Gender Inequality in Expenditure: Great Britain 1978 to 2019[*]

Great Britain 1978 to 2019

Richard Blundell[1,2], Heidi Karjalainen[2], Valérie Lechene[1,2], Krishna Pendakur[3]

July 1, 2025

1: UCL; 2: IFS; 3: SFU

Abstract

We develop the concept of *relative resource shares*, defined as the fraction of total adult expenditure that women command within a household. To recover relative resource shares from expenditure data, we introduce a new identifying restriction, *weak similarity of preferences across people (WSAP),* a shape restriction on preferences that is weaker than those used in the existing literature. With repeated cross section data on household expenditure, we recover the relative contributions of changes in characteristics and changes in the resource share function to the evolution of relative resource shares over time. We apply this new methodology to estimate within-household gender inequality in Great Britain from 1978 to 2019. Women's relative resource shares are estimated to have increased by 12–13 percentage points over this period, rising from disparity to roughly parity with men. As a consequence, individual-level inequality rose by less than household-level inequality. We find that changes in characteristics related to women's bargaining power play a key role in explaining these changes. We document strong differences between childless women and mothers.

1 Introduction

The aim of much of microeconomic policy is to affect outcomes at the individual level, by modifying the environment of individual decision making. Policy design and evaluation should therefore be based on measures of outcomes at the individual level. Yet, for practical rather than ideological reasons, outcomes are usually measured at the household level. Indeed, even though individual incomes can be measured, because the intra-household distribution of incomes and expenditure cannot be observed, the welfare effects of policies targeting individuals are usually measured under the assumption of equal sharing within households. But if resources are not shared equally inside households, and if policy can affect the within-household distribution, the true effect of policy on individual outcomes and individual welfare is unclear. Suppose for instance that child benefits are spent by parents on their intended beneficiaries (the children); standard methods which assume equal sharing of resources will underestimate the effect of the benefit on expenditure and welfare of children, since the benefit will be assumed to have been shared equally between household members.

In addition to the possibility that intended policy outcomes are changed via intra-household redistribution, the market can also have a similar effect. Suppose a policy incentivises certain types of individuals, for example women, to acquire more education, with the intention to improve their outcomes, since returns to education are generally found to be positive. The beneficial welfare effects of such a policy could be attenuated by general equilibrium effects if the returns to education within marriage (or other types of cohabitation relationships) decrease as more women become educated.

To develop our ideas we envision individuals as bundles of characteristics: education, experience, age, temperament, and so on. As such, the attributes of individuals are valued on (at least) two markets: the labour market and the marriage market, where their characteristics determine respectively the wage and the within-household resource share. The wage is the price paid for the individual's bundle of characteristics at the equilibrium of the labour market. The *resource share*, defined as the fraction of household resources that goes to each individual in the marriage, is a very similar object to the wage. It is the price paid for a bundle of characteristics at the equilibrium of the marriage market.

At a point in time, resource shares are determined by the equilibrium between supply and de-

2

mand in the marriage market. As equilibrium objects, resource shares evolve over time. Returns to characteristics in the resource share (RS) function are their *prices.* As the distribution of characteristics in society changes through time, the value of a given bundle of characteristics changes. When measuring changes in resource shares over time, we need to disentangle the effect of *differences in characteristics* (which affect the *value* of the RS) from the effect of *changes in their returns* (which affect the RS *function*).

We adapt a recently developed tool from Lechene, Pendakur and Wolf (2022) to measure the share of household resources which accrue to each type of adult in the household. In doing so we remove the assumption of equal sharing within households. We show our approach is applicable with standard household budget data. Like many papers in this literature, we base our empirical strategy on the demands for assignable goods — those whose consumption is observed at the person level. In our case, we observe clothing expenditure for both adult men and adult women and use those data to recover resource shares. We present both theoretical and empirical contributions to the existing literature on measuring resource shares.

This paper extends the theoretical framework around measurement of individual resource shares from previous literature in several directions.

First, we introduce *relative resource shares*, which equal the fraction of adult expenditure consumed by women. This is a tool for estimating gender inequality inside households – because gender inequality focusses on adult men and women, we do not need to identify children's consumption inside households.

Second, because we have fewer parameters to identify than approaches that seek to recover resource shares for all types of individuals inside households, including children, we are able to use weaker identifying assumptions than previous papers. We provide a semi-parametric identifying restriction on preferences, called Weak Similarity Across People (WSAP), under which relative resource shares may be identified from cross sectional demand behaviour. Similar to Lechene, Pendakur and Wolf (2022), for the case with linear Engel curves, we provide a linear reduced form from which structural parameters can be identified.

