AgEcon SEARCH
RESEARCH IN AGRICULTURAL & APPLIED ECONOMICS

The World's Largest Open Access Agricultural & Applied Economics Digital Library

This document is discoverable and free to researchers across the globe due to the work of AgEcon Search.

Help ensure our sustainability.

Give to AgEcon Search

AgEcon Search
<http://ageconsearch.umn.edu>
aesearch@umn.edu

*Papers downloaded from **AgEcon Search** may be used for non-commercial purposes and personal study only. No other use, including posting to another Internet site, is permitted without permission from the copyright owner (not AgEcon Search), or as allowed under the provisions of Fair Use, U.S. Copyright Act, Title 17 U.S.C.*

Queen's Economics Department Working Paper No. 1019

Consumption Inequality and Intra-Household Allocations

Jeremy Lise

Department of Economics, Queen's University

Shannon Seitz

Department of Economics, Queen's University

Department of Economics
Queen's University
94 University Avenue
Kingston, Ontario, Canada
K7L 3N6

6-2004

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

June 9, 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 the level of 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 in the UK since the 1970s.

JEL Classification: D12, D13, D63, J12, J22

Keywords: Collective Model, Consumption Inequality, Marital Sorting

*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, Bev Lapham, Arthur Lewbel, Frank Lewis, Hao Li, Huw Lloyd-Ellis, Mary MacKinnon, Krishna Pendakur, and participants at the Zeuthen Workshop at the University of Copenhagen, HEC Montreal, Queen's University, Simon Fraser University, Australia National University, the 2004 SOLE Annual Meetings, the 2004 CEA Annual Meetings, Workshop on Consumer Behavior and Welfare, 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; Blundell and Preston, 1998; Krueger and Perri, 2003). This shift has occurred for two main 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, 1997; 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 comprehensive measures of individual level consumption for households with more than one member are not available in the data. 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. 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.¹

This criticism levelled against the use of adult equivalence scales may seem to be a minor point until one considers the following. First, empirical tests of the “uni-

¹A substantial departure from this literature is recent work by Browning, Chiappori, and Lewbel (2004), where the assumption that household members split consumption equally is relaxed in the construction of adult equivalence scales.

tary” 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 have changed over time, then adult equivalence scales may provide an inaccurate picture of the trends in consumption inequality.

Our paper takes a first step toward addressing this issue and in doing so makes three contributions to the literature on consumption inequality. First, we construct and estimate a static 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.³ Previous estimates of collective models indicate that the share of consumption received by members of the household is strongly related to their earnings, or more accurately, their earnings potential (Browning et al., 1994; Browning and Chiappori, 1998; Chiappori et al., 2002; Blundell et al., 2002; Donni, 2001, 2003).⁴ Under relatively weak identification assumptions, the model allows us to infer the level of consumption allocated to each member of the household, which is necessary for the measurement of individual level consumption inequality.

Second, we use estimates of the our model to demonstrate how consumption inequality *within* households relates to consumption inequality *between* households.

²In 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 consumption allocation is independent of the source of income.

³See Haddad and Kanbur (1990), Kanbur and Haddad (1994), and Beaupri (2001) for studies of inequality within the household that do not adopt the collective framework.

⁴Further, Lundberg, Pollak and Wales (1997) provide evidence from a natural experiment in the UK that suggests the source of income in the household affects the consumption allocation.

The model is estimated on a sample of one and two person households from the UK Family Expenditure Survey (FES) for the years 1968 to 2001. We have two main findings. First, measures 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 assumption implicit in adult equivalence scales is valid *only* for households in which the wife has the same earnings as her husband. This result highlights the restrictiveness of the assumptions underlying standard adult equivalence scales. Our analysis demonstrates that relaxing these assumptions yields very different implications regarding our measures of consumption inequality.

The third contribution of our paper is to provide evidence on the importance of several potential explanations for the rise in consumption inequality between households and the fall in inequality within households since the 1970s in the UK. 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.⁵ 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 time can explain a large fraction of the rise in consumption inequality between 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 are presented in Section 5. Section 6 presents a decomposition of consumption inequality and considers the importance of several explanations for the trends in consumption inequality. Section 7 concludes.

