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

First Version: September 15, 2003
This Version: April 21, 2004

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, increased marital sorting on earnings and the subsequent rise in the share of women's income in the household may have important implications 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 intra-household allocations underestimates cross sectional consumption inequality by 30% and overstates the trend by two-thirds. Our findings also indicate that increases in marital sorting on wages and hours worked can simultaneously explain virtually all of the decline in within household inequality and a substantial fraction of the rise in between household inequality.

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

1 Introduction

A large body of research aims to measure changes in the distribution of economic welfare. 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 (Cutler and Katz, 1992; Krueger and Perri, 2003; Blundell and Preston, 1998) for many reasons. 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 an accurate measure of well-being as compared to consumption.

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 equivalence 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 equivalence scales implies a very restrictive model of the household in which husbands and wives split consumption equally, regardless of the source of the income.

This criticism levelled against the use of adult equivalence scales may seem to be a

[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 versus 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 of inequality, not the trends.

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 sources of income within households changed over time, then adult equivalence scales may 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 three 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. Previous estimates of the collective models 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 regarding how consumption inequality *within* households relates to consumption inequality *across* households. The model is estimated using the UK Family Expenditure Survey (FES) for the years 1968 to 2001. We have two main findings. First, measures

[2]In fact, the unitary models are less restrictive than the model generating adult equivalence scales, as the unitary model does not require equal sharing, only that the allocation is independent of the source of income.

of consumption inequality that ignore the potential for intra-household inequality may underestimate individual-level inequality by 30% and overestimate the rise in consumption inequality by two-thirds. Second, the equal sharing implicit in adult equivalence scales is essentially valid *only* for households in which the wife has the same earnings as her husband. Together, the results highlight the importance of inequality *within* households as a major component of consumption inequality in the economy.

The third contribution of our paper is to provide evidence on the importance of several potential explanations for the rise in consumption inequality across households and the fall in inequality within households. While changes in the demographic composition of the population appear to play a limited role, an increase in marital sorting has profound effects on the trends in consumption inequality. In particular, the rise in marital sorting observed in the data has the potential to account for all of the fall in within household inequality and at the same can explain a large fraction of the rise in consumption inequality across households.

The remainder of the paper is organized as follows. Section 2 describes in detail the stylized facts on earnings and consumption inequality, wages, and labor supply that provide the motivation for our study. Section 3 outlines the theoretical framework and the identification strategy for estimating the rule to allocate consumption to individuals within a household. Section 4 describes the data set and the strategy for estimating the model. The estimation results and decompositions of consumption inequality are presented in Section 5. Section 6 considers the importance of several explanations for the trends in consumption inequality. Section 7 concludes.

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 the stylized facts. We are particularly interested in the following three features of the data.

1. There has been a large rise in earnings inequality at the individual level. The first panel in Figure 1 documents the trend in the Gini index for the distribution of individual earnings. In particular, the Gini index has risen by 12% over the past 30 years. This rise in earnings inequality has been well documented in the literature (e.g. Blundell and Preston, 1998).

2. Although the level of earnings inequality at the individual level is much higher than the level of earnings inequality at the household level, household inequality rose at a much faster rate: the Gini index for the latter rose by 41% between 1968 and 2001. What can account for this trend? The following evidence suggests sorting may be an important part of the answer.

3. The correlation between the earnings of husbands and wives increased dramatically over time. Figure 2 suggests the correlation between the earnings of husbands and wives increased dramatically over the sample period. This is due to both the fall in the gender wage gap and the rise in labor supply of women. Figure 3 graphs the wife's potential share of household earnings, measured as the female's wage divided by the sum of the female and male's wage in the household.[3] 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

[3]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 upon request.

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 labor force participation rates for husbands and wives in two person households. The gap in participation rates declined dramatically from 35.7 percentage points in 1968 to 1.5 percentage points in 2001. This change in participation rates represents a 95.9% decline in the gender gap in participation.

4. As reported by Blundell and Preston (1998), there has been a corresponding rise in consumption inequality. To account for economies of scale, we construct a standard measure of individual-level consumption by dividing total household consumption by the square root of household size. The Gini index for this measure of consumption is presented in the second panel of Figure 1. The level of income inequality is higher than the level of consumption inequality but the rise in inequality is higher for consumption than for earnings.[4]

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. What implication does this evidence have for our understanding of consumption inequality? In the next section, we present a model of intra-household allocations that will enable us to address this question.

3 Theoretical Framework

As illustrated in Section 2, the share of household income provided by wives has increased dramatically over the past 30 years. An extensive literature on intra-

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

household allocations suggests that the source of 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. Recent work by Browning, Chiappori, and Lewbel (2004) use a similar framework to study the allocation of consumption within households. The collective framework is also less restrictive than any particular form of bargaining, as the only restriction on the intra-household allocation 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. Labor supply is denoted l^g and the total time available to agents is normalized to one, i.e. $l^g = 1 - L^g$. The joint consumption of public goods is a primary gain to marriage and an important component in the

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

measurement of consumption inequality and is thus included here.[6] Denote total household non-labor income net of savings Y.[7] Labor earnings are denoted $w^g(l)$ and include any after tax income that depends directly on the labor supply decision. In particular, $w^g(l)$ includes unemployment insurance benefits paid to individuals who are not working. We construct labor earnings in this fashion as unemployment benefits are paid directly to one person in the household and likely affect allocations differently than does shared non-labor income. 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:

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

$$\text{subject to } w^g(l) + C^g + CP = Y.$$

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 that contains 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. As in Browning, Chiappori, and Lewbel (2004), 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)

[6] Chiappori, Blundell and Meghir (2002) establish conditions under which the collective model with public goods is identified but do not estimate the model.

[7] This measure of non-labor income is consistent a model with inter-temporal substitution, time separable preferences and an exogenous wage profile, as in the two-stage budgeting process of Deaton and Muelbauer (1980).

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:

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

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

The Pareto weight, λ, represents the female's bargaining power within the household, and will typically be a function of full-time labor income $w^f(1), w^m(1), Y$ 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 (see Chiappori, Blundell, and Meghir (2002) for details). 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:

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

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

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 in (2) produce identical labor supplies and consumption demands, under the assumption 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

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.[8] In particular, we assume that individuals can choose from H 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 is likely important in explaining the long-term trends in consumption inequality.[9] Further assume that L^f, L^m, Y, $w^f(l)$, and $w^m(l)$ 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.

The question we aim to address in this paper is whether measures of consumption inequality from the collective model differ from current measures in the literature based on equivalence scales. 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

[8] 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.

[9] This assumption is not necessary for identification.

additional restriction on preferences. As in Vermeulen (2003) and Browning, Chiappori, and Lewbel (2004), 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 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.[10] The utility an individual of gender g derives 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 ε_h^g is an unobserved preference component that is assumed to be distributed *iid* across individuals and labor supply alternatives. 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 couples and singles.[11]

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.[12] 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

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

[11] 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.