Third, we show how to recover relative resource shares with time varying resource share functions from repeated cross sectional data. Previous contributions to the literature on estimation of resource shares over time have not allowed for the resource share function to vary. In a seminal paper, Lise and Seitz (2011) use consumption data from repeated cross sections to identify the level of resource shares over time in the UK. We build on the approach of Lise and Seitz (2011) and later work by Bargain, Donni and Hentati (2023), both of whom identify resource shares and scale economies over time under the assumption that the model is time invariant. Our approach allows for time varying resource share functions, at the cost of not identifying scale economies. We do this by extending the Lechene, Pendakur and Wolf (2022) model to allow for time and un-

3

observed prices, and developing an approximation to the resource share function with time varying coefficients. Under these flexible assumptions, it is possible to identify the resource share function in each time period, which in turn allows us to assess the relative contributions of changes in characteristics and changes in the resource share function in driving the changes in the estimated relative resource shares.

In addition to these theoretical contributions, this paper makes three empirical contributions. First, we estimate resource shares to assess how the within-household distribution, especially gender inequality, changed over time. Second, we use Oaxaca decomposition to document how changes in the characteristics and changes in the resource share function contribute to the changes in levels of relative resource shares through time. Third, we use these estimates to study the evolution of individual-level consumption inequality in the UK over the period 1978 to 2019.

Resource shares are estimated separately for households with and without children. These are of course not permanently separated populations, but rather the same households can be childless in some periods and have children in others. However, in line with most analysis of inequality, we treat these two types of households separately. It is also worth highlighting that our data is cross-sectional and we estimate average resource shares at different points in time, rather than for example over the life cycle.

We find that women's relative resource shares increased by about 12 percentage points for women in couples with children and by 13 percentage points for women in childless couples over the period of 1978 to 2019. The evolution of relative resource shares through time is different for childless women and for mothers. For women with children, a long period of stagnation of the relative resource shares, from 1978 until about 2008, is followed by an increase in the point estimates. For childless women, a period of significant increase from 1978 until about 1995 is followed by a period of stagnation until about 2008, at which point, resource shares increase, although not statistically significantly.

The overall evolution of the relative resource shares that we describe is similar to the findings of Lise and Seitz (2011) and Bargain, Donni and Hentati (2023), who find increases in women's resource shares over the long term in the UK. We show that this evolution is partly due to improving characteristics of women (higher levels of education, relative wages, and child benefit receipt); and partly contextual (as captured by the changing resource share function which reflects market equilibrium, time, social change). This is different from Lise and Seitz (2011) and Bargain, Donni and Hentati (2023), who do not allow for time effects, meaning that they attribute the entirety of the change in resource shares to changes in characteristics, without allowing for the possibility that the mapping between characteristics and resource shares is itself changing over time.

For childless women, decomposition of the changes in relative resource shares shows that both improving characteristics and the relative resource share function itself were pushing the relative

resource shares upwards until 1995. After that, even though women's characteristics continued to improve, the relative resource share function itself declined, resulting in childless women's relative resource shares in 2019 to be equal to their level of 1995. One potential interpretation for these patterns is through a general equilibrium lens: the value of being a childless single woman rose over the 1970s through 1990s, due to declining gender inequality in the labour market, declining stigma of singlehood and childlessness, and increasing reproductive freedom. This pushed up the relative resource share function. However, by the mid 90s, these processes slowed, so that the relative resource share function was no longer increasing.

For mothers, the point estimate of the relative resource shares increased by 12 percentage points over the period. The period can be decomposed into two sub-periods, from 1978 to 2007, with constant average relative resource shares and 2007 to 2019, with a (statistically insignificant) increase in relative resource shares.

Although the increase in relative resource shares after 2007 is insignificant for both childless women and mothers, the fact that it is increasing for both is suggestive of an increase of women's power in marital relationships generally. In the UK, the financial crisis of 2008, linked to the 2007-2009 Great Recession in the US, saw the economy shrink, and unemployment and uncertainty increase. These features of the economic situation impact on the marriage market, by increasing the potential gains to marriage, as partners can share unemployment and income risk, and generally mitigate the consequences of uncertainty. This could explain not only an increase in relative resource shares at the average but also an increase in the variability of relative resource shares, which we observe in the data.

Finally, we use the estimated relative resource shares, and measures of household consumption, to assess the effect of the intra-household distribution on individual-level consumption inequality in the UK over the period 1978 to 2019. We show that the well-documented increase in household-level market income inequality and household-level net income inequality in the 1980s and 1990s were partly offset by declining gender inequality shown by our estimated relative resource shares. In particular, individual-level inequality only rose by about two-thirds as much as did household-level inequality.

