Consumption Inequality and Intra-Household Allocations*

Jeremy Lise
Department of Economics
Queen's University
lisej@qed.econ.queensu.ca

Shannon Seitz
Department of Economics
Queen's University
seitz@post.queensu.ca

December 17, 2003

Abstract

The current literature on consumption inequality treats all adults within the household equally, making the implicit assumption that all consumption inequality is between, not within, households. However, the rise in the share of women's income in the household may have important implications for the distribution of consumption within the household and thus for consumption inequality measured at the individual level. We use an extension of the collective framework of Chiappori to estimate a rule for assigning resources to individual household members. We then construct a measure of individual level inequality by looking at implied changes in intra-household allocations and explore the implications of our framework for the measurement of *individual level*, versus *household level* consumption inequality. Our analysis, which is based on households comprising one or two adults, suggests that the conventional approach of ignoring intrahousehold allocations underestimates cross sectional consumption inequality by one-third and overstates the trend by forty per cent.

*Department of Economics, Dunning Hall, Queen's University, Kingston, Ontario, Canada, K7L 3N6. We would like to thank Garry Barrett, Richard Blundell, Martin Browning and participants at the Zeuthen Workshop at the University of Copenhagen, HEC Montreal, Queen's University, Australia National University and the PhD conference in Business and Economics in Perth, Australia for many helpful comments.

1 Introduction

The primary goal of a large body of research is to measure changes in the distribution of economic welfare across individuals in the population. This is a straightforward exercise in theory, but is quite difficult to implement in practice. The way much of the literature approaches the study of inequality is to equate the well-being of individuals with observed measures of well-being, such as income or earnings. Good measures of income and earnings are typically available for a representative cross-section of the population, allowing for the study of income and earnings inequality. However, in recent years the study of inequality has shifted from the study of income inequality to the study of consumption inequality for many reasons (Cutler and Katz, 1992; Krueger and Perri, 2003; Blundell and Preston, 1998). First, consumption tends to be more closely related to utility than income. Second, individuals have the capacity to smooth consumption over time through borrowing and lending (Deaton, 1996; Crossley and Pendakur, 2002). As such, current income may not be as accurate a measure of well-being as compared to measures of consumption.

Measures of consumption are widely considered an important indicator for measuring the well-being of individuals. Although important, the study of consumption inequality has proved difficult, as it is not possible to get comprehensive measures of individual level consumption for households with more than one member. To overcome this difficulty, most studies of consumption inequality have used adult equivalency scales to convert measures of household consumption into measures of individual consumption. The drawback of this approach is that it implicitly assumes that there is no inequality within the household or that intra-household inequality is time invariant.[1] In particular, the use of adult equivalency scales has as its basis a very

[1]Haddad and Kanbur (1990) and Kanbur and Haddad (1994) consider the importance of intra-household inequality, but only within a given period. The obvious implication of incorporating intra-household inequality is that measured inequality is unambiguously higher than when it is assumed to be zero. They consider the possible re-ranking that could occur when doing between group comparisons (for example urban verses rural). Beaupri (2001) considers the effect on income inequality of the degree of earnings pooling within a household. The emphasis is again on the level

restrictive model of the household in which husbands and wives split consumption equally, regardless of the source of the income in the household.

This criticism levelled against the use of adult equivalency scales may seem to be a subtle point until one considers the following. First, empirical tests of the "unitary" model of the household, where the consumption allocation does not vary with the source of income in the household are routinely rejected in favor of bargaining (Manser and Brown 1980; McElroy and Horney, 1981) or collective models (Chiappori 1988, 1992).[2] Second, there has been a sizable increase in women's wages and labor supply over the last half century. As a result, the share of household earnings provided by the wife has increased substantially. If consumption allocations depend on the source of income and the source of income has changed a lot over time, then adult equivalency scales likely provide an inaccurate picture of the trends in consumption inequality.

Our paper takes a first step at addressing this issue and in doing so makes two contributions to the literature on consumption inequality. First, we construct and estimate a model of intra-household allocations to examine how changes in the source of income in the household translate into changes in individual-level consumption allocations. The model we consider is a version of the collective model introduced in the seminal work of Chiappori (1988, 1992). This model is ideal for the study of consumption inequality as it places very few restrictions on the intra-household allocation process. Estimates of the collective model in the literature indicate that the share of consumption received by members of the household is strongly related to their earnings, or more correctly, their earnings potential (Browning et al., 1994; Browning and Chiappori, 1998; Chiappori et al., 2002; Blundell et al., 2002; Donni, 2001, 2003). The model we estimate allows us to infer the shares of income allocated to each member of the household under relatively weak identification assumptions.

Our second contribution is to use estimates from our model to make inferences

of inequality, not the trends.

[2]It should be noted that the unitary models are in fact less restrictive than the model generating adult equivalency scales. In particular, the unitary model does not predict equal sharing but only that the allocation is independent of the source of income.

regarding how consumption inequality *within* households relates to consumption inequality *across* households. The model is estimated using the Family Expenditure Survey for the UK, which covers 1968 to 2001 and includes data on household expenditures, wages and labor supply. We have three main findings. First, our estimates suggest that measures of consumption inequality that ignore the potential for intra-household inequality may underestimate individual-level inequality by 40%. Second, our parameter estimates suggest that the equal sharing assumption commonly imposed in the literature is valid for households in which the wife has the same earning power as her husband *and* chooses to work identical hours. Third, inequality tends to decrease across cohorts but increase within cohort over time, largely reflecting temporal and life-cycle changes in the gender wage gap and female labor supply. Together, the results highlight the importance of within household inequality as a major determinant of consumption inequality in the economy.

2 Stylized Facts

In this section, we outline the main stylized facts regarding consumption and income inequality in the UK between 1968 and 2001. Many of these findings have been documented in the literature for Canada (Pendakur, 1999), the United States (Johnson and Shipp, 1997 and Krueger and Perri, 2003) and Australia (Barrett, Crossley and Worswick, 2000). The data we use to conduct our analysis comes from the UK Family Expenditure Survey (FES). The FES contains annual information on household consumption expenditures over the period 1968 to the present, which allows the study of changes in consumption inequality over a long period of time. We restrict the sample to individuals between the ages of 22 and 65 and eliminate students, retirees and the self-employed in the construction of our statistics.

We are interested in documenting two features of the data in particular: the rise in earnings and consumption inequality for households and the rise in the share of household earnings contributed by females. As is common in the literature, we start

by constructing a standard measure of individual-level consumption by dividing total household consumption by the square root of household size. We calculate the mean log deviation for the years 1968 to 2001 and present the trends in earnings and consumption inequality in Figure 1. The fact that the rise in consumption inequality coincides with a much larger increase in income inequality is well documented.[3] For the UK, the mean log deviation in consumption is approximately one-tenth the size of the mean log deviation in earnings, as illustrated in Figure 1. However, Figure 2 indicates the growth in consumption inequality (58%) between 1968 and 2001 is double the growth in earnings inequality (24%).

Interestingly, at the same time as income and consumption inequality are rising, a narrowing of the gaps in earnings between men and women is also present. One major implication of the falling gender wage gap is that the potential for women to provide a greater share of labor income for the family has grown over time. The distribution of potential income generated by the wife in two-adult households is illustrated in Figure 3. The measure of potential income we consider is the female's wage divided by the sum of the female and male's wage in the household.[4] This Figure shows that all quartiles as well as the 90^{th} percentile increased over the last three decades. Overall, potential earnings of wives increased by 13.5%. It is also interesting to consider how the *actual* earnings of husbands relative to wives changed over time. Figure 4 highlights the dramatic change in women's contribution to household labor income between 1968 and the present: women's share of earnings in the household increased by 93% over this time period. This difference partly reflects the increase in female's wages relative to those of men, but also the large changes in female and male labor force participation rates and hours worked since the 1960s. Figure 5 illustrates the convergence in employment rates for husbands and wives in two person households.

[3]Krueger and Perri (2003) find that the large rise in income inequality in the US since the 1970s while consumption inequality remained roughly constant.