[12] 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.

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:

$$C_h^g = w^g(l) + y \tag{3}$$

for single individuals,

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

for married women and

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

for married men.

Only differences in utility between labor supply choices matter in the model; thus the parameters must be estimated relative to a base case. 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:

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

where $\tilde{w}^g(l_h) = w^g(l_h) - w^g(l_0)$ and $[\tilde{w}^g(l_h)]^2 = [w^g(l_h)]^2 - [w^g(l_0)]^2$. 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:

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

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:

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

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.[13]

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

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

where m is an indicator 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, the sharing rule, and the assumption that preferences for consumption are the same regardless of marital status:

$$\begin{aligned} \phi_0 &= \frac{\Pi_{lym}^f}{\Pi_{ly}^f} + 1 = \frac{\Pi_{wlym}^f}{\Pi_{wly}^f} + 1 = -\frac{\Pi_{lym}^m}{\Pi_{ly}^m} = -\frac{\Pi_{wlym}^m}{\Pi_{wly}^m}, \\ \phi_k &= \frac{\Pi_{z_k lym}^f}{\Pi_{ly}^f} = \frac{\Pi_{z_k wlym}^f}{\Pi_{wly}^f} = -\frac{\Pi_{z_k lym}^m}{\Pi_{ly}^m} = -\frac{\Pi_{z_k wlym}^m}{\Pi_{wly}^m}, \quad k = 1 \ldots K. \end{aligned} \tag{9}$$

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

4 Empirical Specification

4.1 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, on wages and labor supply for individuals within households,

[13] 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.

and on demographics 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 elsewhere.[14] 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.[15]

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 and the self-employed. Households in which one of the individuals is in the top one per cent of the wage distribution are also excluded. The resulting sample is composed of $87,668$ individuals.[16] 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. Other income is defined as total household expenditures minus net labor income. We use this expenditure based definition of non-labor income, as it is consistent with the assumptions of a two stage budgeting process, time separable preferences, and separability of public goods consumption from leisure and private consumption as in the model.[17]

[14]Battistin (2003) documents reporting errors in the US Consumer Expenditure Survey due to survey design.

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

[16]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.

[17]In estimation, household expenditures on public goods are subtracted 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 decision, including zero hours, we need to assign an earnings level to all individuals. For those who are working we use the usual hourly wage, defined as weekly earnings divided by usual weekly hours. For non-participants we use a predicted wage, computed based on a reduced form selection-corrected wage equation.[18] After tax earnings are subsequently computed 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. The resulting income 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, $z_1 = w^f/(w^f + w^m)$. Individuals may also be entitled to income related to earnings when working zero hours, for instance unemployment benefits, thus we also predict unemployment benefits for those who are working based on the *Official Yearbook of the United Kingdom (1968-2001).*

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 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 convention. If an individual is observed to be working either part-time or full-time we use the reported number of hours to measure labor supply and usual take home

[18]The log of the wage is estimated as a function of age, birth cohort, year, quarter, and regional dummies, 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 non-labor income, marital status, and the age, education, and the labor income of the spouse. Results are available upon request.

pay in constructing the consumption. In cases for which we do not observe the labor supply state, we calculate after tax earnings based on 20 hours for the part-time choice and 40 hours for the full-time choice. Constructing the individual consumption in this way ensures our measure of total private consumption in the household is consistent with that observed 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 (Seitz, 2003), 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, and the age gap between spouses as distribution factors in estimation.

4.2 Econometric Specification

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 type I extreme value. 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 estimation, heterogeneity in preferences for leisure is introduced through the vector X which includes age, birth cohort, education, region, and quarter and year to control for seasonality and cyclical effects.[19] The parameters β_l and β_{ll} are assumed

[19]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.

to be linear functions of the observed characteristics, 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 observed characteristics and β is a vector of parameters.

Estimation proceeds in two steps. First, we estimate a standard selection corrected wage equation and predict wages for individuals that are not working. Second, we estimate the discrete labor supply choice, treating wages as known.

5 Results

We begin with estimates of the sharing rule parameters, the parameters that allow us to infer the share of consumption attributed to each adult in the household. As discussed in Section 3.3, with quadratic utility and under 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 estimates of the unrestricted reduced form. The unrestricted reduced form estimates for two specifications of the model are presented in Table 2. The first column of the Table presents estimation results from the case in which the only distribution factor is the share of women's potential earnings in household potential earnings. The second column presents results from a model where a second distribution factor, the age difference between spouses, is included.

The estimated sharing rule parameters constructed from the different model restrictions described by Equation 9 are qualitatively similar across gender, remarkably so for men. In both specifications, the positive sign on ϕ_1 indicates that an increase in the female's share of potential earnings increases her share of total consumption in the household. The negative sign on ϕ_2 suggests that the share of consumption women receive is decreasing in the relative age of their husbands. The sharing rule parameters for the second set of restrictions in Equation 9 are larger in absolute value for both the intercept and the distribution factors. Upon closer examination of the

reduced form results, we find the reason for this difference across the estimates is due primarily to the fact that the estimated value of the denominator, Π^f_{wly}, is relatively small. This parameter captures the effect of the interaction between non-labor income and earnings for women on the labor supply decision. Since many women are not working, we need to impute earnings for 39% of the women in the data. Most of the information used to predict wages is also included directly in the reduced form model for hours; as a result, the predicted wage includes very little information. As a result, the parameter estimate is likely biased towards zero.[20] It should be noted that this set of restrictions is less precise; as a result, it has less weight in the minimum distance estimation used to obtain the sharing rule estimates as discussed below.

The test statistics associated with several tests of the model restrictions are presented in the bottom four rows of Table 2. A Wald test on the model with one distribution factor firmly rejects the full set of restrictions. The Wald test on the full set of restrictions from the model with two distribution factors, however, suggests the model is not strongly rejected. We subsequently test whether the sharing rule parameters estimated from the restrictions within gender are the same. The test statistics, presented in Rows 2 and 3 of the bottom panel of Table 2, indicate the within gender restrictions are not rejected by the data. We also test whether each of the individual restrictions from the female's problem are the same as the corresponding restrictions from the male's problem in Rows 4 and 5. In each case, the test statistics indicate the model restrictions are not rejected at conventional significance levels. Overall, the test statistics provide support for our version of the collective model.

We next compute consistent estimates of the sharing rule parameters from the unrestricted estimates by minimizing the distance between the reduced form and structural parameters, using the estimated covariance matrix from the reduced form to construct the weighting matrix. The results of this exercise are presented in Table 3.[21] The estimates suggest that a 10% increase in the share of potential earnings

[20] We are currently estimating the wage equation and hours decision jointly, which should reduce this problem to some extent.