⁵In this instance, the degree of marital sorting is measured by the correlation between characteristics, such as education and wages, across spouses.

2 Trends in Consumption and Earnings Inequality in the UK

In this section, we outline the main stylized facts regarding consumption and income inequality in the UK between 1968 and 2001.⁶ The data we use to conduct our analysis comes from the UK Family Expenditure Survey (FES). The FES contains information on household consumption expenditures and earnings over the period 1968 to the present. In the construction of the following stylized facts, we restrict the sample to individuals between the ages of 16 and 65 and eliminate students, retirees and the self-employed. We are particularly interested in the following four features of the data:

1. There has been a large rise in earnings inequality between individuals. The first panel in Figure 1 documents the trend in the Gini index for the distribution of individual and household earnings. In particular, the Gini index has risen by 12% over the past 30 years. This rise in earnings inequality in the UK has been well documented in the literature (e.g. Blundell and Preston, 1998).
2. Although earnings inequality between individuals is much higher than earnings inequality between households, the latter rose much more rapidly: the Gini index for inequality between households rose by 41% between 1968 and 2001.
3. 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.⁷

⁶Many of these findings have been documented in the literature for Canada (Pendakur, 1998), the United States (Johnson and Shipp, 1997 and Krueger and Perri, 2003) and Australia (Barrett, Crossley and Worswick, 2000).

⁷Krueger and Perri (2003) find a large rise in income inequality in the US since the 1970s while

4. As illustrated in Figure 2, the correlation between the earnings of husbands and wives increased dramatically over time. This is due to both the fall in the gender wage gap and the rise in female labor supply. Figure 3 highlights the dramatic change in the gender wage gap and in women’s contribution to household labor income between 1968 and the present. The dashed line represents the female’s share of potential income, defined as the share of labor earnings that would be contributed by the wife if both spouses worked full-time. The solid line represents women’s share of actual household earnings.⁸ Overall, potential earnings of wives increased by 13.5%, and women’s share of earnings in the household increased by 93% over the sample period. The latter partly reflects the increase in women’s wages relative to those of men, but also the large changes in female and male employment rates and hours worked since the 1960s.

In summary, the evidence presented here highlights the fact that there has been a large rise in earnings and consumption inequality *between* households while at the same time there has been a fall in inequality in the earnings distribution *within* households. In the next section, we present a model of intra-household allocations that allows us to explore the implications of this evidence for the measurement of consumption inequality.

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-household allocations suggests that the source of income plays an important role in determining how resources are allocated within households. 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

consumption inequality remained roughly constant.

⁸For 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.

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. 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, expenditures on private consumption and expenditures on public consumption for an agent of gender g , $g \in \{m, f\}$ by L^g , C^g , and P , 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 measurement of consumption inequality and is thus included here.⁹ Denote total household non-labor income net of savings by Y . 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.

⁹Chiappori, Blundell and Meghir (2002) establish conditions under which the collective model with public goods is identified.

Preferences for single agents are described by $U^g(u^g(L^g, C^g), P)$, 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:

$$\begin{aligned} \max_{L^g, C^g, CP} \quad & U^g(u^g(L^g, C^g), P) \\ \text{subject to} \quad & C^g + P = w^g(l) + Y. \end{aligned}$$

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:

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

As with singles, we assume that private consumption and leisure (C^g, L^g) are separable from consumption of the public good (P) for married couples. Preferences for a married person of gender g can be described by:

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

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:

$$\begin{aligned} \max_{L^f, L^m, C^f, C^m, P} \quad & \lambda V^f(v^f(L^f, C^f), P) + (1 - \lambda) V^m(v^m(L^m, C^m), P) \\ \text{subject to} \quad & C^f + C^m + P = w^f(l) + w^m(l) + Y. \end{aligned} \tag{1}$$

The Pareto weight, λ , represents the female's bargaining power within the household, and will typically be a function of full-time labor income $(w^f(1), w^m(1))$, non-labor income (Y) and other “distribution factors” $(\mathbf{z})$ that influence household bargaining power, but do not have an effect on individual preferences, as in Chiappori, Fortin, and Lacroix (2002).