[4]For households with missing wage data due to non-participation, we include a predicted wage based on a standard selection-corrected wage equation. Results are available from the authors upon request.

The gap in employment rates declined dramatically from 35.7 percentage points in 1968 to 1.5 percentage points in 2001. This change in employment rates represents a 95.9% decline in the gender gap in employment.

In summary, the evidence presented here highlights the fact that there has been a large rise in earnings and consumption inequality *across* households while at the same time there has been less inequality in the earnings distribution *within* households. Together, these phenomena may have very important implications for our understanding of consumption inequality. In the next section, we present a model of intra-household allocations that will enable us to study individual-level consumption allocations using data on labor supply and household-level consumption.

3 Theoretical Framework

As illustrated in Section 2, the fraction of income provided by women in two-person households has increased dramatically over the past 30 years. An extensive literature on intra-household allocations suggests that the source of the income plays an important role in determining how resources are allocated within households.[5] If household members do not share the same preferences, variation in the sources of family income may alter the consumption enjoyed by each member of the household. Therefore, it is important to consider consumption inequality in a framework that allows changes in individual incomes to affect consumption allocations within the household. To this end, we study a collective model of household decision making based on the framework of Chiappori (1988) and refined in Chiappori (1992), Browning et al. (1994), Browning and Chiappori (1998), Chiappori et al. (2002) and Blundell et al. (2002). This framework is ideal for the study of consumption inequality within households as it is less restrictive than the unitary model, which assumes all individuals in the household share a common set of preferences. It is also less restrictive than any particular form of bargaining, as the only restriction on the intra-household allocation

[5]See, for example, Lundberg, Pollak and Wales, (1997).

process is that households reach Pareto efficient allocations.

We start with a description of the problem faced by single agents. We then describe the intra-household allocation decision of married couples. Finally, we outline the model restrictions that allow for the identification of the share of consumption allocated to each household member.

3.1 Single Agents

Assume all single individuals have preferences over leisure and consumption. Denote leisure, private consumption and public consumption for an agent of gender g, $g \in \{m, f\}$ by L^g, C^g, and CP, respectively. The joint consumption of public goods is a primary gain to marriage and is an important component in the measurement of consumption inequality and is thus included here. Chiappori, Blundell and Meghir (2002) establish conditions under which the collective model with public goods is identified but do not estimate the model. Denote total household non-labor income (net of savings) Y and market wages w^g. Preferences for single agents are described by $U^g(u^g(L^g, C^g), CP)$, where it is assumed preferences over private consumption goods and leisure are separable from preferences over public consumption goods. Single person households choose labor supply and consumption to maximize utility, subject to the budget constraint:

$$\max_{L^g, C^g, CP} U^g(u^g(L^g, C^g), CP)$$

$$\text{subject to } w^g L^g + C^g + CP = w^g T + Y.$$

The right hand side of the budget constraint represents full income, where T is the total amount of time available for allocation between work and leisure. Defining Y as other income net of savings is consistent with inter-temporal substitution with time separable preferences and an exogenous wage profile, as in the two-stage budgeting process of Deaton and Muelbauer (1980).

3.2 Married Couples

Consider a two member household, where each member has distinct preferences over own leisure, own private consumption, and household public consumption. Denote by C a Hicksian composite good composed of private and public consumption:

$$C = C^f + C^m + CP.$$

We assume that married individuals have the same preferences over private consumption as single individuals, although preferences over leisure and public goods may depend on marital status. This assumption is crucial for the identification of the intra-household allocation process and is discussed in detail in Section 3.3. As with singles, assume that private consumption and leisure (C^g, L^g) are separable from consumption of the public good (CP) for married couples. Preferences for a married person of gender g can be described by:

$$V^g(v^g(L^g, C^g), CP),$$

where $v^g(L^g, C^g)$ captures preferences over private consumption and leisure. Under the assumptions that preferences are egoistic and that allocations are Pareto efficient, the household's allocations are the solution to the problem:

$$\max_{L^f, L^m, C^f, C^m, CP} \lambda V^f(v^f(L^f, C^f), CP) + (1-\lambda)V^m(v^m(L^m, C^m), CP) \qquad (1)$$

$$\text{subject to } w^f L^f + w^m L^m + C^f + C^m + CP = (w^f + w^m)T + Y.$$

The Pareto weight, λ, represents the female's bargaining power within the household, and will typically be a function of w^f, w^m, Y and other "distribution factors" that influence household bargaining power, but do not have an effect on individual preferences, as in Chiappori, Fortin, and Lacroix (2001).

Chiappori (1992) shows that the intra-household allocation problem faced by a husband and wife can be decentralized by considering a two stage process. In the first stage the husband and wife decide on the level of public good consumption

(CP) and on how to divide the remaining non labor income $y = Y - CP$. The assumption that consumption of the public good is separable from leisure and private consumption is key in allowing the allocation of public consumption to occur in the first stage (Chiappori, Blundell, and Meghir, 2002). Define the sharing rule $\phi(y, \mathbf{z})$ as the amount of non-labor income that is assigned to the wife. Then $y - \phi(y, \mathbf{z})$ is non-labor income assigned to the husband.

In the second step, each household member chooses his or her own private consumption and leisure, conditional on the level of public consumption and the budget constraint determined in the first stage:

$$\max_{L^g, C^g} v^g(L^g, C^g) \tag{2}$$

$$\text{subject to } w^g L^g + C^g = w^g T + \phi^g(y, \mathbf{z})$$

where $\phi^f = \phi(y, \mathbf{z})$ and $\phi^m = y - \phi(y, \mathbf{z})$. The Pareto problem represented in (1) and the sharing rule interpretation (2) are both representations of households behaving as if individuals are maximizing their own utility, with the household allocation being Pareto efficient. Both representations produce identical labor supplies and consumption demands, assuming an efficient level of public consumption is chosen in the first stage.

3.3 Identification of the Sharing Rule in the Case of Quadratic Preferences

Assume that L^f, L^m, Y, w^f, and w^m are observed in the data. As is consistent with our empirical exercise, C, and CP are observed although the distribution of consumption between the husband and wife of private consumption (C^f and C^m) is not observed.

Our treatment of households extends the models of Blundell, Chiappori, Magnac, and Meghir (2002) and Vermeulen (2003) to allow for households in which both spouses do not necessarily work full time and in which both spouses make labor force

participation decisions.[6] In particular, we assume that individuals can choose from H possible discrete labor supply possibilities, in addition to non-participation. This is not a very restrictive assumption, as the discrete choice of hours can be any integer value of weekly hours, and has been used in recent work by van Soest (1995) and Keane and Moffitt (1998). In addition, it provides a natural way to incorporate the participation decision, a margin that we think is likely important when looking at the long-term trends in consumption inequality. This assumption is not crucial to identification.

The main question we aim to address in this paper is how our measures of consumption inequality at the individual level differs from current measures in the literature that ignore inequality within the household. To provide an answer to this question, it is necessary to obtain an estimate of the full sharing rule to uncover the share of income allocated to each household member for consumption. In this case, the first order conditions of the sharing rule are not sufficient for identification. We therefore need to impose an additional restriction on preferences. As in Vermeulen (2003) we assume that married individuals have the same preferences for private consumption as single individuals, but possibly different preferences for leisure. This assumption still allows for the possibility that marriage allows for some complementarities in leisure time across spouses and for complementarities between leisure and consumption for married couples.

Let preferences for private consumption and leisure be represented by a quadratic direct utility function, a flexible form representing a second-order Taylor series expansion in leisure and consumption.[7] The utility an individual of gender g derives

[6] Blundell et al. (2002) model the labor force decision of the wife as continuous and of the husband as discrete; either he works full time or not at all. Vermeulen (2003) considers the case where males are assumed to work full-time and females face a discrete labor supply choice which includes the option of non-participation.