[21] See the Appendix for further details on the minimum distance estimation exercise.

attributed to the wife results in a 16% increase in the share of non-labor income she receives. This result is consistent with an increase in the wife's threat point within a bargaining model. The estimate of ϕ_2 indicates that an increase in the husband's age by 1 year results in a 0.4% decrease in the wife's share of non-labor income. While small in magnitude, this finding suggests that older spouses tend to have more bargaining power in marriage.

One of the main goals of this paper is to determine whether measures of consumption inequality using standard adult equivalence scales provide an accurate estimate of consumption inequality across individuals. Recall, adult equivalence scales typically assume that husbands and wives share in household consumption equally. It is of interest to determine under what conditions our model would yield the same measures of consumption inequality as measures using adult equivalence scales. To this end, we determine under what conditions does equal sharing arises in our model. Assuming husbands are two years older than their wives, we use the sharing rule estimates to determine what value of the female's share in potential household earnings satisfies:

$$\frac{1}{2} = \hat{\phi}_0 + \hat{\phi}_1 \cdot \frac{w^f}{w^f + w^m} + \hat{\phi}_2 \cdot 2.09.$$

Using estimates for ϕ_0, ϕ_1, and ϕ_2 of -0.31, 1.59, and -0.004 respectively yields 51%. In other words, the model predicts that consumption is split equally across the husband and wife when they have the approximately same earnings![22] It is worth emphasizing that this 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 individual's preferences over private consumption goods are the same when single as when married.

[22] To be precise, husbands and wives will split consumption equally when they have approximately the same wages and hours.

5.1 Sensitivity Analysis

In this section, we consider whether our results are robust to several modifications. The first robustness check we consider is whether the results are sensitive to our definitions of public and private consumption. We first estimate the model under the assumption that there are no public goods and then sequentially add housing, heat and lighting, household durables, transport and services to public good consumption.[23] The results of this exercise are reported in Table 4. With the exception of the zero public goods case, the parameter estimates are quite robust across specifications: an increase in the wife's share of potential household earnings of 10% results in an increase in her consumption share of between 12% and 17%. Under the most restrictive assumption, that no goods are public in the household, the model predicts women receive 40% of the consumption in households where both spouses choose the same hours of work. As the fraction of public goods in household expenditures increases, women receive a greater share of consumption. This result reflects, in part, the fact that a larger portion of consumption in the household is public and thus split equally across spouses.

The next specification we estimate allows for differences in the sharing rule parameters for each birth cohort in our pooled sample. The sample covers a long time period and a wide age range in every year; we thus estimate sharing rules for each ten-year cohort in the data. The parameter estimates are presented in Column 1 and Columns 2 in Table 5 for the models with one distribution factor and two distribution factors, respectively. With the exception of the 1900 and 1960 birth cohorts (which are relatively small samples), the parameter estimates and the predicted share of consumption assigned to wives when earnings are equal across spouses are quite similar across the cohorts. For the cohorts between 1910 and 1950, an increase in the wife's share of potential household earnings increases her share of non-labor income between 13% and 23% and the estimated effect of an increase in the husband's age by one year fall within the range of -0.6% and 0.5%. The fact that the sharing rule

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

parameter estimates are quite similar across specifications is surprising considering the large changes in divorce costs and gains to marriage over time.

The final robustness check we perform is to add unobserved heterogeneity in preferences to the model for two reasons. First, we want to allow for the possibility that the preferences shocks are correlated across labor supply choices. Second, we to allow for additional flexibility in estimating preferences over leisure. We specify $\beta_l = X_i\beta_l + u_{hi}$ and $\beta_{ll} = X_i\beta_{ll} + u_{hhi}$, 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. Results from this specification for the model with one distribution factor are presented in Column 3 of Table 5. Incorporating unobserved preference heterogeneity appears to reduce both ϕ_0 and ϕ_1 slightly but does not change the implications of the model. In particular, the effect of a 10% increase in potential household earnings attributed to wives results an increased transfer of between 11% and 21% for the 1910 to 1950 cohorts, which is close to the range reported for Column 1 above.

5.2 Sharing Rule Estimates of Consumption Inequality

In this section, we compare the inequality measure implied by our model to a conventional measure of consumption inequality. For the purposes of this analysis, we use estimates of the model with two distribution factors and no unobserved preference heterogeneity to construct our benchmark sharing rule.[24] Next, we use this sharing rule to divide non-labor income between the husband and wife in each household. We subsequently construct private consumption based on the individuals' share of non-labor income and his or her personal net labor earnings, where private consumption is constructed as in equations (4) and (5). Our *sharing rule* measure of individual consumption, for those in married couples, is then equal to individual private

[24] See Column 2 of Table 3.

consumption, plus household public consumption

$$\begin{aligned} C^f &= CP + w^f(l_h) + [\hat{\phi}_0 + \hat{\phi}_1 \frac{w^f(1)}{w^f(1) + w^m(1)} + \hat{\phi}_2(age_m - age_f)] \cdot y \\ C^m &= CP + w^m(l_h) + [1 - \hat{\phi}_0 - \hat{\phi}_1 \frac{w^f(1)}{w^f(1) + w^m(1)} - \hat{\phi}_2(age_m - age_f)] \cdot y. \end{aligned}$$

Single individuals simply consume their entire labor and non-labor income. 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. In 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 equivalence scale to assign household consumption to individual members.[25] 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 also compute the Mean Logarithmic Deviation (MLD) in consumption, defined as

$$I_\alpha(C) = \frac{1}{n} \sum_{i=1}^{n} \log \left(\frac{\mu_C}{C_i^g} \right). \tag{10}$$

The index of total inequality using the MLD can be additively decomposed into within and between household inequality:

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

[25]It should be noted that the correlation between our equal division consumption inequality measure and a measure of inequality using equivalence scales is 0.99.

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 total inequality index $I_\alpha^T(C)$. We can then recover intra-household inequality as

$$I_\alpha^W(C) = I_\alpha^T(C) - I_\alpha^B(C). \tag{11}$$

Using Equation (11) we calculate the inequality index using both the sharing rule and the equal division constructions of individual consumption. We then recover the index of within household inequality as the difference between the two MLD indices.

The time-series trend of total and between household household inequality for the years 1968 to 2001 is presented in Figure 8. The Gini index measures are presented in the first panel, and the MLD measures of inequality are presented in the second panel. Inequality was 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, ignoring intra-household inequality underestimates consumption inequality in 1968 by approximately 30% and 15% when inequality is measured in using the MLD and Gini index, respectively. The large magnitude of these differences in consumption inequality measures highlights the need for further study of this issue.