The remainder of the paper is as follows: In section (2), after a quick tour of household models, we introduce the main theoretical building blocs of our identification strategy. Technical material and theorems relating to the identification strategy are presented in a series of appendices. The following section (3) lays out the empirical implementation of the model. We then turn to the data in (4). Results concerning relative resource shares are in section (5) and those concerning inequality are in section (6). Section (7) concludes.

2 Model

We start by providing the necessary background on theories of household models (2.1) and show how, in these models, the assumption of efficiency of decision making within households naturally leads to the concept of the resource share in (2.2). In this section, we also describe the special role that *assignable* goods play in household models. We then introduce our new concept of the *relative* resource share and show how it relates to gender inequality in (2.3). We describe Engel curves in household models in (2.4), and show how we identify relative resource shares from Engel curve data on assignable goods in (2.5).

2.1 Household models

For most of the 20th century, economic models of households featured household utility functions or representative utilities. This was the case for instance in Becker (1965, 1981) in his pioneering models of what are now called unitary households, or households whose decisions can be represented by a maximization of a single utility function against a well-defined budget constraint. Household-level observed heterogeneity, such as family size and structure, was also incorporated into this unitary framework, e.g., in Jorgenson and Slesnick (1987), Blackorby and Donaldson (1993), Blundell, Duncan and Pendakur (1998) and Pendakur (1999). During the 1980s, Patricia Apps among others led the development of fully structural models of collective households that treated households as economic environments in which individuals—who have utility functions— live and interact. Examples of such models include McElroy and Horney (1980), Apps and Rees (1988) and Apps and Savage (1989). These models were very specific in the sense that they used particular models of individual utility functions, bargaining processes between individuals and household scale economies in consumption.

In the 1990s and 2000s, P.A. Chiappori (1988) and coauthors brought *efficient* collective household models to the foreground. The key insight here was that if the household is assumed to reach an efficient allocation and if scale economies are presumed to be embodied in public goods within the household, then the exact utility functions that people have and the exact process by which individuals bargain with each other do not have to be specified. Instead, a set of *generic* results was generated regarding *all* household models in which agents are presumed to collectively reach the Pareto frontier and in which all goods are purely public or purely private (or, as shown by Cherchye et al 2009, a mix of these). These results are presented in Chiappori (1988), Bourguignon and Chiappori (1994), Browning, Bourguignon, Chiappori, and Lechene (1994); Browning and Chiappori (1998), Chiappori, Blundell, Meghir (2002), Cherchye et al (2009) and Chiappori and Ekeland (2009).

More recently, a set of analogous results is available for the case where agents are presumed

to reach the Pareto frontier and household scale economies arise from sharing of goods, where the extent of sharing may be different for every good. This latter model was described by Browning, Chiappori and Lewbel (2002, 2013), and elaborated on by Lewbel and Pendakur (2008), Dunbar, Lewbel and Pendakur (2013) and Lechene, Pendakur and Wolf (2022). The demand functions for this set of models are particularly easy to work with, and form the basis of the work presented in this paper.

These models provide a framework that allows us to identify resource shares—which are not typically directly observable—from micro data on household expenditures. They therefore provide a new interpretation of patterns in data. Standard household budget survey data records most expenditure at the household level and some expenditure at the individual level (for instance men's clothing and women's clothing). Such data allows the empirical researcher to tell directly whether men have greater expenditure on clothing than do women, but we cannot say anything on overall resource shares of men and women based on data alone. With the help of a model and a theoretical framework, we can interpret spending patterns in such a way so as to reveal overall resource shares of men and women.

2.2 Efficiency, resource shares and assignable goods

A key assumption of much of the literature on identification in household models since the 1990s is that household members reach an efficient allocation, that is, the assumption that household members together reach the Pareto frontier. This is a powerful assumption. Chiappori (1988) showed that if we assume efficiency, then we can call upon general equilibrium theory from the 1950s about efficient economies and use that theory to describe households. In particular, one result from that foundational literature in Economics is that efficient economies can be *decentralised* to individual-level problems. In the context of a household, this implies that the following optimization programmes, $O1$ and $O2$ are behaviourally equivalent.

A first optimisation programme, $O1$, is a *centralised* household-level optimisation, in which $H()$ is a general household objective function, defined on quantities consumed Q, but also potentially on prices P and household budget y:

$$\begin{cases} \max_q H(P,y,Q) \\ \text{s.t } P'Q = y \end{cases}$$

This program is unitary if and only if H does not depend on P, y; otherwise it is not. If it is unitary, then we may take H to be a "household" utility function (whatever that means). Solutions of this programme are demand functions F, yielding quantities demanded by the household as functions

of market prices and household budget $Q = F(P, y)$.[1]

