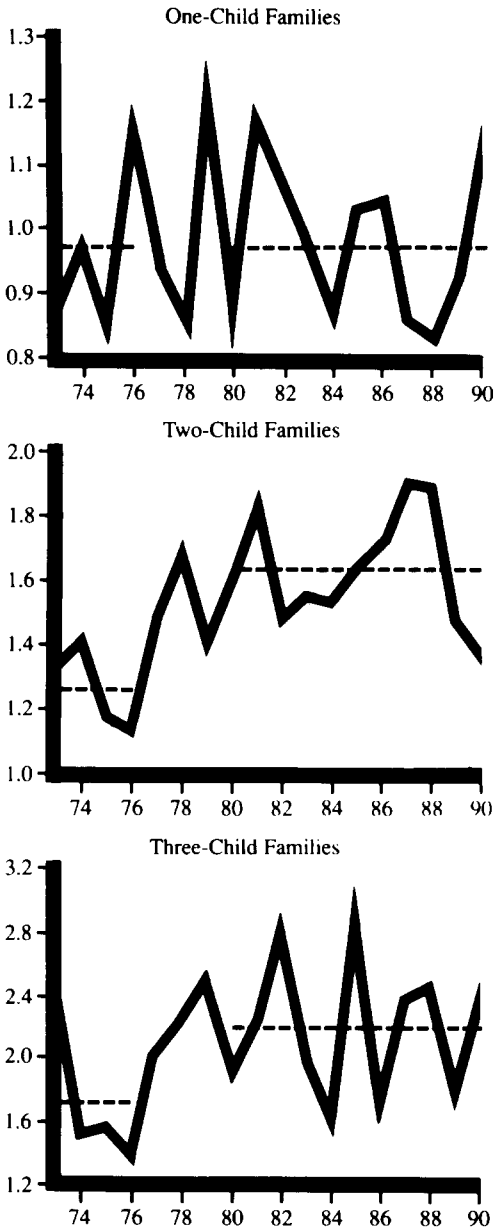


Figure 1a

Ratio of children's clothing expenditure to men's clothing expenditure, 1973-90

Note: Each ratio is calculated from the weighted average of cell expenditure means across income categories. The straight lines are the pre-change (1973-76) and post-change (1980-90) means of the data.

significance level when the broad income measure is included, and at the 52 percent significance level with the narrow definition. This is of course what we would expect in the absence of a policy change and in the absence of an underlying trend in the ratio of children's to men's clothing expenditures unrelated to the policy change we are considering. Hence, the insignificance of these intertemporal dummies adds credibility to our inference that the change in the child benefit caused significant and substantial changes in consumption patterns.¹²

Tables 6 and 7 present the results from an analogous set of models with a different indicator of the consumption pattern as the dependent variable: the ratio of expenditures on women's clothing to expenditures on men's clothing. The models in Table 6 use the broad income measure; those in Table 7 use the narrow measure. The variables reflecting the age distribution of children in the household are not included in the results presented in these tables; when they were included, their effects on the expenditure ratio were always small and insignificant. Increasing family size tends to reduce expenditures on women's clothing relative to men's clothing, but changes in household income have no significant effect on this ratio.

The change in clothing consumption revealed by the intertemporal dummies is remarkably similar to that revealed in the first set of models. Once again, the postpolicy change in consumption is insignificant for one-child families, but positive and significant for two- and three-child families. Significant effects persist, though they are weaker, when the postpolicy period is shortened to four years, but the division of the sample period into early 1980s and late 1980s does not reveal any continuing trend in the expenditure ratio beyond the date of the policy change. Figure 1b, which plots the average expenditure ratio over time, also indicates a one-time shift in relative expenditures on women's clothing. Since the labor force participation of women continued to rise in the United Kingdom during the 1980s, this argues against a labor supply-based interpretation of the increase in relative spending on women's clothing over the late 1970s. The increase in relative expenditures on both children's and women's clothing following a policy change that transferred considerable purchasing power from husbands to wives strongly supports our conjecture that families do not pool their resources.

The effects of the child benefit on the clothing expenditures of two- and three-child families are quite substantial. Are they unreasonably large, given the magnitude of the income redistribution within the family? For the prepolicy change period (1973–76), the mean values for annual expenditures on men's, women's, and children's clothing are £108, £145, and £119 respectively. Holding expendi-

12. Alternative specifications of the model yield essentially identical results. A male unemployment rate included to capture any business-cycle-related changes in clothing expenditures had a small and insignificant effect on both expenditure ratios. Though unemployment in the United Kingdom rose dramatically from the beginning of our sample through 1981, it then fell substantially in the late 1980s, while the expenditure ratios remained at high levels.