[7] A more general discussion of identification will be available in a future version of the paper.

from labor supply choice h is:

$$\begin{aligned} u_h^g &= v^g(l_h, C_h^g) + \omega^g(CP) + \varepsilon_h^g \\ &= \beta_l^g l_h + \beta_{ll}^g l_h^2 + \beta_{cl}^g l_h C_h^g + \beta_c^g C_h^g + \beta_{cc}^g (C_h^g)^2 + \beta_{cP}^g CP_h^g + \beta_{ccP}^g (CP_h^g)^2 + \varepsilon_h^g, \end{aligned}$$

where $l_h = T - L_h$ is labor supply, and ε_h^g is an unobserved preference component that is assumed to be distributed *iid* across individuals and labor market states. This specification allows preferences for leisure to differ between married and single men and women, but restricts preferences for consumption to be the same for both married and singles.[8]

Assume the sharing rule is linear in the distribution factors:

$$\begin{aligned} \phi(y, \mathbf{z}) &= \left(\phi_0 + \sum_{k=1}^{K} \phi_k z_k\right) y \\ &= (\mathbf{z}^\top \phi) y, \end{aligned}$$

where there are K distribution factors plus a constant in the vector $\mathbf{z}$ and where y is non-labor income net of expenditures on the public good.[9] We can condition on household expenditures on the public good for both singles and married couples under the assumptions that households make efficient decisions in the first stage and that preferences over public goods are separable from preferences over consumption and leisure (Deaton and Muelbauer, 1980), defining y as non-labor income net of expenditures on public goods.

The budget constraints for the second stage of the budgeting process can be expressed as:

$$C_h^g = w^g l_h + y \tag{3}$$

for single individuals,

$$C_h^f = w^f l_h + (\mathbf{z}^\top \phi) y \tag{4}$$

[8]The interaction between consumption and leisure represented by $\beta_{cl}^g \; lC_h^g$ is limited to consumption of the private good to maintain separability from public goods consumption.

[9]Vermeulen (2003) writes the sharing rule as a function of total household expenditure, rather than of total non-labor income. While his definition is convenient for estimation, it is not as consistent with the two stage decision process discussed in Section 3.2, since total household expenditure is a function of the spouse's labor supply decision.

for married women and

$$C_h^m = w^m l + (1 - \mathbf{z}^\top \phi) y \tag{5}$$

for married men.[10]

For the decision process, only differences in utility matter, and thus the parameters must be estimated relative to some base utility. We assume that the choice of not working ($h = 0$) is the base case. After substituting the budget constraint into the utility function, the difference between working $h > 0, \forall h \in \{1, 2, ..., H\}$ and not working ($h = 0$) for single men and women can be expressed as:

$$u_h^g - u_0^g = \beta_l^g l_h + \beta_{ll}^g l_h^2 + \beta_{cl}^g l_h w^g l_h + \beta_c^g w^g l_h + \beta_{cc}^g \left[(w^g l_h)^2 + 2 w^g l_h y \right] + \varepsilon_h^g - \varepsilon_0^g. \tag{6}$$

Consider next the problem of a married woman. The difference between working $h > 0, \forall h \in \{1, 2, ..., H\}$ and not working ($h = 0$) is described by:

$$\begin{aligned} u_h^f - u_0^f &= \beta_l^f l_h + \beta_{ll}^f l_h^2 + \beta_{cl}^f l_h w^f l_h + \beta_{cl}^f (\mathbf{z}^\top \phi) l_h y \\ &\quad + \beta_c^f w^f l_h + \beta_{cc}^f (w^f l_h)^2 + 2\beta_{cc}^f (\mathbf{z}^\top \phi) w^f l_h y + \varepsilon_h^f - \varepsilon_0^f. \end{aligned} \tag{7}$$

Finally, consider the problem of a married man, where the difference between working $h > 0, \forall h \in \{1, 2, ..., H\}$ and not working ($h = 0$) is described by:

$$\begin{aligned} u_h^m - u_0^m &= \beta_l^m l_h + \beta_{ll}^m l_h^2 + \beta_{cl}^m l_h w^m l_h + \beta_{cl}^m l_h y - \beta_{cl}^m (\mathbf{z}^\top \phi) l_h y \\ &\quad + \beta_c^m w^m l_h + \beta_{cc}^m (w^m l_h)^2 + 2\beta_{cc}^m w^m l_h y - 2\beta_{cc}^m (\mathbf{z}^\top \phi) w^m l_h y + \varepsilon_h^m - \varepsilon_0^m. \end{aligned} \tag{8}$$

The parameters $\beta_l^g, \beta_{ll}^g, \beta_{cl}^g, \beta_c^g$, and β_{cc}^g are directly identified. The parameters capturing preferences over the public consumption good ($\beta_{cP}^g, \beta_{ccP}^g$) can not be identified as the utility from consumption of the public good is the same regardless of the labor supply decision.[11]

[10] For the purposes of notational simplicity we do not include taxes explicitly. $w^g l_h$ should be interpreted as labor income net of taxes, and $l_h w^g l_h$ should be interpreted as hours worked times net labor income.

[11] One implication is that we will be able to estimate the sharing rule but not fully recover preferences. As a result, we cannot make welfare comparisons.

Equations (6), (7) and (8) can be expressed for all individuals in reduced form as:

$$\begin{aligned} u_h^g - u_0^g &= \Pi_l^g l_h + \Pi_{ll}^g l_h^2 + \Pi_{ly}^g l_h y + \Pi_{lwl}^g l_h w^g l_h + \Pi_{wl}^g w^g l_h + \Pi_{(wl)^2}^g \left(w^g l_h\right)^2 \\ &+ \Pi_{wly}^g w^g l_h y + \Pi_{lym}^g l_h y m + \mathbf{\Pi}_{zlym}^g \cdot \mathbf{z} l_h y m + \Pi_{wlym}^g w^g l_h y m \\ &+ \mathbf{\Pi}_{zwlym}^g \cdot \mathbf{z} w^g l_h y m + \varepsilon_h^g - \varepsilon_0^g, \end{aligned}$$

where m is a dummy variable that takes a value of 1 for married individuals and 0 for single individuals.

The system above implies a series of over-identifying restrictions on the assumptions of the collective model, the functional form for preferences and the sharing rule, and our assumption that preferences for consumption are the same regardless of marital status:

$$\begin{aligned} \phi_0 &= \frac{\Pi_{lym}^f}{\Pi_{ly}^f} + 1 = \frac{\Pi_{wlym}^f}{\Pi_{wly}^f} + 1 = -\frac{\Pi_{lym}^m}{\Pi_{ly}^m} = -\frac{\Pi_{wlym}^m}{\Pi_{wly}^m}, \\ \phi_k &= \frac{\Pi_{z_k lym}^f}{\Pi_{ly}^f} = \frac{\Pi_{z_k wlym}^f}{\Pi_{wly}^f} = -\frac{\Pi_{z_k lym}^m}{\Pi_{ly}^m} = -\frac{\Pi_{z_k wlym}^m}{\Pi_{wly}^m}, \quad k = 1 \ldots K. \end{aligned}$$

In the following section, we outline our strategy for estimating the model and testing the above restrictions using consumption data from the UK.

4 Data

The data we use to conduct our analysis comes from the UK Family Expenditure Survey (FES). This data is ideal for the study of consumption inequality for three reasons. First, it contains detailed information on household private and public consumption expenditures, wages and labor supply for individuals within households, and demographic information including age, sex, education (from 1978 onward) and region of residence. Second, the FES has fewer problems with measurement issues than the leading contenders in the US and other countries do. Battistin (2003) documents reporting errors in the US Consumer Expenditure Survey due to survey design. The FES uses a weekly diary to collect data on frequently purchased items and uses

recall questions to collect data on large and infrequent expenditures. Finally, the FES contains annual information over the period 1968 to the present, which allows the study of changes in consumption inequality over a long period of time.[12]

Our sample is composed of single person households and couples without children. We exclude households with children in this paper to abstract from the intra-household allocation of resources for children's consumption. This is obviously an important issue, but one we leave to future work. We restrict the age range in the sample to individuals between the ages of 22 and 65 and eliminate students so as to eliminate individuals that are not making labor supply decisions. We also exclude individuals who are self employed due to endogeneity of wages. Additionally we exclude households in which one of the individuals is in the top one per cent of the wage distribution. The resulting sample is composed of $87,668$ individuals.[13] Descriptive statistics for our entire sample, and the years 1968 and 2001 are presented in Table 1.