Given the assumption of efficiency, an observationally equivalent optimisation programme, *O2*, is a set of *decentralized* individual-level optimisations over shadow quantities for individuals facing shadow budget constraints:

$$\begin{cases} \max_{\widetilde{Q}^j} U^j(\widetilde{Q}^j) \\ \text{st } \widetilde{P}^{j\prime} \widetilde{Q}^j = \widetilde{y}^j \end{cases}$$

where $\widetilde{}$ denotes shadow objects (which may be unobserved), and j superscripts indicate people in the household. In this observationally equivalent representation, it is as though the household allocates a shadow budget $\widetilde{y}^j$ to each person j, who then spends their shadow budget at shadow prices $\widetilde{P}^j$. Solutions of the programme *O2* are shadow demand functions $\widetilde{F}$, yielding quantities demanded by the individuals as functions of shadow prices and shadow budgets $\widetilde{Q}^j = \widetilde{F}\left(\widetilde{P}^j, \widetilde{y}^j\right)$. The shadow quantities demanded are the arguments to the individuals's utility functions. Consequently, we may use shadow budgets $\widetilde{y}^j$ and shadow prices $\widetilde{P}^j$ to do welfare analysis at the person level in collective household models. A key feature of this welfare analysis is that individual utilities are increasing in shadow budgets, so if we find inequality in shadow budgets, it is related to inequality in individual utilities within the household.

We note that models like these easily accomodate caring preferences, where the utility level of one member enters the utility function of another. In *O1*, the household optimizes the household level problem accounting for such "double-counting".

In *O2*, the household allocates larger shadow budgets to members about whom other members care. Chiappori and Mazzocco (2017) state that caring preferences in *O1* always have a decentralized representation, and therefore always have a representation in *O2* with appropriate shadow budgets. Thus, these models allows parents to love their children, and spouses to love each other.

Since *O1* and *O2* are observationally equivalent, the observed quantities purchased by the household are also the quantities required to satisfy every member's shadow demand.

Let h denote household type, where for example $h = s$ is a household composed of a single individual and $h = mf$ is a couple. Let Q_h be the observed quantity vector purchased at market prices P by a household of type h when the household budget is y_h, $Q_h = F_h(P, y_h)$. Let $\widetilde{Q}_h^j$ be the shadow demand of person j given the preferences they have if they live in a household of type h, $\widetilde{Q}_h^j = \widetilde{F}_h^j\left(\widetilde{P}_h^j, \widetilde{y}_h^j\right)$. The equivalence between the centralised programme *O1* and the decentralised programme *O2* means that the household satisfies the shadow, person-level, demands $\left\{\widetilde{Q}_h^j\right\}_j$ of

---

[1]Refer to online appendix 9 for notations used in the paper.

the decentralised programme above, by buying observed quantities Q_h

$$Q_h = G_h \left(\left\{ \widetilde{Q}_h^j \right\}_j \right) \tag{1}$$

where G_h is a *household consumption technology* vector function that accounts for scale economies or shareability or publicness in household consumption. Because singles don't get scale economies, the function G_s is the identity function.

For non-shareable (or, private) goods:

$$Q_h = \sum_j \widetilde{Q}_h^j$$

For such goods, in order to satisfy the demands of each member, the household must purchase the sum of the shadow quantities demanded by all household members.

For shareable and/or public goods, $Q_h < \sum_j \widetilde{Q}_h^j$. Due to scale economies in household consumption, or to publicness of consumption, the household need only purchase an amount which is less than the sum of all household members's quantities demanded. We offer additional details in online Appendix 14.

Assignable goods are goods for which individual consumption is observed. Suppose there is one assignable good for each type of individual j, and suppose it is non-shareable. Use lower case for its quantity, price and demand function. For the non-shareable assignable good of individual j, the quantity demanded by the household is denoted q_h^j and is equal to the shadow quantity demanded by the individual, denoted $\widetilde{q}_h^j$:

$$q_h^j = \widetilde{q}_h^j \tag{2}$$

Let $f_h^j(P, y)$ be the element of the vector of demand functions of the household corresponding to the assignable good of person j and let $\widetilde{f}_h^j(P, y)$ be the element of the vector of shadow demand functions of individual j corresponding to the assignable good. Therefore, for a non-shareable assignable good, the following equality holds

$$f_h^j(P, y_h) = \widetilde{f}_h^j \left(\widetilde{P}_h^j, \widetilde{y}_h^j \right)$$

Shadow prices of consumption within the household, $\widetilde{P}_h^j$, are determined by scale economies (shareability of goods) and/or the publicness of each good. These within-household shadow prices differ from market prices, and depend on market prices and, possibly, on household income. For the non shareable assignable good, denote the market price as p and the shadow price as $\widetilde{p}_h^j$. Because

this good is non shareable, its shadow price is equal to its market price.[2]

$$\widetilde{p}_h^j = p.$$

For shareable (and/or public) goods, shadow prices are lower than market prices due to scale economies/sharing and/or publicness. Therefore the other elements of $\widetilde{P}_h^j$ do not satisfy an equality like this, and these shadow prices are lower than market prices.