Second, the rise in consumption inequality under equal division, or between household inequality, may be over-stated by as much as 65%, as illustrated by the trend in the MLD presented in Figure 9. The reason the sharing rule measure of inequality differs so markedly from the equal division measure is due to the fall in within-household inequality. In particular within-household inequality fell by 15% between 1968 and 2001.[26] The stylized facts presented in Section 2 allude to possible reasons for the

[26]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

decline in within household inequality, such as the fall in the gender wage gap and the rise in female labor supply. In the next section, we assess the importance of these and other explanations for the trends in consumption inequality over time.

6 What accounts for the trends in consumption inequality?

In this section, we examine several explanations for the rise in consumption inequality observed in the data. We conduct a series of thought experiments to illustrate the potential importance of each explanation. Some of our thought experiments entail changing the distribution of education over time. Unfortunately, the education levels of respondents is only available beginning in 1978, so we focus on the years 1978 to 2001. This should not have a large impact on our conclusions, as the major changes in inequality occurred over the 1980s. Results are presented in Table 6. The first panel of the Table presents the values of the Gini index and MLD in consumption for various measures of consumption inequality. The first two rows contain the benchmark inequality measures for 1978 and 2001. Subsequent rows present the inequality measures under various scenarios discussed in detail below. The second panel of the Table presents the percentage change in the observed inequality measures attributed to each explanation we consider.

The first hypothesis we consider is that the rise in consumption inequality is capturing cohort effects due to the changes in the age structure of the population. Most inequality measures, such as those presented in this paper, are based on repeated cross-sectional data, which 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). To assess the importance of the changing age structure in the population for our measures of consumption inequality,

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.

we re-weight the 1978 data so that the age structure is the same as that in 2001, holding all else constant. The implications of this experiment are presented in the third row of the top panel of Table 6. Surprisingly, the results of this exercise suggest that changes in the age distribution between 1978 and 2001 had virtually no effect on consumption inequality.

The second explanation we consider is the large change in household composition that occurred alongside the rise in inequality. In particular, with delays in marriage and a rise in divorce rates, the fraction of households with one adult increased relative to the fraction of two adult households. Although single person households have no within household inequality by definition, it is still the case that there may exist substantial inequality across single adult households. To assess the importance of changing household composition, we re-weight the 1978 data so that the fraction of married couples, the fraction of single women, and the fraction of single men in the population to match the proportions in 2001. The results of this exercise, suggest that the change in household composition can explain up to 30% of the change in household inequality according to the sharing rule estimates and 22% of the change in consumption inequality when measured using adult equivalence scales. Together, a combination of a changing age distribution and the change in household composition over time can explain a little more than 30% of the change in the Gini index over time, most of this effect coming through household composition.

The next set of experiments we consider involve changes in the distribution of educational attainment over time. In our sample the average age at which individuals completed full time schooling rose from 15 in 1978 to 17 in 2001. The results of these experiments suggest that changes in the distribution of education levels over time are not a major factor in explaining the trends in total or between consumption inequality. In fact, the results suggest that total consumption inequality would have fallen if the only change over time was sorting on education. Interestingly, accounting for the change in sorting over time does play a big role in explaining the fall in within household inequality. Because of the rise in sorting within households, women had

levels of education more similar to their spouses, which served to increase both their market wages and their share of potential income in the household. Considering the correlation between the rise in sorting on education and the fall in intra-household inequality, we next assess the role of wages, hours, and sorting in further detail.

According to the stylized facts, two of the most salient trends over time are the closing of the gender gap in wages and the rise in female labor supply. We next assess the role of both factors in accounting for the rise in consumption inequality across households and the fall in consumption inequality within households. We consider both the effects of changes in the wage and hours distributions on aggregate as well as increases in sorting on wages and hours as observed in the data. In the first exercise, we re-weight the wage distribution by gender so that the histogram of log wages is the same in 1978 as in 2001.[27] In the second experiment, we re-weight the joint spousal distribution of wages in the same fashion. Both experiments are subsequently repeated for labor supply. The results are presented in the bottom four rows of both panels of Table 6.

What is really interesting about the results on wage and hours sorting is that they can simultaneously explain both the rise in consumption inequality across households and the fall in consumption inequality within households: sorting on wages alone can explain approximately 40% of the rise between household inequality and 78% of the fall in within household inequality. With respect to sorting on hours, approximately one-third of the rise in between household inequality and virtually all of the fall in within household inequality can be explained by increased sorting within marriage. Regardless of the measure of consumption inequality considered, the thought experiments conducted above illuminate the dramatic role of sorting in determining distribution of consumption across individuals. These results are complementary to those of Fernández and Rogerson (2001), among others, on sorting and income inequality and suggest an important avenue for future research.

[27]The histograms used to re-weight the wage distributions have 10 bins each.

7 Conclusions

Our paper makes three 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 30% and may overstate the rise in individual consumption inequality by 65%. Most importantly, the results of our analysis highlight the importance of intra-household allocations for our understanding of consumption inequality and its implications. Finally, we consider several potential explanations for the rise in consumption inequality across households and the fall in consumption inequality within households. The results indicate that changes in sorting on education, wages and hours play prominent roles in explaining the inequality trends over time.

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.

Blundell, R., and I. Preston (1998): "Consumption Inequality and Income Uncertainty," *Quarterly Journal of Economics*.

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.

Browning, M., P.-A. Chiappori, and A. Lewbel (2003): "Estimating Consumption Economies of Scale, Adult Equivalence Scales, and Household Bargaining Power," Unpublished Manuscript.

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.

Crossley, T., and K. Pendakur (2002): "Consumption Inequality," Unpublished Manuscript.

Cuttler, D., and L. Katz (1992): "Rising Inequality? Changes in the Distribution of Income and Consumption in the 1980's," *AEA Papers and Proceedings*, 82(2).

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

Deaton, A., and J. Muellbauer (1980): *Economics and Consumer Behavior*. Cambridge University Press.

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.

Fernández, R., and R. Rogerson (2001): "Sorting and Long-Run Inequality," *Quarterly Journal of Economics*, 116(4), 1305–41.

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 Decision Making: A bargaining Analysis," *International Economic review*, 21(1), 31–44.

McElroy, M., and M. Horney (1981): "Nash-Bargained Household Decisions: Toward A Generalization of the Theory fo Demand," *International Economic Review*, 22(2).

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

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

Office for National Statistics (1968–2001): *The Official Yearbook of the United Kingdom*.

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.

Seitz, S. (2003): "Accounting for Racial Differences in Marriage and Employment," Unpublished Manuscript.

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.