We define consumption and non-labor income measures as follows. Total consumption is defined as total household expenditures. Public consumption is defined as expenditures on housing, light and power, and household durable goods. Private household consumption is total expenditures net of public consumption. We use an expenditure based definition of non-labor income which is consistent with a two stage budgeting process with time separable preferences and separability of public goods consumption from leisure and private consumption; other income is defined as total household expenditures minus net labor income. In our estimation we subtract household expenditures on public goods from other income, resulting in non-labor income net of public goods consumption. In addition to the separability assumptions, wage profiles are assumed to be exogenous. This rules out the possibility of job-specific human capital accumulation.

In order to construct the level of consumption corresponding to each labor supply

[12]We presently exclude the year 1997 from our analysis due to a missing data problem.

[13]The sample size in 1968 is 2,584 and the sample size in 2001 is 2,757. The sample sizes do not vary markedly across years: the smallest sample is 2,502 in 1979 and the largest is 2,932 in 2000.

decision we need to assign a wage to all individuals. For those who are working we use their usual hourly wage, defined as weekly earnings divided by usual weekly hours. For non-participants we use a predicted wage, computed based on a standard reduced form selection-corrected wage equation. The log of the wage is estimated as a function of age, birth cohort, year, quarter, and regional dummies, (with the year effect normalized to be orthogonal to a time trend and sum to zero over the sample period as in Deaton (1997)), plus the age at which full time education was completed, and its square. The selection equation is identified by the exclusion from the wage equation of household nonlabor income, marital status, and the age, education, and the labor income of the spouse.[14] The resulting wage measure is treated as known and is used to construct the within household distribution factor defined as the *potential* share of household labor income contributed by the wife, $z = w^f/(w^f + w^m)$.

For each possible labor supply choice we construct the corresponding consumption level, using earnings net of taxes. We calculate after tax earnings by converting weekly wage income to an annual base, deducting the appropriate personal allowance and then applying the appropriate tax rate. Personal allowances and marginal tax rates are from the Board of Inland Revenue (1968–2001). All monetary values are expressed in 1987 pounds.

Labor supply is measured by a discrete variable that takes on three values: not participating, working part-time and working full-time. Full time is defined as working 35 hours per week or more, and part-time is defined as 1 to 34 hours per week. The choice of these ranges is based on the the hours histograms in Figures 6 and 7, which suggest a full-time definition of 35 hours a week or more. The average hours worked in the part-time category is approximately 20 hours per week, and approximately 40 hours per week in the full-time category.

In order to ensure consistency between the number of hours worked in each of the three states and the corresponding consumption level we adopt the following

[14]Results for this wage equation are available from the authors upon request.

convention. If an individual is observed to be working either part-time or full-time we use their actual number of hours as their labor supply for the corresponding choice, and their usual take home pay in constructing the corresponding consumption. For the labor supply states we do not observe, we use 20 and 40 hours for part-time and full-time labor supply, and calculate after tax earnings based on these hours and the appropriate tax rates. Constructing the individual consumptions in this way ensures that within a household, private consumption adds up to what we observe in the data.

Likely candidates for the distribution factors are the wife's potential share of total household labor income ($w_i^f/(w_i^f + w_i^m)$), the local sex ratio, and an index of the generosity to the wife of local divorce legislation (Chiappori, Fortin, and Lacroix, 2002). At present, we consider the wife's share of potential labor earnings, presented in Figure 3, as a distribution factor in estimation.

4.1 Estimation Results

The model of Section 3.3 can be estimated using a multinomial logit under the assumption that the disturbances ε_{ih} are independent and identically distributed with type I extreme value distribution. Let d_{ih}^g denote an indicator equal to 1 if individual i makes labor supply choice h and zero otherwise. The contribution of individual i to the likelihood function is the probability of observing individual i making labor force decision h, which has the form:

$$\begin{aligned} \Pr(d_{ih}^g = 1) &= \Pr(u_{ih}^g > u_{ij}^g, \forall j \neq h; j, h \in \{0, 1, ..., H\}) \\ &= \frac{\exp(v^g(L_{ih}, C_{ih}; X_i, z_i))}{\sum_{j=0}^{H} \exp(v^g(L_{ij}, C_{ij}; X_i, z_i))}. \end{aligned}$$

In the empirical estimation we allow for heterogeneity in preferences for leisure through the vector X which includes age, birth cohort, education, region, and quarter and year to control for seasonality and cyclical effects. In order to break the collinearity between age, birth cohort and year we follow Deaton (1997) and transform the

year dummy variables so that the coefficients are orthogonal to a time trend and sum to zero over the period 1968 to 2001. The parameters β_h and β_{ll} are assumed to be linear functions of these observables, so that for individual i we have

$$\begin{aligned} \beta_l &= X_i\beta_l \\ \beta_{ll} &= X_i\beta_{ll}, \end{aligned}$$

where X_i is a vector of observables and β is a vector of parameters.[15]

Our estimation proceeds in two steps. First we estimate a standard selection corrected wage equation from which we predict wages for those individuals we do not observe working in the data. Second we estimate the multinomial logit, treating wages as known.[16]

As discussed in Section 3.3, with quadratic utility and with the assumption that preferences over private consumption are the same for married and single individuals, we can construct each of the sharing rule parameters in four different ways from estimation of the unrestricted reduced form. The unrestricted reduced form estimates, presented in Table 2, are used to recover the parameters of the sharing rule. Estimates of the sharing rule parameters, along with their standard errors and 95% confidence intervals (constructed using the delta method) are presented in Table 3. In addition, a series of tests of the assumption of common preferences for private consumption between married and single individuals are presented.

The estimated sharing rule parameters constructed from the different restrictions are remarkably similar, especially for men. The point estimate for the intercept falls

[15]In addition to observed heterogeneity in tastes for leisure, unobserved preference heterogeneity is likely to be very important. Future versions of this paper will allow for unobserved heterogeneity and specify $\beta_h = X_i\beta_h + u_{hi}$ and $\beta_{hh} = X_i\beta_{hh} + u_{hhi}$, with $u_{hi} \sim N(0, \sigma_h^2)$ and $u_{hhi} \sim N(0, \sigma_{hh}^2)$ (see Train (2003)). The contribution to the likelihood function then becomes

$$\Pr(d_{ih}^g = 1) = \int\int \frac{\exp(v^g(L_{ih}, C_{ih}; X_i, z_i, u_{hi}, u_{hhi}))}{\sum_{j=0}^{H} \exp(v^g(L_{ij}, c_{ij}; X_i, z_i, u_{hi}, u_{hhi}))} dF(u_{hi})dF(u_{hhi}),$$

which does not have a closed form solution, but can be estimated using Simulated Maximum Likelihood.

[16]For both efficiency reasons, and to obtain appropriate standard errors we will jointly estimate the wage equation and the discrete choice labor supply decision in future versions.

in the range -0.423 to -0.276, and the point estimate for the slope ranges from 1.366 to 1.705. Together, the unrestricted results suggest that an increase in the female's share of potential earnings increases her share of total consumption in the household. This result highlights the potential importance of the rise in female wages in reducing intra-household consumption inequality.

A Wald test on the unrestricted model rejects the restriction that preferences for private consumption are the same across marital status for both men and women: the calculated χ^2 statistic with 6 degrees of freedom is 78.02, which exceeds the critical value of 1.635 at the 95 per cent confidence level. Testing the restrictions separately for men and women reveals that the assumption is rejected for women, but it is not rejected for men. The outcome of this test, unfortunately, does not reveal whether the functional form for preferences is rejected or the model restriction on preferences over private consumption or both. It seems that further exploration with flexible functional forms and perhaps different identifying assumptions are required for women. For the purposes of constructing measures of individual consumption for our inequality decomposition we use the sharing rule estimates constructed from the male preference parameters. Experimentation with other representations of the utility function and other identification strategies is left for future work.