We define the *resource share*, denoted η_h^j, to be the fraction of the total shadow budgets that goes to individual j if they live in a household h:

$$\eta_h^j = \frac{\widetilde{y}_h^j}{\Sigma_k \widetilde{y}_h^k} \tag{3}$$

Shadow budgets are chosen by the household as functions of market prices and household budget. So, resource shares $\eta_h^j(P,y)$ are also functions of prices P and budgets y. Thus, for example, $\eta_{mfc}^f(P,y)$ gives the resource share of an adult female, in a household of type mfc composed of a couple with a child, facing market prices P and with a household budget y. With abuse of notation, we will use the same greek letter η to denote either the resource share function $\eta_h^j(P,y)$ or its value η_h^j.

Because shadow budgets are chosen by the household for each member given market prices and the household budget, they are in general functions of prices and the budget. In all general-purpose collective household models that we are aware of (including the pure public/pure private model of Chiappori's 1992, the mixed public/private model of De Rock et al 2007, and some versions of the general consumption technology model of Browning et al 2013), the shadow budgets of the individuals in the household add up to the total budget of the household: $\Sigma_j \widetilde{y}_h^j = y$.

2.3   Relative resource shares and gender inequality

In this paper, we are interested in gender inequality. Consequently, we don't need to identify as much about resource shares as previous papers that have sought to identify the resource shares of all household members (for instance Dunbar, Lewbel and Pendakur (2013) or Lechene, Pendakur and Wolf (2022)). Instead, we seek to identify the *relative resource share* of women in a household of type h, R_h, defined to be the women's fraction of adult household resources (those consumed by

---

[2]Note that "private goods" and "non shareable goods" are identical, whereas "shareable goods" and "public goods" are goods which have in common some jointness in consumption, but are different, both in practice and conceptually. More on this in online appendix 14.

both adult men and adult women):

$$R_h = \frac{\widetilde{y}_h^f}{\widetilde{y}_h^m + \widetilde{y}_h^f} = \frac{\eta_h^f}{\eta_h^m + \eta_h^f}. \tag{4}$$

The children's resource share η^c (in households where they exist) is not our object of interest. Because the relative resource share doesn't depend on children's shadow budgets $\widetilde{y}_h^c$, it can be identified under weaker conditions than those that identify all the resource shares in a household. We target the relative resource share, R_h, in our theoretical work on identification and in our empirical work: it gives the share of the total (shadow budgets) of adults enjoyed by women in the household.

For male-female couples, if $R_h < 1$, then gender inequality favours men, if $R_h = 1$, there is no gender inequality, and if $R_h > 1$, then gender inequality favours women. For households with more than 1 man or woman, if R_h is less than the fraction of female adult members, then gender inequality favours men. The relative resource share R_h has less information about the structural model than the full set of resource shares does. In particular, it does not hold information about the children's share of household resources (η_h^c) and focusses directly on gender inequality amongst the adults.

In households comprised just of adult men and women ($h = mf$), the relative resource share R_{mf} is sufficient to identify the resource shares η_{mf}^j of both men and women, but in households with children ($h = mfc$), knowledge of the relative resource share R_{mfc} reveals gender inequality but is not sufficient to identify the resource shares η^j of any household member.

Unlike many identification results for collective household models (see especially Chiappori 1992), we identify both the level of the relative resource share function and its response to covariates (rather than just its response to covariates). For childless couples in particular, this means we identify the resource share and thus the shadow budget of both the man and the woman, which is equivalent to identifying the "sharing rule" (including its level) as defined in Chiappori (1992).

The relative resource share may be interpreted as a measure of gender inequality regardless of how children enter the household model. If children are people with utilities and resource shares, then the relative resource share tells how adult men and women share the pie that is left over after children take their piece. If children are household attributes or public goods, then the relative resource share tells how adult men and women share the total expenditure of the household. Either way, it speaks directly to gender inequality in the household.

2.4 Engel curves

Suppose, following Browning et al (2013), that the shadow price vector $\widetilde{P}^j = \widetilde{P}$ is the same for all household members.[3] This implies in a cross-section of households that face a common market price vector (P), all the individuals in those households also face a common price vector $(\widetilde{P})$. Consequently, we can use Engel curve data, that is cross-sectional data on households facing a common market price vector.