A Minimum Distance Estimator of Structural Parameters

The structural parameters

$$\theta = \left(\phi_0, \phi_1, \phi_2, \beta^f_{cl}, \beta^f_{cc}, \beta^m_{cl}, \beta^m_{cc}, \beta^f_{lX}, \beta^m_{lX}\right)^\top.$$

can be consistently estimated by using a minimum distance estimator (MDE) (see Chamberlain (1984)). We define the MDE as

$$\hat{\theta} = \arg\min_{\theta} \left(\hat{\Pi} - f\left(\theta\right)\right)^\top V^{-1} \left(\hat{\Pi} - f\left(\theta\right)\right),$$

where the function f imposes the structural restrictions on the reduced form, and V is the covariance matrix of the reduced form parameter estimates. For the case in which the sharing rule is a linear function of three distribution factors the structure of the model implies the following restrictions on the reduced form parameters:

$$\left(\hat{\Pi} = f(\theta)\right) \Longrightarrow \left(\begin{array}{ccc}
\hat{\Pi}^f_{ly} & = & \beta^f_{cl} \\
\hat{\Pi}^f_{lym} & = & \beta^f_{cl}\left(\phi_0 - 1\right) \\
\hat{\Pi}^f_{z_1 lym} & = & \beta^f_{cl}\phi_1 \\
\hat{\Pi}^f_{z_2 lym} & = & \beta^f_{cl}\phi_2 \\
\hat{\Pi}^f_{wly} & = & \beta^f_{cc} \\
\hat{\Pi}^f_{wlym} & = & \beta^f_{cc}\left(\phi_0 - 1\right) \\
\hat{\Pi}^f_{z_1 wlym} & = & \beta^f_{cc}\phi_1 \\
\hat{\Pi}^f_{z_2 wlym} & = & \beta^f_{cc}\phi_2 \\
\hat{\Pi}^m_{ly} & = & \beta^m_{cl} \\
\hat{\Pi}^m_{lym} & = & -\beta^m_{cl}\phi_0 \\
\hat{\Pi}^m_{z_1 lym} & = & -\beta^m_{cl}\phi_1 \\
\hat{\Pi}^m_{z_2 lym} & = & -\beta^m_{cl}\phi_2 \\
\hat{\Pi}^m_{wly} & = & \beta^m_{cc} \\
\hat{\Pi}^m_{wlym} & = & -\beta^m_{cc}\phi_0 \\
\hat{\Pi}^m_{z_1 wlym} & = & -\beta^m_{cc}\phi_1 \\
\hat{\Pi}^m_{z_2 wlym} & = & -\beta^m_{cc}\phi_2 \\
\hat{\Pi}^f_{lX} & = & \beta^f_{lX} \\
\hat{\Pi}^m_{lX} & = & \beta^m_{lX}
\end{array}\right),$$

$\hat{\theta}$ is distributed asymptotically normal as:

$$\sqrt{n}\left(\hat{\theta} - \theta\right) \rightarrow_d N\left(0, \left(G^\top V^{-1} G\right)^{-1}\right),$$

where $G(\theta) = \frac{\partial f(\theta)}{\partial \theta^\top}$.

Table 1: Descriptive Statistics from the FES

	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		919	
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	56.21	47.77
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: Unrestricted Estimates of the Sharing Rule.

	One Distribution Factor		Two Distribution Factors	
$\phi_0^{fcl} = \Pi^f_{lym}/\Pi^f_{ly} + 1$	−0.534	***	−0.517	***
	(0.056)		(0.055)	
$\phi_0^{fcc} = \Pi^f_{wlym}/\Pi^f_{wly} + 1$	−1.707	***	−1.634	***
	(0.536)		(0.520)	
$\phi_0^{mcl} = -\Pi^m_{lym}/\Pi^m_{ly}$	−0.393	***	−0.388	***
	(0.089)		(0.088)	
$\phi_0^{mcc} = -\Pi^m_{wlym}/\Pi^m_{wly}$	−0.327	**	−0.329	**
	(0.160)		(0.158)	
$\phi_1^{fcl} = \Pi^f_{z_1lym}/\Pi^f_{ly}$	2.297	***	2.300	***
	(0.179)		(0.178)	
$\phi_1^{fcc} = \Pi^f_{z_1wlym}/\Pi^f_{wly}$	5.796	***	5.794	***
	(1.687)		(1.691)	
$\phi_1^{mcl} = -\Pi^m_{z_1lym}/\Pi^m_{ly}$	1.714	***	1.765	***
	(0.175)		(0.175)	
$\phi_1^{mcc} = -\Pi^m_{z_1wlym}/\Pi^m_{wly}$	1.635	***	1.726	***
	(0.342)		(0.341)	
$\phi_2^{fcl} = \Pi^f_{z_2lym}/\Pi^f_{ly}$			−0.009	***
			(0.003)	
$\phi_2^{fcc} = \Pi^f_{z_2wlym}/\Pi^f_{wly}$			−0.038	**
			(0.019)	
$\phi_2^{mcl} = -\Pi^m_{z_2lym}/\Pi^m_{ly}$			−0.011	***
			0.004	
$\phi_2^{mcc} = -\Pi^m_{z_2wlym}/\Pi^m_{wly}$			−0.015	**
			(0.009)	

Tests	df	χ^2	p-value	df	χ^2	p-value
$\phi^f = \phi^m$	6	19.49	0.003	9	21.73	0.010
$\phi^{fcl} = \phi^{fcc}$	2	5.80	0.055	3	6.27	0.099
$\phi^{mcl} = \phi^{mcc}$	2	1.05	0.592	3	1.92	0.590
$\phi^{fcl} = \phi^{mcl}$	2	8.65	0.013	3	8.26	0.041
$\phi^{fcc} = \phi^{mcc}$	2	6.09	0.048	3	5.81	0.121

Note: The sharing rule has the form: $\phi = \phi_0 + \phi_1 \left(\frac{w^f}{w^f + w^m}\right) + \phi_2(age^m - age^f)$. Each sharing rule parameter (ϕ_0, ϕ_1, and ϕ_2) can be recovered from the restrictions on the reduced form estimates in equation 9. Standard errors in parentheses. ***, **, and * indicate the coefficient is statistically different from zero at the 10%, 5% and 1% significance levels, respectively

Table 3: Minimum Distance Sharing Rule Estimates

	(1)	(2)
ϕ_0	-0.317	-0.310
	(0.021)	(0.021)
ϕ_1	1.584	1.592
	(0.058)	(0.057)
ϕ_2		-0.004
		(0.001)
$\phi_0 + \frac{1}{2}\phi_1$	0.475	0.485
	(0.012)	(0.013)

Note: Standard errors in parentheses.