Before proceeding to a discussion of consumption inequality, it is of interest to consider the conditions under which our collective model predict consumption is split equally across men and women. To this end, we use the sharing rule estimates constructed from the male preferences parameters and determine what value of the female's share in potential household earnings satisfies:

$$\frac{1}{2} = \hat{\phi}_0 + \hat{\phi}_1 \cdot \frac{w^f}{w^f + w^m}.$$

Using estimates for ϕ_0 and ϕ_1 of 0.280 and 1.582, respectively yields 49.4%. In other words, the model predicts non-labor income is split equally between the husband and wife when the female's share of potential earnings in the household is 49.4%! For consumption to be equally split the husband and wife must also choose identical hours

of work. It is worth emphasizing that this highly intuitive result is derived not from a model in which equal sharing is assumed: the only assumptions imposed in estimation are that households make Pareto Efficient decisions, that public consumption is separable from private consumption, and that the husband's preferences over private consumption goods are the same when single as when married.

4.2 Decomposing the Trends in Consumption Inequality

To uncover the degree of consumption inequality in the household, we proceed as follows. First we estimate the empirical model of Section 3.3 for the purpose of recovering the sharing rule $\phi(\mathbf{z})$. Next we use this sharing rule to divide non-labor income between the husband and wife in each household and construct private consumption based on the individuals' share of non-labor income and his or her personal net labor earnings. Private consumption is constructed as in equations (4) and (5). Our measure of individual consumption is then equal to individual private consumption, plus household public consumption. For single individuals consumption is simply total consumption expenditure. For married individuals our measure of consumption is

$$\begin{aligned} C^f &= CP + \tau(w^f L_h^f) + (\phi_0 + \phi_1 z)y \\ C^m &= CP + \tau(w^m L_h^m) + (1 - \phi_0 - \phi_1 z)y, \end{aligned}$$

where $\tau(w^g L_h^g)$ is labor income net of tax and $z = w^f/(w^f + w^m)$ is the potential share of non-labor income due to the wife.

For comparison purposes, we construct another measure of individual consumption, equal division, which assumes that all consumption is divided equally between the husband and wife. For the equal division case, individual consumption is calculated as household public consumption plus one half of household private consumption. In both the *sharing rule* and the *equal division* case, we double count public consumption. This accomplishes the same end as using an equivalency scale to assign household consumption to individual members. The advantage of our method is that

it becomes very clear why households have economies of scale: both individuals in the household can consume the public good.

Having constructed these two measures of individual consumption, we can construct a time series of inequality measures, and decompose them into changes in between and within household inequality. While the Gini coefficient is probably the best known and most widely used inequality index, it does not allow overall inequality to be exactly decomposed into within and between group contributions. As this is the main objective of this paper we use an inequality measure belonging to the "Generalized Entropy" family which has the property that overall inequality is the weighted sum of the inequality within each group, in our case each household, and the inequality between the groups. This decomposition is discussed in Shorrocks (1984), from which we draw the following discussion.

The Generalized Entropy inequality index is a single parameter family that can be expressed in terms of the sum of ratios of individual consumption C^g to the mean consumption μ_C:

$$I_\alpha(C) \begin{cases} \frac{1}{n}\frac{1}{\alpha(\alpha-1)}\sum_{i=1}^{n}\left\{\left(\frac{C_i^g}{\mu_C}\right)^\alpha - 1\right\} & \text{if } \alpha \neq 0, 1, \\ \frac{1}{n}\sum_{i=1}^{n}\frac{C_i^g}{\mu_C}\log\left(\frac{C_i^g}{\mu_C}\right) & \text{if } \alpha = 1, \\ \frac{1}{n}\sum_{i=1}^{n}\log\left(\frac{\mu_C}{C_i^g}\right) & \text{if } \alpha = 0. \end{cases} \tag{9}$$

The limiting cases of $\alpha = 1$ and $\alpha = 0$ are respectively the Theil index and the Mean Logarithmic Deviation (MLD). We use the mean logarithmic deviation version of the generalized entropy index as it is convenient, and is more sensitive to changes at the bottom end of the distribution than higher values of α.

In our measurement of consumption inequality, the overall index can be additively decomposed into within and between household inequality:

$$I_\alpha(C) = I_\alpha^W(C) + I_\alpha^B(C),$$

where $I_\alpha^W(C)$ is the within household inequality and $I_\alpha^B(C)$ is the between household inequality. Since under the assumption of equal division within household inequality

is zero, we can calculate $I_\alpha^B(C)$ by using equal division. Using individual consumption constructed with the sharing rule we obtain the overall inequality index $I_\alpha(C)$. We can then recover intra-household inequality as

$$I_\alpha^W(C) = I_\alpha(C) - I_\alpha^B(C). \tag{10}$$

Using Equation (10) we calculate the inequality index using both the sharing rule and the equal division construction of individual consumption. We then recover the index of within household inequality as the difference between the two indices. The time-series trend of total, between household, and within household inequality for the years 1968 to 2001 is presented in Figure 8. Inequality was basically stable from 1968 to 1980 at which time it increased substantially until around 1990, and has been falling slightly from 1990 through 2001. Of particular interest are two findings. First, our estimates of consumption inequality under the sharing rule assumption are approximately one-third higher than a measure of inequality that ignores intra-household allocations.[17]

Second, our results indicate that the rise in women's labor supply and the fall in the gender wage gap over time tempered the rise in consumption inequality measured under the assumption of equal division: the rise in consumption inequality under equal division may be over-stated by as much as 40%. The reason the sharing rule measure of inequality differs so markedly from the equal division measure is due to the large fall in within-household inequality. In particular, due to the rise in female wages and labor supply, within-household inequality fell by 20% between 1968 and 2001.[18]

[17] At present we have not constructed standard errors, preventing us from making statements about the statistical significance of the differences.

[18] The compression of marginal tax rates also appear to have played a role in generating the sharp rise in between household inequality during the 1980s. The top and bottom marginal tax rates are plotted in Figure 10, where the top marginal rate falls from 83 per cent in 1978 to 60 per cent in 1979, and then falls again to 40 per cent in 1988. The increase in between household consumption inequality is closely linked to the increase in after tax income inequality that occurred over the 1980s. The changes in marginal tax rates had the effect of increasing between household inequality substantially while having only a modest effect on within household inequality. This result is not surprising, as we would expect changes in marginal taxes to have a larger effect across households than within households due to positive assortive matching on education.

4.3 Extensions

4.3.1 Cohort analysis

The first extension to our analysis we consider is a decomposition of consumption inequality by cohorts. This is an important exercise to undertake for two reasons. First, comparisons based on repeated cross-sectional data confound between and within cohort effects. In particular, comparisons of inequality at a point in time involve comparing individuals at different stages in the life-cycle (Blundell and Preston, 1998). Second, a rise in inequality over time may simply reflect increases in inequality within a birth cohort, due to the accumulation of persistent shocks over the life-cycle (Deaton and Paxton, 1994). The results of our cohort analysis are presented in Figure 11. The results indicate that the decline in within household inequality described above is due primarily to a cohort effect: the gender gap in wages was higher and labor market opportunities for women were lower for older cohorts. The second finding of interest is that although inequality fell across cohorts, within household inequality increased over time within cohorts.

4.3.2 Changes in household composition over time

The next extension we consider is the importance of changing household composition. In particular, we consider the effect of the rise in households with one adult in our current measure of consumption inequality. Although single adult households have no within household inequality by definition, it is still the case that there may exist substantial inequality across single adult households. To address this issue, we present measures of the mean log deviation in consumption for a sample that excludes singles. Figure 12 suggests that the level of consumption inequality is approximately 20% lower when singles are eliminated from the analysis. This suggests heterogeneity across single adult households is an important source of inequality. Figure 13 suggests, interestingly, that the rise in inequality is not related to the rise in singles, as the MLD in consumption relative to 1968 is the same regardless of whether singles are included.

4.4 Sensitivity Analysis

In this section, we discuss the robustness of our results to several variations in the model and empirical exercise. The first robustness check we consider is whether the results are sensitive to our definitions of public and private consumption from the data. We consider an estimate with no public goods and then sequentially add housing, heat and lighting, household durables, transport and services.[19] The estimates of the sharing rule are quite consistent across specifications. Second, we estimate the model including a second distribution factor: the difference between the husband's and wife's age. Again, we find the model is rejected for female's but not for men.