For a non shareable assignable good, equation (2) relates the quantities demanded by the household, at market prices P and household budget y to the shadow quantity demanded by individual j, at shadow prices $\widetilde{P}$ and shadow budget $\widetilde{y}$:

$$q^j(P, y) = \widetilde{q}^j\left(\widetilde{P}, \widetilde{y}\right)$$

Pre-multiply both sides of the equation above by the market price of the assignable good p and divide by household income y, to get the budget share for the assignable good:

$$\frac{p q^j(P, y)}{y} = \frac{p \widetilde{q}^j\left(\widetilde{P}, \widetilde{y}\right)}{y}$$

This equation can be rewritten to relate (observed) household budget shares to (shadow) individual budget shares. Note that, by definition of the resource share, $\widetilde{y} = \eta y$. Substituting y by $\frac{\widetilde{y}}{\eta}$ on the right hand side of the equation above yields:

$$\frac{p q^j(P, y)}{y} = \eta^j(P, y) \frac{p \widetilde{q}^j\left(\widetilde{P}, \widetilde{y}\right)}{\widetilde{y}} \tag{5}$$

Let $\widetilde{w}_h^j(\widetilde{P}, \widetilde{y}) = \frac{p \widetilde{q}^j(\widetilde{P}, \widetilde{y})}{\widetilde{y}}$ be the budget share that person j would choose if they had the preferences they have in household h and faced the shadow budget constraint $\left(\widetilde{P}, \widetilde{y}\right)$. On the left hand side is the budget share of the good, as arising from the optimisation of the centralised programme $O1$, w_h^j, while on the right hand side, we have the product of the resource share of individual j (η_h^j), by the shadow budget share of the assignable private good of j, ($\widetilde{w}_h^j$), as arising from the optimisation of the decentralised programme $O2$ for j.

A final step is to note that, within a given price regime, we can use the Engel curve version of

---

[3]Browning et al (2013) show that this is implied by a restriction on the household consumption technology given in equation (2). Specifically, if $Q_h = G_h\left(\left\{\widetilde{Q}_h^j\right\}_j\right) = A\left(\Sigma_j \widetilde{Q}_{hj}^j\right)$ for some matrix A, then $\widetilde{P}^j = \widetilde{P} = AP$. This household consumption technology allows each good to have a different degree of shareability (via diagonal elements of A), and for that shareability to be joint across goods (via off-diagonal elements of A). These parameters of the model do not have to be specified, known or estimated to identify relative resource shares.

equation (5) to identify the resource shares. With some abuse of notation, we simply drop the price argument from the Engel curve functions to get

$$w_h^j(y) = \eta_h^j(y)\widetilde{w}_h^j\left(\eta_h^j(y)y\right) \tag{6}$$

Without additional structure, resource shares $\eta_h^j(y)$ are not identified by Engel curve data on $w_h^j(y)$. For example, with $h = s, mf$ (singles and male-female couples) and $j = m, f$ (males and females), there are 4 Engel curves $w_h^j(y)$ that we may observe (on the left-hand side: $(w_s^f, w_s^m, w_{mf}^f, w_{mf}^m)$), but 5 unknown structural functions on the right-hand side: $\left(\eta_{mf}^f, \widetilde{w}_s^f, \widetilde{w}_s^m, \widetilde{w}_{mf}^f, \widetilde{w}_{mf}^m\right)$.

Browning, Chiappori and Lewbel (2013) solve this identification problem by assuming that $\widetilde{w}_h^j$ do not vary with h, so that there are only 3 unknown structural functions. Dunbar et al (2013) solve this identification problem by assuming that $\widetilde{w}_h^j$ have a component that does not vary with j for $h = mf$, similarly reducing the number of unknown structural functions. In this paper, we follow a strategy similar to that of Dunbar et al (2013), but we generalize their restriction to allow for greater preference variation across j within a household type h.

2.5 Identification of relative resource shares

In this section, we show that with a simple linear structure for the unobserved individual assignable good budget share functions $\widetilde{w}_h^j$, a shape restriction on how preferences vary across people, and a restriction on the resource share function, we can identify the relative resource share R_h from Engel curve data on assignable goods. In this case, the model implies a linear reduced form for observed household-level assignable good budget share functions w_h^j and the reduced-form coefficients are sufficient to construct relative resource shares. In the online Appendix, we show that the model is identified semi-parametrically, that is for unknown and nonparametric $\widetilde{w}_h^j$, under similar conditions.

Like Browning, Chiappori and Lewbel (2013) and Bargain, Donni and Hentati (2023), we make use of the observed behaviour of singles. Unlike those authors, we do not assume that the preferences for any good are identical for people who live alone and people who live in households.

First, we assume that the model is an efficient collective household as described in section (2.2), and that we have data on the non shareable assignable goods budget shares of males and females (people types m, f) as single individuals and in couples (household types s, mf). This means that the household demands exactly what each individual demands as in equation (2).