The size of the intertemporal effect on expenditures did not vary significantly by income, though some redistribution across income classes should have occurred as a result of the policy change. Including the interim years 1977–79 in the sample period (with a dummy variable representing this period) reduced the standard errors on many of the time-constant variables, but did not change the estimated effect of the child benefit.

Table 6
Ratio of Women's to Men's Clothing Expenditures: Broad Income Measure

Variable	Sample Time Period		
	1973-76 1980-90	1973-76 1980-83	1980-83 1987-90
Two-child families	-0.35 (2.6)	-0.35 (2.5)	-0.27 (1.6)
Three-child families	-0.41 (2.3)	-0.39 (3.7)	-0.16 (0.7)
Late period X one-child (D_1)	0.03 (0.2)	0.15 (1.0)	-0.16 (0.8)
Late period X two-child (D_2)	0.25 (2.4)	0.23 (1.8)	0.14 (0.9)
Late period X three-child (D_3)	0.45 (2.6)	0.38 (1.9)	0.22 (0.7)
Income/10	0.02 (0.9)	-0.004 (0.1)	-0.005 (0.1)
Intercept	1.59 (11.3)	1.69 (9.2)	1.85 (9.0)
R^2	0.40	0.39	0.22
H_0	0.008	0.06	0.55
Observations	181	118	82

Notes: See Table 4.

Table 7
Ratio of Women's to Men's Clothing Expenditures: Narrow Income Measure

Variable	Sample Time Period		
	1973-76 1980-90	1973-76 1980-83	1980-83 1987-90
Two-child families	-0.35 (2.5)	-0.34 (2.5)	-0.26 (1.6)
Three-child families	-0.40 (2.3)	-0.38 (2.2)	-0.16 (0.7)
Late period X one-child (D_1)	0.02 (0.2)	0.17 (1.1)	-0.16 (0.8)
Late period X two-child (D_2)	0.24 (2.2)	0.25 (1.9)	0.15 (1.0)
Late period X three-child (D_3)	0.45 (2.5)	0.40 (1.9)	0.23 (0.8)
Income/10	0.10 (0.7)	-0.12 (0.5)	-0.06 (0.3)
Intercept	1.63 (13.7)	1.73 (11.4)	1.86 (10.6)
R^2	0.40	0.39	0.22
H_0	0.02	0.06	0.54
Observations	181	118	82

Notes: See Table 4.

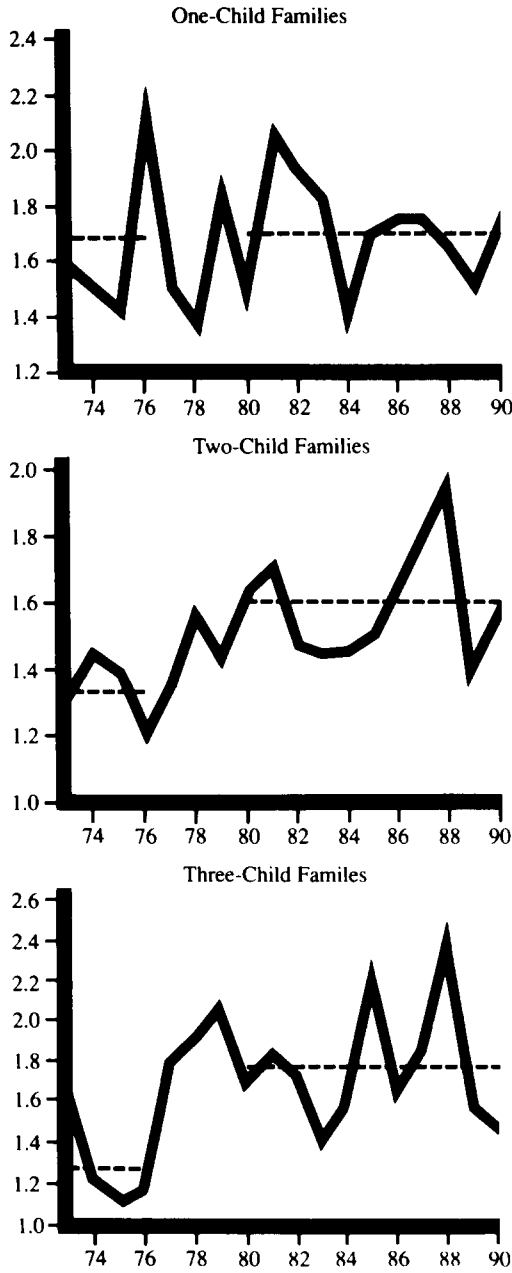


Figure 1b

Ratio of women's clothing expenditure to men's clothing expenditure, 1973–90

See note to Figure 1a.

tures on men's clothing at this level, the coefficients on D_2 imply an increase of approximately £52 in expenditures on children's clothing, and an increase of £27 in expenditures on women's clothing.¹³ Since the annual income redistributed from husbands to wives was, on average, in excess of £400 for a two-child family, these would seem to be reasonable magnitudes. The new Child Benefit was intended to cause a reallocation of resources within families in favor of women and children. Our analysis of clothing expenditures, though limited to a single expenditure category, suggests that it may have succeeded in doing so.