5 Conclusions

Our paper makes two contributions to the literature on consumption inequality. First, we construct and estimate a model of intra-household allocations to examine how changes in the source of income in the household translate into changes in individual-level consumption allocations. Second, we use estimates from our model to make inferences regarding how consumption inequality *within* households relates to consumption inequality *across* households. Our estimates suggest that measures of consumption inequality that ignore the potential for intra-household inequality may underestimate the level individual consumption inequality by 40% and may over-state the rise in individual consumption inequality by 40%. The results of our analysis highlight the importance of intra-household allocations for our understanding of consumption inequality and its implications.

[19]Full estimation results are available from the authors upon request.

Table 1: FES Descriptive Statistics.

	Male				Female			
	Single		Married		Single		Married	
1968	Mean	SD	Mean	SD	Mean	SD	Mean	SD
Age (22 to 65)	50.28	12.25	49.08	13.19	54.40	9.79	47.10	13.71
No hours dummy	0.13	0.34	0.10	0.30	0.40	0.49	0.46	0.50
Part time dummy	0.03	0.17	0.03	0.17	0.18	0.38	0.21	0.41
Full time dummy	0.84	0.37	0.87	0.34	0.42	0.49	0.33	0.47
Hourly wage	3.11	1.76	3.36	1.57	2.46	1.59	2.00	0.98
Total Expend.	94.42	70.18	147.35	98.55	76.06	50.94	147.35	98.55
Housing Expend.	27.02	43.29	40.52	57.01	27.08	19.34	40.52	57.01
Observations	174		1033		344		1033	
Observed wages	148		919		204		555	
2001								
Age (22 to 65)	43.86	12.10	49.01	12.90	47.81	12.70	46.94	12.96
No hours dummy	0.32	0.47	0.30	0.46	0.39	0.49	0.31	0.46
Part time dummy	0.07	0.26	0.06	0.24	0.13	0.33	0.24	0.43
Full time dummy	0.60	0.49	0.64	0.48	0.48	0.50	0.44	0.50
Hourly wage	6.00	2.94	6.02	2.88	4.93	2.06	4.31	1.79
Total Expend.	134.69	113.14	248.05	162.68	130.64	87.30	248.05	162.68
Housing Expend.	48.52	42.88	83.06	80.70	56.21	47.77	83.06	80.70
Observations	508		886		477		886	
Observed wages	336		615		283		592	
1968 to 2001								
Age (22 to 65)	43.92	13.39	48.60	13.51	50.41	13.31	46.51	13.60
No hours dummy	0.28	0.45	0.19	0.40	0.43	0.49	0.35	0.48
Part time dummy	0.05	0.22	0.04	0.20	0.16	0.37	0.25	0.43
Full time dummy	0.68	0.47	0.77	0.42	0.41	0.49	0.40	0.49
Hourly wage	4.94	2.43	4.74	2.27	3.86	1.96	3.33	1.66
Total Expend.	118.37	99.10	192.51	129.48	99.53	75.63	192.51	129.48
Housing Expend.	39.57	41.62	59.61	62.93	39.07	37.91	59.61	62.93
Observations	10,958		31,871		12,967		31,871	
Observed wages	7,663		25,208		7,271		20,291	

Table 2: Discrete Choice Labor Supply Estimates.

	Coef.	Std. Error		Coef.	Std. Error
Π^f_l	0.300	0.061	Π^m_l	0.405	0.117
Π^f_{ll}	−6.373	1.470	Π^m_{ll}	−8.660	2.785
Π^f_{lm}	0.007	0.003	Π^m_{lm}	0.034	0.006
Π^f_{llm}	−0.001	0.000	Π^m_{llm}	−0.000	0.000
Π^f_{ly}	−1.689	0.058	Π^m_{ly}	−1.263	0.039
Π^f_{lym}	2.194	0.071	Π^m_{lym}	−0.349	0.082
Π^f_{lyzm}	−2.306	0.102	Π^f_{lyzm}	1.996	0.158
$\Pi^f_{wl^2}$	0.127	0.052	$\Pi^f_{wl^2}$	−1.492	0.075
Π^f_{wl}	−0.048	0.002	Π^m_{wl}	0.007	0.002
$\Pi^m_{(wl)^2}$	0.143	0.006	$\Pi^m_{(wl)^2}$	0.138	0.005
Π^f_{ywl}	0.211	0.015	Π^m_{ywl}	0.145	0.007
Π^f_{ywlm}	−0.301	0.020	Π^f_{ywlm}	0.041	0.016
Π^f_{ywlzm}	0.359	0.028	Π^m_{ywlzm}	−0.229	0.034
Number of obs = 263,004 (87,668 individual × 3 labor force states)					
Log likelihood = -47,409.023			Prob > χ^2 = 0.0000		
LR $\chi^2(286)$ = 97,808.24			Pseudo R^2 = 0.5078		

Note: Coefficients for the interactions of labor supply with individual characteristics are ommited for space constraints. They are available on request.

Table 3: Sharing Rule Estimates.

	Coef.	Std. Err	[95% c.i.]	
$\phi_{0a} = \Pi^f_{lym}/\Pi^f_{ly} + 1$	-0.300	0.025	-0.349	-0.250
$\phi_{0b} = \Pi^f_{lwym}/\Pi^f_{lwy} + 1$	-0.430	0.068	-0.563	-0.297
$\phi_{0c} = -\Pi^m_{lym}/\Pi^m_{ly}$	-0.276	0.069	-0.411	-0.141
$\phi_{0d} = -\Pi^m_{lwym}/\Pi^m_{lwy}$	-0.284	0.116	-0.511	-0.057
$\phi_{1a} = \Pi^f_{lyzm}/\Pi^f_{ly}$	1.366	0.076	1.217	1.514
$\phi_{1b} = \Pi^f_{lwyzm}/\Pi^f_{lwy}$	1.705	0.178	1.356	2.053
$\phi_{1c} = -\Pi^m_{lyzm}/\Pi^m_{ly}$	1.580	0.134	1.318	1.842
$\phi_{1d} = -\Pi^m_{lwyzm}/\Pi^m_{lwy}$	1.583	0.246	1.099	2.066

Tests	df	χ^2	p-value
$\phi_{0a} = \phi_{0b} = \phi_{0c} = \phi_{0d}$	6	78.02	0.00
and $\phi_{1a} = \phi_{1b} = \phi_{1c} = \phi_{1c}$			
$\phi_{0a} = \phi_{0b}$ and $\phi_{1a} = \phi_{1b}$	2	8.66	0.01
$\phi_{0c} = \phi_{0d}$ and $\phi_{1c} = \phi_{1d}$	2	0.10	0.95

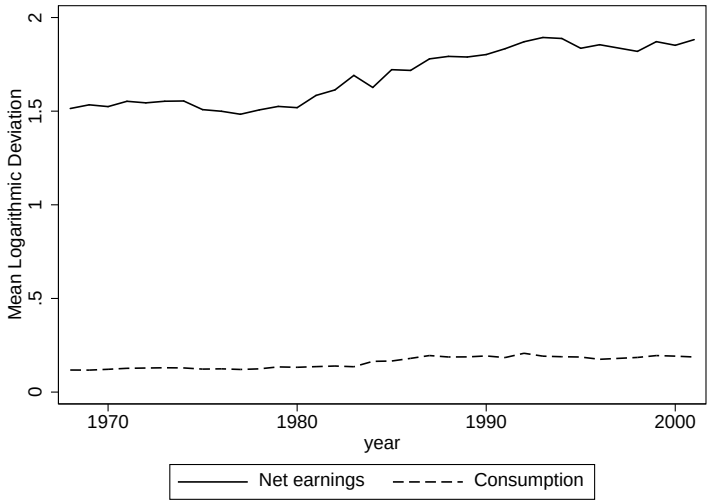

Figure 1: **Trends in consumption and earnings inequality.**
Own calculations from the FES.