Second, we assume the shadow price vector $\widetilde{P}^j = \widetilde{P}$ is the same for all household members as described in section (2.4). This implies that in a cross-section of identical households facing the same price vector, all the individuals living in those households face the same shadow price vector. It also implies that, for the non shareable assignable goods, household Engel curves are given by

13

(6).

Third, we assume that a reference type of person (we will say f) has PIGLOG preferences (see Muellbauer 1975) so that this person's shadow budget shares for non shareable assignable, (and actually all) goods, w_h^f, are linear in the log of the shadow budget $\widetilde{y}$:

$$\widetilde{w}_h^f(\widetilde{y}) = \alpha_h^f + \beta_h^f \ln(\widetilde{y}).$$

This assumption is not necessary for identification, but, in combination with the next two restrictions, it gives the model a linear reduced form so that estimation is easy and transparent. In online Appendix 12, we show that identification is possible without any restriction on the preferences of the reference person.

Fourth, we impose a shape restriction on how preferences vary across people called weakened similarity across people (WSAP). Under this assumption, preferences for other types of people (in our case m, c) are similar but not identical to those of the reference person (f). They are related by translation in w_h^j and $\ln(\widetilde{y})$ and by stretching in $\ln(\widetilde{y})$, and we detail the family of preferences and utility functions permitted in online Appendix 12. However, in the case where reference preferences are PIGLOG, WSAP implies that all people have PIGLOG preferences

$$\widetilde{w}_h^j(\widetilde{y}) = \alpha_h^j + \beta_h^j \ln(\widetilde{y}).$$

and that the slope of Engel curves for a person j living in a household h, β_h^j, is multiplicatively decomposable as the product of two terms, one which only depends on the type of household h and one which only depends on the type of individual j:

$$\beta_h^j = \delta^j \beta_h.$$

This implies that the ratio of structural slopes $\beta_h^f / \beta_h^m = \delta^f / \delta^m = \delta^j$ is independent from the household type in which people live.

We call this restriction on preferences *weakened* similarity across people, WSAP, because in this PIGLOG context, the assumption of similarity across people (SAP restriction) of Dunbar, Lewbel and Pendakur (2013) would require δ^j the same for all people.

We normalize

$$\delta^m = 1$$

and denote

$$\delta = \delta^f$$

so that $\beta_h^m = \beta_h$ and $\beta_h^f = \delta \beta_h$.

Fifth, we assume (like Dunbar, Lewbel and Pendakur (2013)) that resource shares are independent of the household budget y, satisfying $\eta_{mf}^j(p,y) = \eta_{mf}^j$. Then, because we are holding prices constant and studying behavior at the Engel-curve level, we have that

$$\eta_{mf}^j(p,y) = \eta_{mf}^j(y) = \eta_{mf}^j$$

This fifth assumption, combined with the PIGLOG preferences $\widetilde{w}_h^j(\widetilde{y})$ and our household model, gives a reduced form for household assignable good budget shares that is linear in the variables.

Under the five assumptions given above, for singles $\eta_s^j(y) = 1$, so that the structural parameters for singles' Engel curves are directly revealed by the reduced form. That is,

$$w_s^j(y) = a_h^j + b_s^j \ln(y)$$

where $a_s^j = \alpha_s^j$ and $b_s^j = \delta^j \beta_s$. Given that $\delta^m = 1$ and $\delta = \delta^f$, we identify

$$\delta = b_s^f / b_s^m$$

from the slopes of singles' Engel curves.

For individuals in households, the observed Engel curve is equal to the product of the resource share by the shadow Engel curve (equation (6)). Substituting the third, fourth and fifth identifying restrictions into (6) yields the structural assignable goods budget share functions of households:

$$W_h^j(y) = \eta_h^j \left(\alpha_h^j + \delta^j \beta_h \ln\left(\eta_h^j y \right) \right) \tag{7}$$

We may rewrite this as a linear reduced form

$$W_h^j(y) = a_h^j + b_h^j \ln y \tag{8}$$

where

$$a_h^j = \eta_h^j \alpha_h^j + \eta_h^j \beta_h^j \ln\left(\eta_h^j \right)$$

and

$$b_h^j = \eta_h^j \delta^j \beta_h$$

Taking the ratio of slopes of Engel curves for women and men in couples, we get:

$$\frac{b_{mf}^f}{b_{mf}^m} = \frac{\eta_{mf}^f \delta^f}{\eta_{mf}^m \delta^m} = \frac{\eta_{mf}^f}{\eta_{mf}^m} \delta$$

15

This equation says that if the ratio of slopes of Engel curves for individuals in couples (the left hand side) is greater than the ratio of slopes for singles (δ), it means that the resource share of women in couples is more than 50% ($\eta_{mf}^f/\eta_{mf}^m > 1$). The value-added of our model framework is that it allows us to interpret the slopes of Engel curves—reduced form objects—in terms of resource shares and other structural parameters.