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 (P) and on how to divide the remaining non labor income $y = Y - P$. The assumption that consumption of the public good is separable from leisure and private consumption is key to 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:

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

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

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 that 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

The question we aim to address in this paper is whether measures of consumption inequality from the collective model differ from measures in the literature based on standard 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 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 and public consumption. This assumption still allows for complementarities in leisure time across spouses and for complementarities between leisure and consumption for married couples.

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.¹⁰ In particular, we assume that individuals can choose from H discrete labor supply possibilities, in addition to non-participation.¹¹ 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. 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 P are observed although the distribution of private consumption between the husband and wife (C^f and C^m) is not observed.

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. The utility a single individual of gender g derives from labor supply choice h is:

$$\begin{aligned} U_h^g &= v^g(l_h, C_h^g) + \omega^g(P) + \varepsilon_h^g \\ &= \alpha_l^g l_h + \alpha_{ll}^g l_h^2 + \alpha_{cl}^g l_h C_h^g + \alpha_c^g C_h^g + \alpha_{cc}^g (C_h^g)^2 + \alpha_P^g P_h^g + \alpha_{PP}^g (P_h^g)^2 + \varepsilon_h^g, \end{aligned}$$

and the utility a married individual of gender g derives from labor supply choice h is:

$$V_h^g = \beta_l^g l_h + \beta_{ll}^g l_h^2 + \alpha_{cl}^g l_h C_h^g + \alpha_c^g C_h^g + \alpha_{cc}^g (C_h^g)^2 + \beta_P^g P_h^g + \beta_{PP}^g (P_h^g)^2 + \varepsilon_h^g,$$

where ε_h^g is an unobserved preference component that is assumed to be distributed *iid*

¹⁰Blundell 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.

¹¹This assumption is not necessary for identification, and is not very restrictive, as the discrete choice of hours can be any integer value of weekly hours.

across individuals and labor supply alternatives. This specification allows preferences for leisure and public consumption to differ between married and single men and women, but restricts preferences for private consumption to be the same for both married couples and singles.

Assume the sharing rule is linear in the distribution factors:

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

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. 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,

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

for married women and

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

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 $h > 0, \forall h \in \{1, 2, \dots, H\}$ and not

working ($h = 0$) for single men and women can be expressed as:

$$U_h^g - U_0^g = \alpha_l^g l_h + \alpha_{ll}^g l_h^2 + \alpha_{cl}^g l_h \cdot \tilde{w}^g(l_h) + \alpha_c^g \tilde{w}^g(l_h) + \alpha_{cc}^g [\tilde{w}^g(l_h)]^2 + 2\tilde{w}^g(l_h) \cdot y + \varepsilon_h^g - \varepsilon_0^g, \quad (6)$$

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 $h > 0, \forall h \in \{1, 2, \dots, H\}$ and not working ($h = 0$) is described by:

$$\begin{aligned} V_h^f - V_0^f &= \beta_l^f l_h + \beta_{ll}^f l_h^2 + \alpha_{cl}^f l_h \cdot \tilde{w}^f(l_h) + \alpha_{cl}^f (\mathbf{z}^\top \phi) \cdot l_h \cdot y \\ &\quad + \alpha_c^f \tilde{w}^f(l_h) + \alpha_{cc}^f [\tilde{w}^f(l_h)]^2 + 2\alpha_{cc}^f (\mathbf{z}^\top \phi) \cdot \tilde{w}^f(l_h) \cdot y + \varepsilon_h^f - \varepsilon_0^f. \end{aligned} \quad (7)$$

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

$$\begin{aligned} V_h^m - V_0^m &= \beta_l^m l_h + \beta_{ll}^m l_h^2 + \alpha_{cl}^m l_h \cdot \tilde{w}^m(l_h) + \alpha_{cl}^m l_h \cdot y - \alpha_{cl}^m (\mathbf{z}^\top \phi) \cdot l_h \cdot y + \alpha_c^m \tilde{w}^m(l_h) \\ &\quad + \alpha_{cc}^m [\tilde{w}^m(l_h)]^2 + 2\alpha_{cc}^m \tilde{w}^m(l_h) \cdot y - 2\alpha_{cc}^m (\mathbf{z}^\top \phi) \cdot \tilde{w}^m(l_h) \cdot y + \varepsilon_h^m - \varepsilon_0^m. \end{aligned} \quad (8)$$