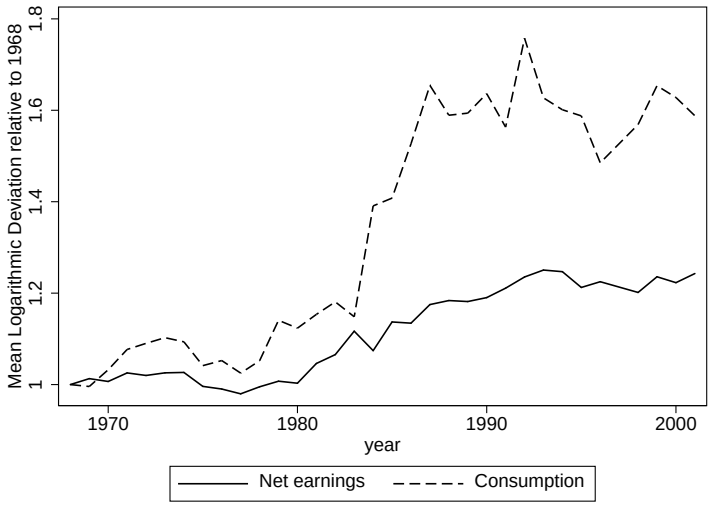

Figure 2: **Growth in consumption and earnings inequality.**
Own calculations from the FES.

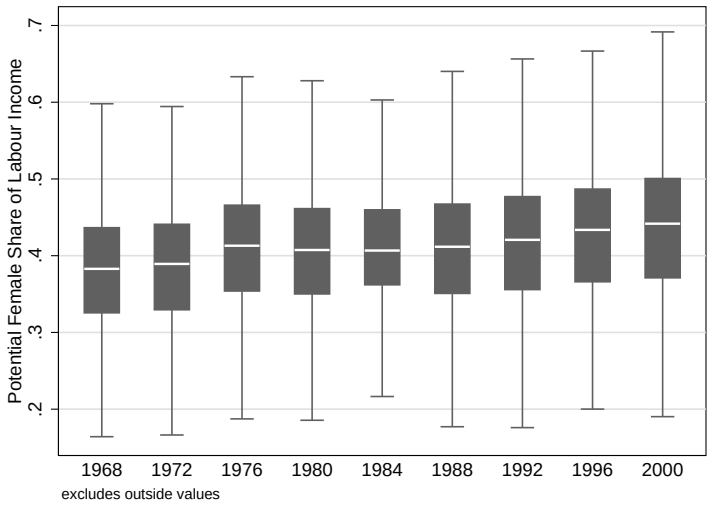

Figure 3: **Fraction of potential household earnings provided by wife.**
Source: Own calculations from the FES.

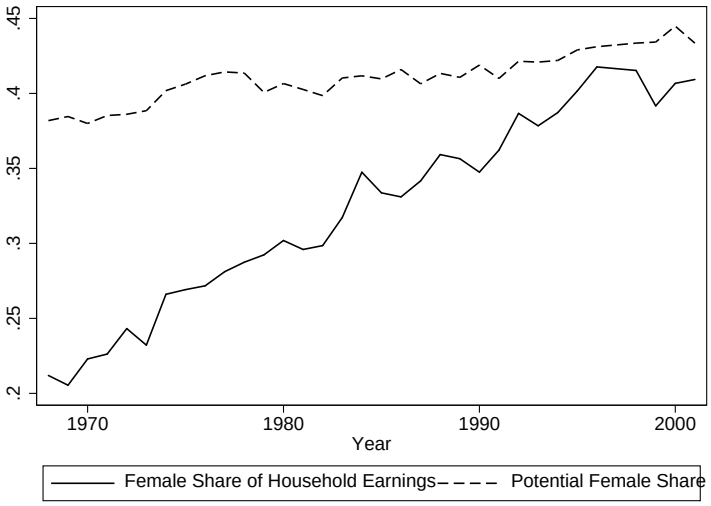

Figure 4: **Fraction of actual household earnings provided by wife.**
Source: Own calculation from the FES.

Figure 5: **Employment rates of husband and wives in childless couples.**
Source: Own calculations from the FES.

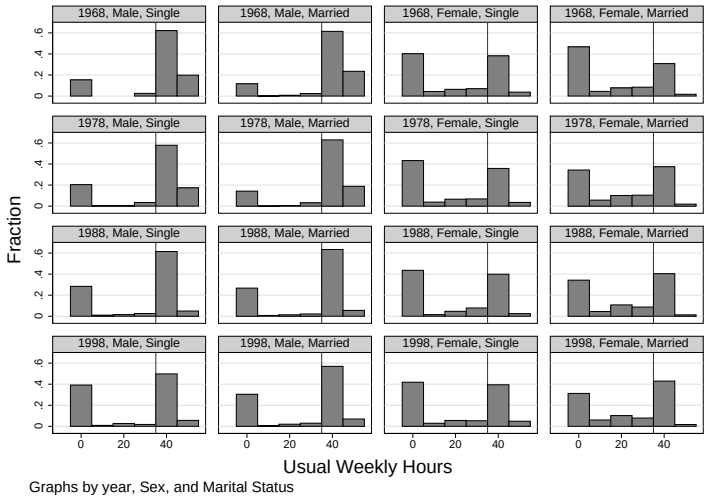

Figure 6: **Histogram of usual weekly hours.**
Source: Own calculation from the FES.

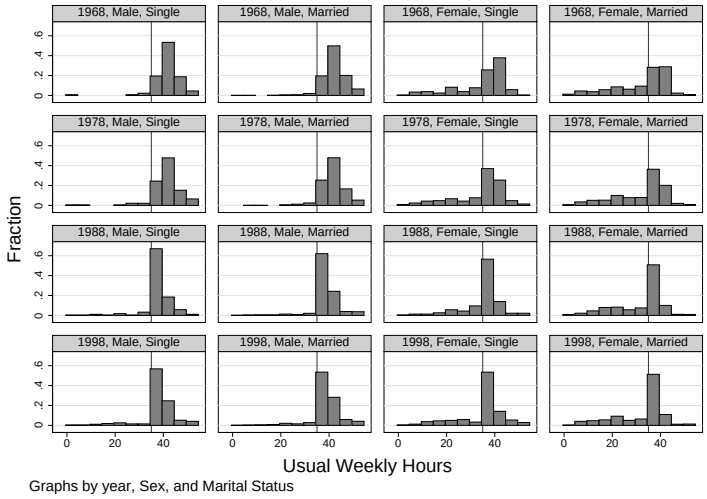

Figure 7: **Histogram of usual weekly hours for participants.**
Source: Own calculation from the FES.

Figure 8: **Within and between household decomposition of inequality trends.**

Source: Own calculation from the FES.

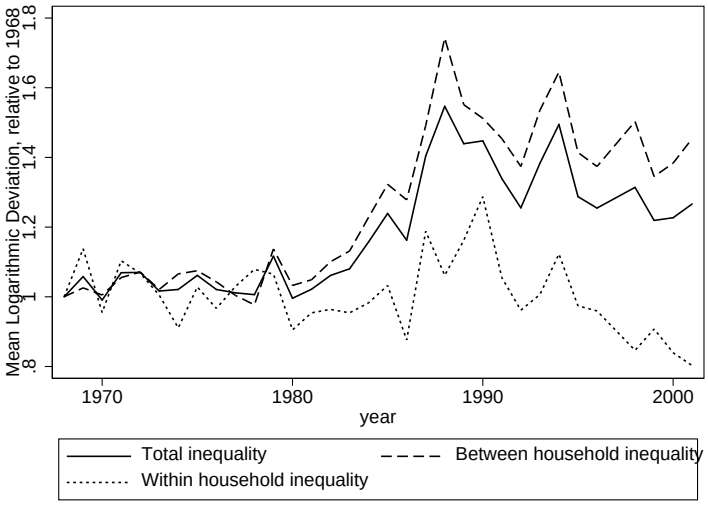

Figure 9: **Relative changes in between and within household inequality trends.**

Source: Own calculation from the FES.