Table 4: Sharing Rule Estimates Under Alternative Measures of Public Goods

	No Public Goods	Public Goods (i) (housing)
ϕ_0	0.008	−0.186
	(0.038)	(0.028)
ϕ_1	0.773	1.220
	(0.091)	(0.073)
$\phi_0 + \frac{1}{2}\phi_1$	0.395	0.424
	(0.018)	(0.015)
% of Total Consumption	0%	17%
	Public Goods (ii) (i + heat)	Public Goods (iii) (ii + durables)
ϕ_0	−0.239	−0.317
	(0.027)	(0.015)
ϕ_1	1.352	1.584
	(0.071)	(0.013)
$\phi_0 + \frac{1}{2}\phi_1$	0.437	0.475
	(0.015)	(0.010)
% of Total Consumption	24%	33%
	Public Goods (iv) (iii + transport)	Public Goods (v) (iv + services)
ϕ_0	−0.310	−0.256
	(0.013)	(0.008)
ϕ_1	1.665	1.562
	(0.038)	(0.025)
$\phi_0 + \frac{1}{2}\phi_1$	0.523	0.525
	(0.010)	(0.007)
% of Total Consumption	47%	56%

Note: Standard errors in parentheses.

Table 5: Minimum Distance Sharing Rule Estimates by Birth Cohort.

Birth Cohort		(1)	(2)	(3)	Birth Cohort		(1)	(2)	(3)
1900	ϕ_0	−0.099	−0.173	−0.260	1940	ϕ_0	−0.138	−0.152	−0.001
$N: 3,134$		(0.11)	(0.09)	(0.11)	N: 15,284		(0.066)	(0.066)	(0.052)
	ϕ_1	0.470	0.778	1.008		ϕ_1	1.474	1.487	1.223
		(0.32)	(0.25)	(0.27)			(0.154)	(0.156)	(0.117)
	ϕ_2		−0.002			ϕ_2		0.005	
			(0.007)					(0.004)	
	$\phi_0 + \frac{1}{2}\phi_1$	0.136	0.216	0.244		$\phi_0 + \frac{1}{2}\phi_1$	0.599	0.591	0.611
		(0.061)	(0.050)	(0.041)			(0.033)	(0.033)	(0.023)
1910	ϕ_0	−0.379	−0.347	−0.139	1950	ϕ_0	−0.361	−0.348	−0.274
N: 12,211		(0.053)	(0.050)	(0.054)	N: 11,692		(0.093)	(0.087)	(0.078)
	ϕ_1	1.799	1.650	1.135		ϕ_1	1.738	1.746	1.462
		(0.154)	(0.143)	(0.122)			(0.240)	(0.228)	(0.191)
	ϕ_2		0.005			ϕ_2		−0.004	
			(0.003)					(0.005)	
	$\phi_0 + \frac{1}{2}\phi_1$	0.520	0.478	0.428		$\phi_0 + \frac{1}{2}\phi_1$	0.507	0.525	0.457
		(0.034)	(0.033)	(0.021)			(0.048)	(0.049)	(0.033)
1920	ϕ_0	−0.591	−0.568	−0.471	1960	ϕ_0	0.188	0.167	0.290
N: 18,660		(0.058)	(0.052)	(0.030)	N: 7,974		(0.153)	(0.158)	(0.169)
	ϕ_1	2.325	2.303	2.067		ϕ_1	0.574	0.590	0.290
		(0.171)	(0.154)	(0.056)			(0.332)	(0.336)	(0.370)
	ϕ_2		−0.006			ϕ_2		0.007	
			(0.003)					(0.008)	
	$\phi_0 + \frac{1}{2}\phi_1$	0.572	0.583	0.563		$\phi_0 + \frac{1}{2}\phi_1$	0.475	0.462	0.435
		(0.036)	(0.034)	(0.010)			(0.056)	(0.056)	(0.050)
1930	ϕ_0	−0.195	−0.204	−0.154	1970	ϕ_0	−0.147	−0.193	0.017
N: 16,219		(0.056)	(0.041)	(0.033)	N: 2,585		(0.184)	(0.264)	(0.179)
	ϕ_1	1.248	1.330	1.071		ϕ_1	0.997	1.210	0.703
		(0.137)	(0.122)	(0.085)			(0.465)	(0.613)	(0.455)
	ϕ_2		−0.000			ϕ_2		0.295	
			(0.003)					(0.059)	
	$\phi_0 + \frac{1}{2}\phi_1$	0.429	0.461	0.382		$\phi_0 + \frac{1}{2}\phi_1$	0.352	0.412	0.369
		(0.052)	(0.031)	(0.016)			(0.089)	(0.130)	(0.098)

Note: Standard errors in parentheses. N indicates the sample size for each birth cohort. Column 3 contains preliminary estimates allowing for unobserved preference heterogeneity for leisure based on 100 random draws. The covariance matrix of the reduced form estimates is based on the numerical Hessian for column 1 and 2 and on the outer product of the gradient for column 3.

Table 6: Decomposition of the Change in Between and Within Consumption Inequality.

	Gini Index		Mean Logarithmic Deviation		
Absolute Change	Total	Between	Total	Between	Within
1978	0.332	0.285	0.197	0.137	0.060
2001	0.372	0.337	0.255	0.204	0.050
Age Distribution	0.333	0.285	0.197	0.137	0.060
Household Composition	0.344	0.296	0.207	0.148	0.059
Age and Household	0.345	0.302	0.207	0.153	0.054
Education	0.333	0.286	0.195	0.138	0.057
Educational Sorting	0.329	0.288	0.190	0.140	0.050
Wages	0.340	0.293	0.202	0.143	0.059
Wage Sorting	0.353	0.311	0.216	0.163	0.053
Labor Supply	0.348	0.304	0.211	0.155	0.055
Labor Supply Sorting	0.348	0.308	0.209	0.159	0.051

	Gini Index		Mean Logarithmic Deviation		
Percentage Change	Total	Between	Total	Between	Within
1978 to 2001 Change	0.040	0.052	0.057	0.067	−0.010
Age Distribution	0.4	0.2	−1.0	0.0	5.7
Household Composition	29.2	21.5	17.1	16.4	12.3
Age and Household	30.7	32.5	16.3	23.1	63.3
Education	2.5	3.1	−3.7	2.3	37.4
Educational Sorting	−7.8	6.1	−13.3	3.8	104.3
Wages	18.2	15.2	8.3	9.4	15.8
Wage Sorting	51.7	50.8	32.3	38.9	78.1
Labor Supply	39.6	36.3	23.4	27.1	48.4
Labor Supply Sorting	40.4	44.2	20.8	32.1	98.3