The parameters β_l^g, β_{ll}^g are directly identified from the data on married individuals. Given $\alpha_{cl}^g, \alpha_c^g$, and α_{cc}^g , identified from data on single individuals, it is straightforward to recover the sharing rule parameters (ϕ) .¹²

The differences in utility described by Equations (6), (7) and (8) can be expressed for single individuals in reduced form as:

$$\begin{aligned} U_h^g - U_0^g &= \Pi_l^g l_h + \Pi_{ll}^g l_h^2 + \Pi_{ly}^g l_h y + \Pi_{wl}^g l_h \tilde{w}^g(l_h) + \Pi_{wl}^g \tilde{w}^g(l_h) + \Pi_{(wl)^2}^g [\tilde{w}^g(l_h)]^2 \\ &\quad + \Pi_{wly}^g \tilde{w}^g(l_h) \cdot y + \varepsilon_h^g - \varepsilon_0^g, \end{aligned}$$

and for married individuals as:

$$\begin{aligned} V_h^g - V_0^g &= \Pi_l^g l_h + \Pi_{ll}^g l_h^2 + \Pi_{ly}^g l_h y + \Pi_{wl}^g l_h \tilde{w}^g(l_h) + \Pi_{wl}^g \tilde{w}^g(l_h) + \Pi_{(wl)^2}^g [\tilde{w}^g(l_h)]^2 \\ &\quad + \Pi_{wly}^g \tilde{w}^g(l_h) \cdot y + \Pi_{lym}^g l_h y + \Pi_{zlym}^g \cdot \mathbf{z} l_h y + \Pi_{wlym}^g \tilde{w}^g(l_h) \cdot y \\ &\quad + \Pi_{zwlym}^g \cdot \mathbf{z} \tilde{w}^g(l_h) \cdot y + \varepsilon_h^g - \varepsilon_0^g, \end{aligned}$$

¹²The parameters capturing preferences over the public consumption good ($\alpha_P^g, \alpha_{PP}^g, \beta_P^g, \beta_{PP}^g$) can not be identified as the utility from consumption of the public good is the same regardless of the labor supply decision. 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.

where the Π s are reduced form parameters. The system above implies a set of over-identifying restrictions for the sharing rule parameters that enable us to test 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:

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

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 private and public consumption expenditures for households, on wages and labor supply for individuals, 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.¹³ 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 information over the period 1968 to the present, which allows the study of changes in consumption inequality over a long period of time.¹⁴

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-

¹³Battistin (2003) documents reporting errors in the US Consumer Expenditure Survey due to survey design.

¹⁴We presently exclude the year 1997 from our analysis due to a missing data problem.

household allocation of resources for children’s consumption. This is obviously an important issue. To this end, our estimates of the sharing rule and the comparison of various inequality measures only apply to households without children. We leave to future work an analysis of consumption inequality for the entire sample of households. 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.¹⁵ Descriptive statistics for our entire sample 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.¹⁶

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.¹⁷ After tax earnings are subsequently computed by converting weekly wage income to an annual base, deducting the

¹⁵The 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.

¹⁶In 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.

¹⁷The 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.

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, so 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 Figure 4, which suggests 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 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 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, 2004), 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:

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

In estimation, heterogeneity in preferences for leisure is introduced through the vector X which includes age, birth cohort, education, marital status, region, and quarter and year to control for seasonality and cyclical effects.¹⁸ The parameters Π_l^g and Π_{ll}^g are assumed to be linear functions of the observed characteristics, so that for individual i we have

$$\begin{aligned} \Pi_l^g &= X_i \Pi_l^g \\ \Pi_{ll}^g &= X_i \Pi_{ll}^g. \end{aligned}$$

Estimation proceeds in two steps. First, we estimate a 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 Estimation 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

¹⁸In 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.

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 reduced form. The sharing rule parameters recovered from the 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, 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_{wly}^f , 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. 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.¹⁹

¹⁹For comparison purposes, we also estimated a version of the model where predicted wages were used in place of actual wages for all individuals and found no substantial changes in the parameter estimates. Results are available from the authors upon request.