V. Conclusion

Our analysis shows that a substantial increase in spending on women's and children's clothing, relative to men's clothing, followed a policy change in the U.K. child benefit scheme that transferred resources from husbands to wives. Holding constant total family income, the income received by each spouse has substantial and significant effects on family expenditure patterns. These findings are consistent with the notion that children do better when their mothers control a larger fraction of family resources.

Our rejection of income pooling within the family is an important supplement to previous tests of the pooling hypothesis based on cross-sectional data. Previous tests have either treated earnings as exogenous or have treated unearned income or some portion of unearned income (for example, property income, transfers) as exogenous. The change in the child benefit scheme and the redistribution of income from husbands to wives that accompanied it were clearly exogenous with respect to the prices and preferences of individual families. Thus, the natural experiment provided by the change in the U.K. child benefit scheme is important new evidence against common preference models of family behavior and income pooling. The most important implications of this result concern, not the potential effects of alternative child allowance schemes on intrafamily distribution, but the potential effects of increased access to market work and market income for women, in both developed and developing countries.

References

- Becker, Gary. 1981. *A Treatise on the Family*. Cambridge: Harvard University Press, enlarged edition 1991.
- Bergstrom, Theodore. 1996. "A Survey of Theories of the Family." In *Handbook of Population and Family Economics*, ed. Mark R. Rosenzweig and Oded Stark. Amsterdam: North-Holland.
- Bourguignon, Francois, Martin Browning, Pierre-Andre Chiappori, and Valerie Lechene. 1993. "Intra Household Allocation of Consumption: A Model and Some Evidence from French Data." *Annales d'Economie et de Statistique* 29:137-56.
- Brown, Joan C. 1984. *Children in Social Security*. Studies of the Social Security System no. 3. Thetford and London: Thetford Press Ltd.

13. Coefficients are from Column 1 of Tables 4 and 6.

- Browning, Martin, Francois Bourguignon, Pierre-Andre Chiappori, and Valerie Lechene. 1994. "Income and Outcomes: A Structural Model of Intrahousehold Allocation." *Journal of Political Economy* 102(6):1067-96.
- Chiappori, Pierre-Andre. 1988. "Rational Household Labor Supply." *Econometrica* 56(1):63-89.
- . 1992. "Collective Labor Supply and Welfare." *Journal of Political Economy* 100(3):437-67.
- Kanbur, Ravi, and Lawrence Haddad. 1994. "Are Better Off Households More Unequal or Less Unequal: A Bargaining Theoretic Approach to 'Kuznets Effects' at the Micro Level." *Oxford Economic Papers* 46(3):445-58.
- Lundberg, Shelly, and Robert A. Pollak. 1993. "Separate Spheres Bargaining and the Marriage Market." *Journal of Political Economy* 101(6):988-1010.
- . 1994. "Noncooperative Bargaining Models of Marriage." *American Economic Review Papers and Proceedings* 84(2):132-37.
- . 1996. "Bargaining and Distribution in Marriage." *Journal of Economic Perspectives* 10(4):139-58.
- Manser, Marilyn, and Murray Brown. 1980. "Marriage and Household Decision Making: A Bargaining Analysis." *International Economic Review* 21(1):31-44.
- McElroy, Marjorie, and M. J. Horney. 1981. "Nash Bargained Household Decisions." *International Economic Review* 22(2):333-50.
- Phipps, Shelley, and Peter Burton. 1992. "What's Mine is Yours? The Influence of Male and Female Incomes on Patterns of Household Expenditure." Working Paper 92-12, Department of Economics, Dalhousie University, Halifax, Canada.
- Samuelson, Paul A. 1956. "Social Indifference Curves." *Quarterly Journal of Economics* 70:1-22.
- Schultz, T. Paul. 1990. "Testing the Neoclassical Model of Family Labor Supply and Fertility." *Journal of Human Resources* 25(4):599-634.
- Thomas, Duncan. 1990. "Intra-Household Resource Allocation: An Inferential Approach." *Journal of Human Resources* 25(4):635-64.
- U.K. Central Statistical Office. 1973-90. *Annual Abstract of Statistics*. London: HMSO.
- U.K. Department of Employment and Productivity. 1973-90. *Family Expenditure Survey*. London: HMSO.
- U.K. Department of Health and Social Security. 1991. *Social Security Statistics*. London: HMSO.
- U.K. House of Commons. May 1975. *Hansard*. London: HMSO.
- U.K. House of Commons. January 1980. *Hansard*. London: HMSO.