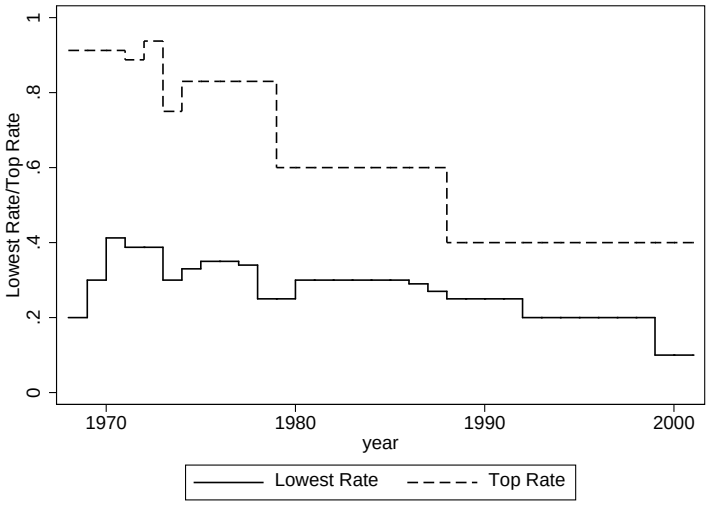

Figure 10: **Bottom and top marginal tax rates.**
Source: UK National Statistics (1968–2001).

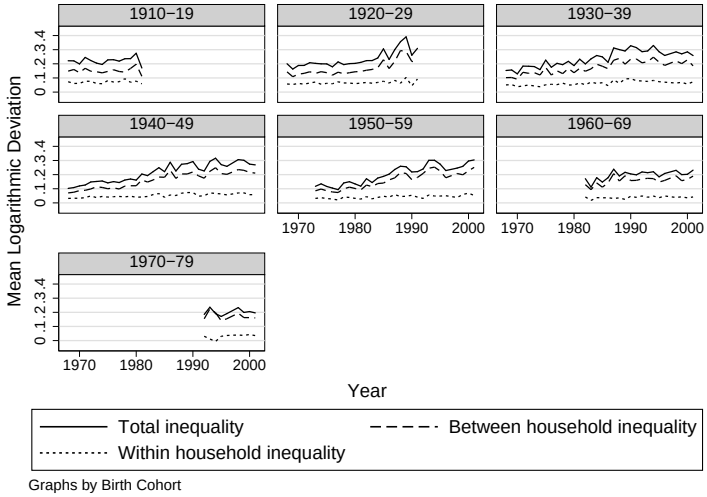

Figure 11: **Consumption Inequality by Birth Cohort.**
Source: UK National Statistics (1968–2001).

Figure 12: **MLD in Consumption - Sample of Couples**
Source: UK National Statistics (1968–2001).

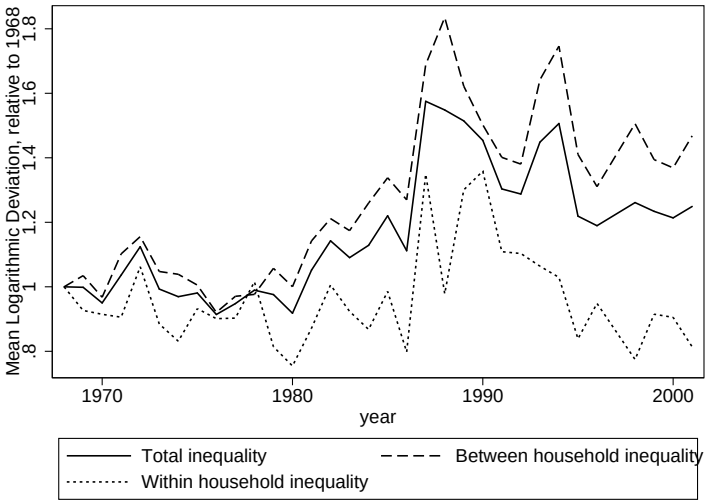

Figure 13: **MLD in Consumption Relative to 1968 - Sample of Couples**
Source: UK National Statistics (1968–2001).

References

Barrett, G., T. Crossley, and C. Worswick (2000): "Consumption and Income Inequality in Australia," *Economic Record*, 76(233), 116–138.

Battistin, E. (2003): "Errors in Survey Reports of Consumption Expenditures," Institute for Fiscal Studies Working Paper.

Beaupre, L. (2001): "Intrahousehold Income Distribution and Inequality Measurement: The Effects of Sharing on Inequality," Unpublished Manuscript.

Blundell, R., P.-A. Chiappori, T. Magnac, and C. Meghir (2002): "Collective Labor Supply: Heterogeneity and Nonparticipation," Unpublished Manuscript.

Browning, M., F. Bourguignon, P.-A. Chiappori, and V. Lechene (1994): "Income and Outcomes: A Structural Model of Intrahousehold Allocation," *Journal of Political Economy*, 102(6), 1067–1096.

Browning, M., and P.-A. Chiappori (1998): "Efficient Intra-household Allocations: A General Characterization and Empirical Tests," *Econometrica*, 66(6), 1241–1278.

Chiappori, P.-A. (1988): "Rational Household Labor Supply," *Econometrica*, 56(1), 63–89.

——— (1992): "Collective Labor Supply and Welfare," *Journal of Political Economy*, 100(3), 437–467.

Chiappori, P.-A., R. Blundell, and C. Meghir (2002): "Collective Labour Supply with Children," Unpublished Manuscript.

Chiappori, P.-A., B. Fortin, and G. Lacroix (2002): "Marriage Market, Divorce Legislation, and Household Labor Supply," *Journal of Political Economy*, 110(1), 37–72.

Deaton, A. (1997): *The Analysis of Household Surveys: A Microeconometric Approach to Development Policy*. World Bank, Baltimore.

Department of Employment (1968–2001): *Family Expenditure Survey*UK Data Archive, Essex, computer file.

Donni, O. (2001): "Collective Female Labor Supply: Theory and Application," Unpublished Manuscript.

——— (2003): "Collective Household Labor Supply: Nonparticipation and Income Taxation," *Journal of Public Economics*, 87, 1179–1198.

Haddad, L., and R. Kanbur (1990): "How Serious is the Neglect of Intra-household Inequality?," *The Economic Journal*, 100, 866–881.

Johnson, D., and S. Shipp (1997): "Trends in Inequality using Consumption-Expenditures: The U.S. from 1960 to 1993," *Review of Income and Wealth*, 43(2), 133–152.

Kanbur, R., and L. Haddad (1994): "Are Better Off Households More Unequal or Less Unequal?," *Oxford Economic Papers*, 46, 445–458.

Keane, M., and R. Moffitt (1998): "A Structural Model of Multiple Welfare program Participation and Labor Supply," *International Economic Review*, 39(3).

Krueger, D., and F. Perri (2003): "Does Income Inequality Lead to Consumption Inequality? Evidence and Theory," Unpublished Manuscript.

Lundberg, S., R. Pollak, and T. Wales (1997): "Do Husbands and Wives Pool Their Resources?," *The Journal of Human Resources*, 32(3), 463–480.

Manser, M., and M. Brown (1980): "Marriage and Household Decisionn Making: A bargaining Analysis," *International Economic review*, 21(1), 31–44.

National Statistics (1968–2001): *Inland Revenue Statistics*. Board of Inland Revenue.

——— (2003): *Focus on Consumer Price Indices*. National Statistics.

Pendakur, K. (1998): "Changes in Canadian Family Income and Family Consumption Inequality Between 1978 and 1992," *Review of Income and Wealth*, 44(2), 259–283.

Shorrocks, A. (1984): "Inequality Decomposition by Population Subgroups," *Econometrica*, 52(6), 1369–1386.

Train, K. (2003): *Discrete Choice Methods with Simulation*. Cambridge.

van Soest, A. (1995): "Structural Models of Family Labor Supply: A Discrete Choice Approach," *The Journal of Human Resources*, 30, 63–88.

Vermeulen, F. (2002): "Collective Household Models: Principles and Main Results," *Journal of Economics Surveys*, 16(4), 533–564.

——— (2003): "A Collective Model for Female Labour Supply with Nonparticipation and Taxation," Unpublished Manuscript.