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 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 consistent with 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 seem to provide some 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 3.²⁰ The estimates suggest that a 10% increase in the share of potential earnings 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.

5.1 Adult Equivalence Scales Revisited

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

²⁰See the Appendix for further details on the minimum distance estimation exercise.

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. We set the age difference between spouses to the average age difference in the data. We then 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!²¹ 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.

5.2 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.²² 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

²¹To be precise, husbands and wives will split consumption equally when they have approximately the same wages and hours.

²²Full estimation results are available from the authors upon request.

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 and have the same wage. 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 is 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 Columns 1 and 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 have 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 parameter estimates are quite similar across specifications is surprising considering the possibility of 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. We do this for two reasons. First, we want to allow for the possibility that the preference shocks are correlated across labor supply choices. Second, we want to allow for additional flexibility in estimating preferences over leisure.²³

²³We specify $\Pi_l = X_l \Pi_l + u_{hi}$ and $\Pi_{hl} = X_{hl} \Pi_{hl} + u_{hhi}$, 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 in an increased transfer of between 11% and 21% for the 1910 to 1950 cohorts, which is close to the range reported in Column 1 above.

6 Consumption Inequality

6.1 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 married individuals, is then equal to individual private consumption, plus household public consumption:

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

and

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

Single individuals 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

²⁴See Column 2 of Table 3.

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.

The distribution of consumption within the household under the sharing rule case for the years 1968 and 2001 is presented in Figure 5. The sharing rule estimates suggest that the mean of the intra-household consumption inequality distribution increased over time and the variance decreased. The dashed line in the right panel of the figure represents the wife's share of consumption under equal sharing. A comparison of equal division and the sharing rule distributions further suggests that equal division may not be an accurate representation of consumption allocations in the majority of households.

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 a well known and widely used inequality index, it does not allow overall inequality to be exactly decomposed into within and between group contributions. As this is one of the main objectives of this paper we also compute the Mean Logarithmic Deviation (MLD)²⁶ in consumption, defined as

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

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

$$I_{\alpha}^T(C) = I_{\alpha}^W(C) + I_{\alpha}^B(C), \quad (11)$$

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

²⁶The MLD is a member of the Generalized Entropy Class, the only class of additively decomposable inequality indices (Shorrocks, 1984).

where $I_\alpha^W(C)$ is within household inequality and $I_\alpha^B(C)$ is between household inequality. Under the assumption of equal division, within household inequality is zero; therefore, 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 using Equation (11).

The time-series trend of total and between household inequality for the years 1968 to 2001 is presented in Figure 6. 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 using the MLD and Gini index, respectively.

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 7. 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. The stylized facts presented in Section 2 allude to possible reasons for 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.

6.2 Accounting 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. We focus on the years 1978 to 2001, 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.

6.3 Wages and Labor Supply

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. It is therefore of interest to assess the extent to which changes in the distribution of wages can account for the rise in consumption inequality. To answer this question, we consider two thought experiments. In the first exercise, we re-weight the data for 1978 so that the wage distribution in 1978 matches that of 2001.²⁷ This experiment captures the fall in the gender wage gap and aggregate changes in the wage distribution, but not changes in sorting on wages within households. To capture the latter, 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 top four rows of both panels of Table 6.

What is 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 39% of the rise in between household inequality and 78% of the fall in within household inequality. With respect to sorting on hours, 32% of the rise in between household inequality and 98% 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 the distribution of consumption across

²⁷In particular, we construct histograms of log wages for 1978 and 2001. The histograms used to re-weight the wage distributions have 10 bins each. We re-weight 1978 data so that the histograms of log wages are the same in both years.

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.²⁸

6.4 Demographics and Household Composition

Next, we consider the hypothesis that the rise in consumption inequality is capturing cohort effects due to the changes in the age structure of the population. To assess the importance of the changing age structure in the population for our measures of consumption inequality, 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 fifth row of the bottom panel of Table 6. 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.²⁹ 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 match the proportions in 2001. The results of this exercise suggest that changes 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

²⁸The 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 8, 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. It appears that changes in marginal tax rates had the effect of increasing between household inequality substantially while having only a modest effect on within household inequality.