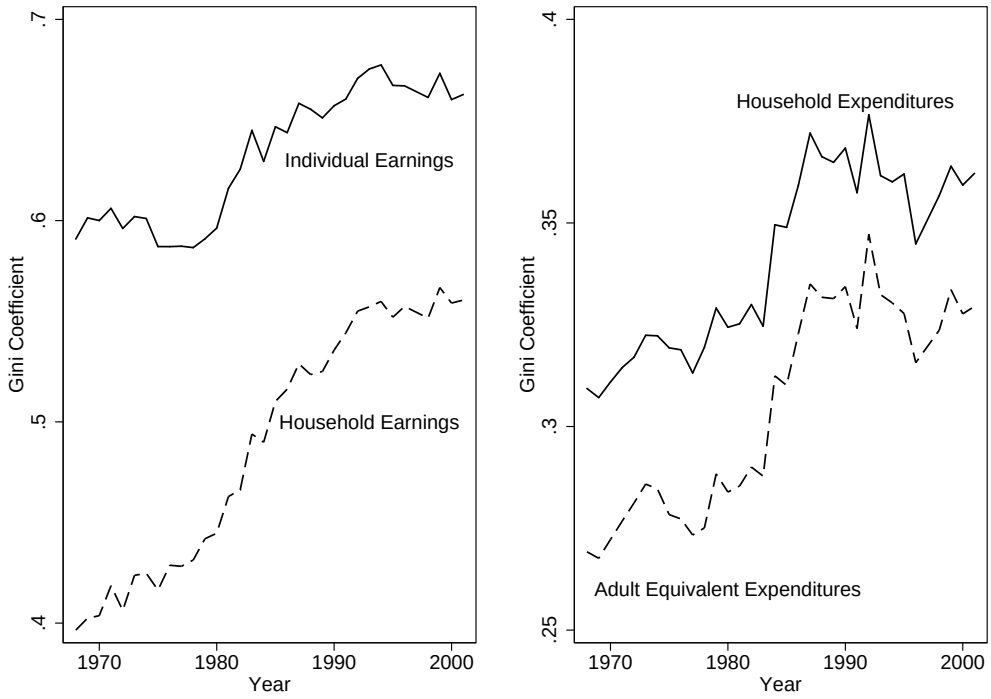

Figure 1: **Trends in the Gini index for earnings.**
Own calculations from the FES.

Figure 2: **Correlation in earnings across husbands and wives.**
Own calculations from the FES.

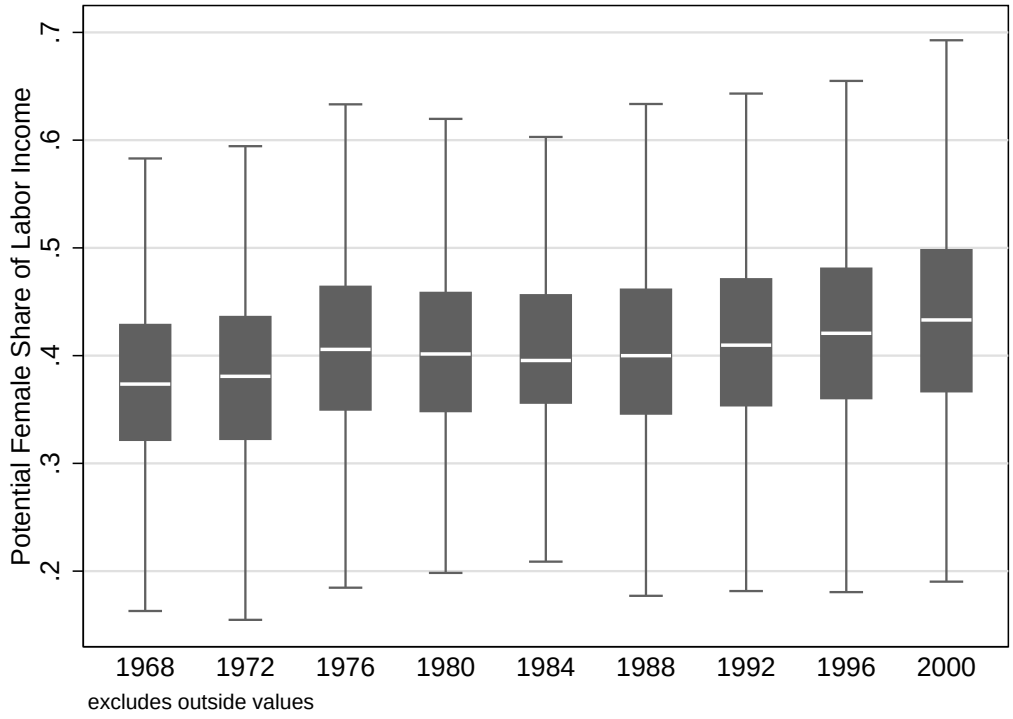

Figure 3: **Fraction of potential household earnings provided by wife.**
Source: Own calculations from the FES.

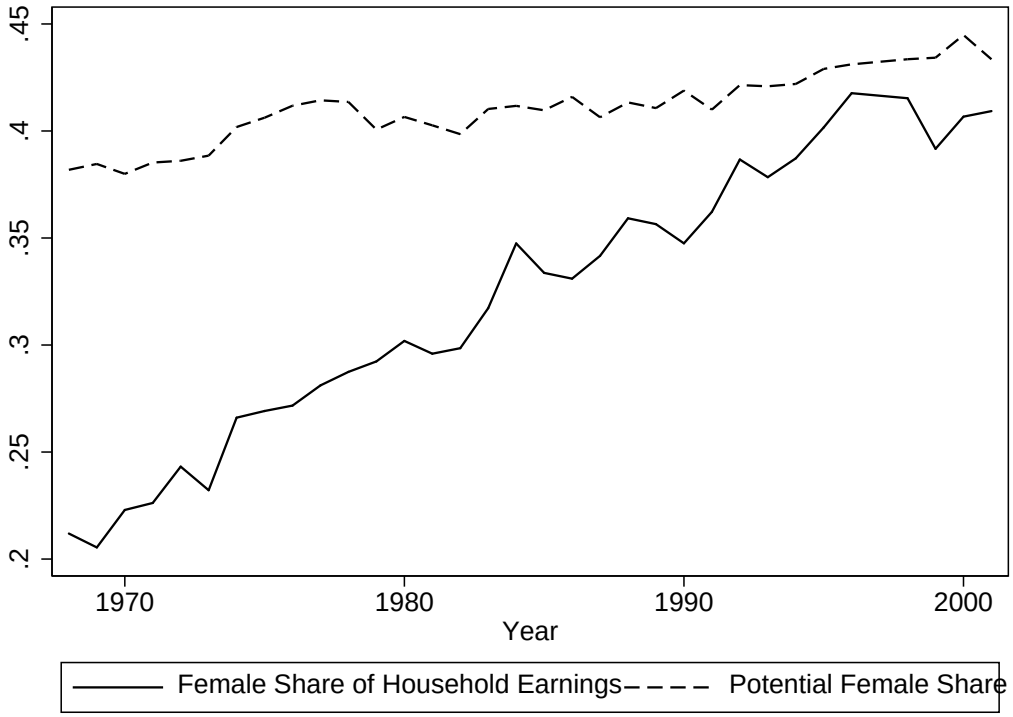

Figure 4: **Fraction of actual household earnings provided by wife.**
Source: Own calculation from the FES.