²⁹Although single person households have no within household inequality by definition, it is still the case that there may exist substantial inequality between single adult households.

explain a little more than 30% of the change in the Gini index over time, most of this effect coming through household composition.

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 *between* households. Our estimates suggest that measures of consumption inequality that ignore the potential for within household inequality may underestimate the level of individual consumption inequality by 30% and may over-state 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. The results also indicate that changes in sorting on wages and hours play prominent roles in explaining the inequality trends over time.

We also highlight the restrictive assumption implicit in adult equivalence scales, namely that consumption is equally split across spouses in two-person households. Our analysis suggests this assumption is only valid for households in which both partners have the same earnings. We illustrate that relaxing this assumption changes our measures of consumption inequality substantially. Browning, Chiappori, and Lewbel (2004) demonstrate that bargaining power and adult equivalence scales can both be identified within a collective framework. They are the first to relax this equal sharing assumption in the construction of adult equivalence scales. Our work provides strong evidence in support of this line of research.

Although our paper represents a first step towards understanding the role of within household inequality in the measurement of inequality, much work remains to be done. For one, the estimates presented above apply only to single person households and

to married couples without children. On a related note, we do not incorporate home production in the model. If non-working spouses spend time producing goods at home and caring for children, it may be the case that we are over-estimating the effect of women's earnings on the sharing rule. To incorporate home production and children, we require additional data on the way non-working time is spent in the household and/or we need to impose additional restrictions on the model. Both issues are worth further exploration.

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 (2004): "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.
- CUTLER, 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 Microeconomic 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.
- 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.
- MCELMROY, 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. (2004): “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.
- 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

$$\theta = \left(\phi_0, \phi_1, \phi_2, \alpha_{cl}^f, \alpha_{cc}^f, \alpha_{cl}^m, \alpha_{cc}^m, \alpha_{lX}^f, \alpha_{lX}^m, \beta_{lXm}^f, \beta_{lXm}^m \right)^\top$$

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

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

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:

$$\left(\hat{\Pi} = f(\theta) \right) \Rightarrow \begin{pmatrix} \hat{\Pi}_{ly}^f & = & \alpha_{cl}^f \\ \hat{\Pi}_{lym}^f & = & \alpha_{cl}^f (\phi_0 - 1) \\ \hat{\Pi}_{z_1lym}^f & = & \alpha_{cl}^f \phi_1 \\ \hat{\Pi}_{z_2lym}^f & = & \alpha_{cl}^f \phi_2 \\ \hat{\Pi}_{wly}^f & = & \alpha_{cc}^f \\ \hat{\Pi}_{wlym}^f & = & \alpha_{cc}^f (\phi_0 - 1) \\ \hat{\Pi}_{z_1wlym}^f & = & \alpha_{cc}^f \phi_1 \\ \hat{\Pi}_{z_2wlym}^f & = & \alpha_{cc}^f \phi_2 \\ \hat{\Pi}_{ly}^m & = & \alpha_{cl}^m \\ \hat{\Pi}_{lym}^m & = & -\alpha_{cl}^m \phi_0 \\ \hat{\Pi}_{z_1lym}^m & = & -\alpha_{cl}^m \phi_1 \\ \hat{\Pi}_{z_2lym}^m & = & -\alpha_{cl}^m \phi_2 \\ \hat{\Pi}_{wly}^m & = & \alpha_{cc}^m \\ \hat{\Pi}_{wlym}^m & = & -\alpha_{cc}^m \phi_0 \\ \hat{\Pi}_{z_1wlym}^m & = & -\alpha_{cc}^m \phi_1 \\ \hat{\Pi}_{z_2wlym}^m & = & -\alpha_{cc}^m \phi_2 \\ \hat{\Pi}_{lX}^f & = & \alpha_{lX}^f \\ \hat{\Pi}_{lX}^m & = & \alpha_{lX}^m \\ \hat{\Pi}_{lXm}^f & = & \beta_{lX}^f \\ \hat{\Pi}_{lXm}^m & = & \beta_{lX}^m \end{pmatrix},$$

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

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

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

Table 1: Descriptive Statistics from the FES

	Male				Female			
	Single		Married		Single		Married	
1968 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		
$\phi_0^{fcl} = \Pi_{lym}^f / \Pi_{ly}^f + 1$	-0.534	***		-0.517	***	
	(0.056)			(0.055)		
$\phi_0^{fcc} = \Pi_{wlym}^f / \Pi_{wly}^f + 1$	-1.707	***		-1.634	***	
	(0.536)			(0.520)		
$\phi_0^{mcl} = -\Pi_{lym}^m / \Pi_{ly}^m$	-0.393	***		-0.388	***	
	(0.089)			(0.088)		
$\phi_0^{mcc} = -\Pi_{wlym}^m / \Pi_{wly}^m$	-0.327	**		-0.329	**	
	(0.160)			(0.158)		
$\phi_1^{fcl} = \Pi_{z_1lym}^f / \Pi_{ly}^f$	2.297	***		2.300	***	
	(0.179)			(0.178)		
$\phi_1^{fcc} = \Pi_{z_1wlym}^f / \Pi_{wly}^f$	5.796	***		5.794	***	
	(1.687)			(1.691)		
$\phi_1^{mcl} = -\Pi_{z_1lym}^m / \Pi_{ly}^m$	1.714	***		1.765	***	
	(0.175)			(0.175)		
$\phi_1^{mcc} = -\Pi_{z_1wlym}^m / \Pi_{wly}^m$	1.635	***		1.726	***	
	(0.342)			(0.341)		
$\phi_2^{fcl} = \Pi_{z_2lym}^f / \Pi_{ly}^f$				-0.009	***	
				(0.003)		
$\phi_2^{fcc} = \Pi_{z_2wlym}^f / \Pi_{wly}^f$				-0.038	**	
				(0.019)		
$\phi_2^{mcl} = -\Pi_{z_2lym}^m / \Pi_{ly}^m$				-0.011	***	
				0.004		
$\phi_2^{mcc} = -\Pi_{z_2wlym}^m / \Pi_{wly}^m$				-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 (0.021)	-0.310 (0.021)
ϕ_1	1.584 (0.058)	1.592 (0.057)
ϕ_2		-0.004 (0.001)
$\phi_0 + \frac{1}{2}\phi_1$	0.475 (0.012)	0.485 (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 (0.038)	-0.186 (0.028)
ϕ_1	0.773 (0.091)	1.220 (0.073)
$\phi_0 + \frac{1}{2}\phi_1$	0.395 (0.018)	0.424 (0.015)
% of Total Consumption	0%	17%
	Public Goods (ii) (i + heat)	Public Goods (iii) (ii + durables)
ϕ_0	-0.239 (0.027)	-0.317 (0.015)
ϕ_1	1.352 (0.071)	1.584 (0.013)
$\phi_0 + \frac{1}{2}\phi_1$	0.437 (0.015)	0.475 (0.010)
% of Total Consumption	24%	33%
	Public Goods (iv) (iii + transport)	Public Goods (v) (iv + services)
ϕ_0	-0.310 (0.013)	-0.256 (0.008)
ϕ_1	1.665 (0.038)	1.562 (0.025)
$\phi_0 + \frac{1}{2}\phi_1$	0.523 (0.010)	0.525 (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.

Absolute Change	Gini Index		Mean Logarithmic Deviation		
	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
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
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

	Gini Index		Mean Logarithmic Deviation		
	Total	Between	Total	Between	Within
1978 to 2001 Change	0.040	0.052	0.057	0.067	-0.010
Percentage change attributable to:					
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
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

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 actual household earnings provided by wife.**
Source: Own calculation from the FES.

Figure 4: **Histogram of usual weekly hours.**
Source: Own calculation from the FES.

Figure 5: Distribution of consumption inequality within households.

Figure 6: Total and between household decomposition of consumption inequality trends.

Source: Own calculation from the FES.

Figure 7: **Relative changes in between and within household inequality trends.**

Source: Own calculation from the FES.

Figure 8: **Bottom and top marginal tax rates.**

Source: UK National Statistics (1968–2001).