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

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

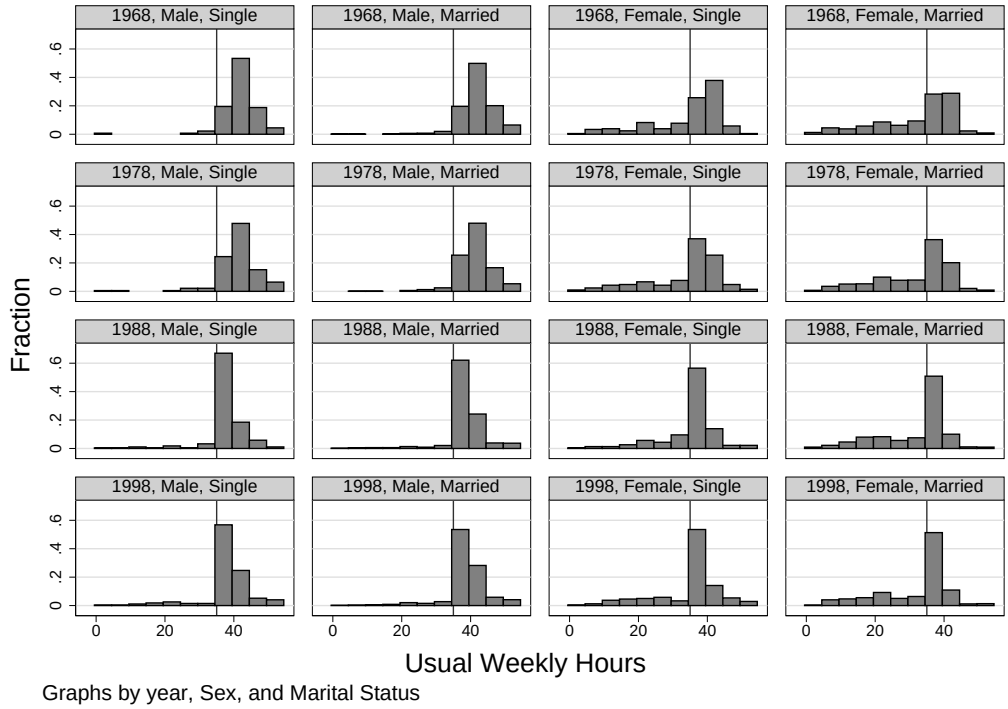

Figure 7: **Histogram of usual weekly hours for participants.**
Source: Own calculation from the FES.

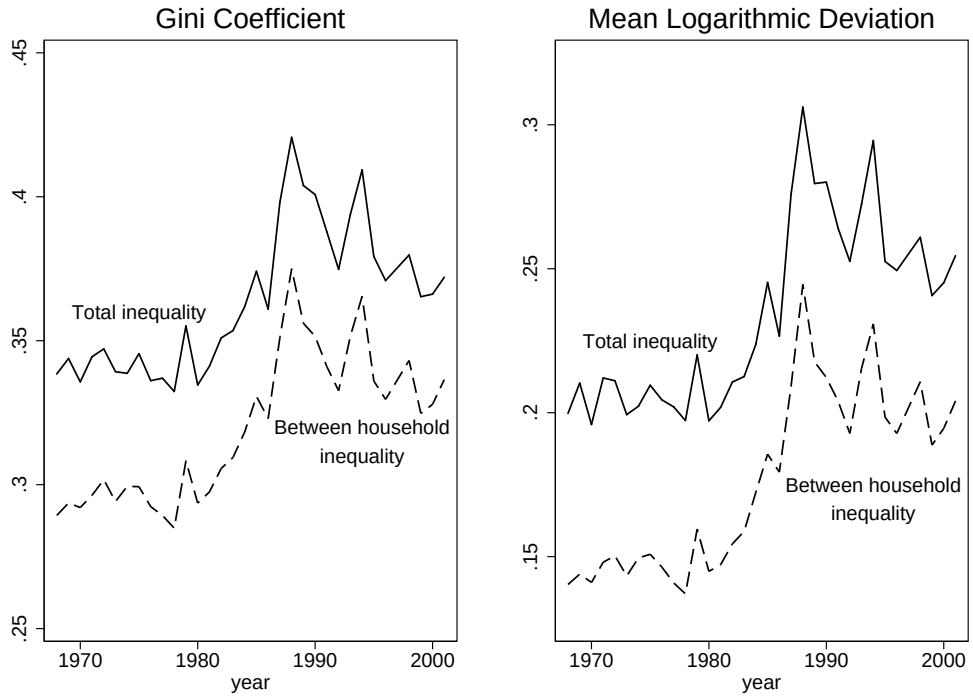

Figure 8: **Total and between household decomposition of inequality trends.**
Source: Own calculation from the FES.

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

Source: Own calculation from the FES.

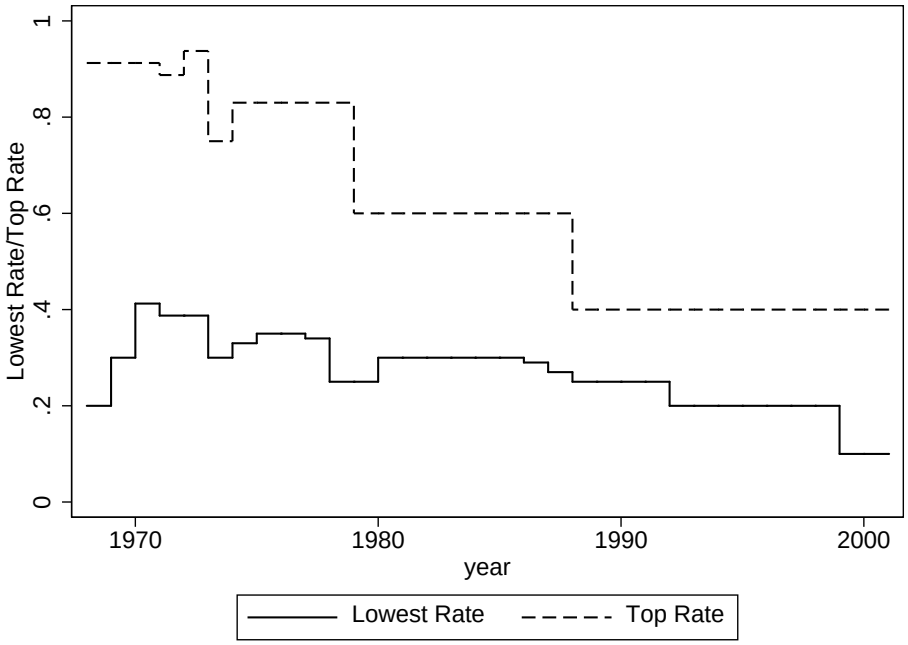

Figure 10: **Bottom and top marginal tax rates.**
Source: UK National Statistics (1968–2001).