The object we are interested in, the relative resource share of women in household type h, can be identified from the Engel curves for assignable goods for singles and individuals in household type h. From equation (4) and simple algebra, the relative resource share R_h is

$$R_h = \frac{\eta_h^f}{\eta_h^m + \eta_h^f} = \frac{b_h^f}{b_h^f + b_h^m \delta} \tag{9}$$

where the b_h^j are the slopes of Engel curves for assignable goods for individuals of type j in households of type h and δ is the ratio of slopes of singles Engel curves.

2.5.1 Remarks on identification

Although the restriction that resource shares are independent from the household budget is not testable in our context, it is often used in this literature and has some empirical support. Menon et al (2013) and Cherchye et al (2015) find that the restriction is not rejected in Italian and UK data, respectively. Hsieh (2025) finds that the restriction is rejected in US data, but that the magnitude of the dependence of η on y is very small. Bargain et al (2022) find that the restriction is rejected in UK data, and that the magnitude of the estimated dependence is large. However, their rejection is conditional on a model where the resource share function is fixed over time. In this paper, we do not impose that restriction.

In our model, relative resource shares are identified from relative *marginal* responses to income in the household Engel curves for private assignable goods. In particular, the reduced form coefficients b_h^j say how much of a luxury the assignable good is for singles and for households. Suppose that, in the data at hand, information on household expenditure on clothing is available for men and women. Suppose that for singles, clothing is twice as much a luxury for women as it is for men, with an Engel curve twice as steep. Now, suppose that this is the same for couples. One would then conclude that, in couples, women's claim to resources is the same as men's, hence women's and men's resource shares are 50%. However, if you found that, in couples, clothing is more than twice as much a luxury for women as for men, you would conclude that women in couples have a larger claim to resources than do men in those couples, hence a resource share higher than 50%.

Our approach allows for a lot of preference variation across people and the households in which they live. The parameters α_h^j vary arbitrarily across people and household types and are not used in the identification of relative resource shares. Suppose that women demand more clothing than men given any budget. This would not imply necessarily that they have larger resource shares, because resource shares are identified from slopes of Engel curves, not from levels of Engel curves. These slopes relate to the degree to which the assignable good is a luxury for a person. Further, suppose that women find clothing to be more of a luxury than do men. This too does not necessarily imply that they have larger resources shares, because relative resource shares are identified by the relative slopes of Engel curves within a household, compared with the relative slopes of Engel curves for singles. This could arise if the non shareable assignable good was more of a luxury for women in household type h relative to men in household type h than for single women relative to single men, $\frac{b_h^f}{b_h^m} < \frac{b_s^f}{b_s^m}$.

Identification of the model requires that the denominator of the expression for the relative resource share is non-zero: [4]

$$b_h^f + b_h^m \delta \neq 0.$$

In our empirical work, we will assess how far the estimate of $b_h^f + b_h^m \delta$ is from 0, to see if we can reasonably take the model to be identified.

Above, we describe the case where a reference person has PIGLOG preferences, and we will use this model in our estimation. However, in Theorem 1 in online Appendix (12), we show that this strong restriction on preferences is not necessary to identify the relative resource share. In fact, no restriction on preferences of the reference person is necessary to identify relative resource shares. WSAP is a shape restriction on how preferences vary across people. In that Appendix, we characterise the class of Engel curves and associated utility functions that are sufficient to identify relative resource shares. This class of Engel curves may have any shape (not just log-linearity), as long as the shapes satisfy certain restrictions across people and household types. We use the term "weakened" similarity across people because the semiparametric restrictions we invoke are weaker than those invoked by the "similarity across people" restriction of Dunbar, Lewbel and Pendakur (2013).

In the case where preferences are PIGLOG, WSAP is identical to the "similar ratios across people" condition defined by Arduini (2025) and similar in spirit to the "similarity over time" restriction used by Sokollu and Valente (2022). However, in the case where preferences are outside the PIGLOG class, those alternative restrictions are not defined but WSAP still has meaning and could be used. For example, if Engel curves are quadratic in the log of expenditure, the restrictions of Arduini (2025) and Sokollu and Valente (2022) are not defined and so cannot be used, but

---

[4]This is analogous to the lack of identification of the SAP estimator when $b_h^f + b_h^m$ is equal to 0 (see Dunbar, Lewbel and Pendakur 2013 and Lechene, Pendakur and Wolf 2022).

WSAP is defined and implies restrictions on quadratic terms across people and households (rather than restrictions on linear terms) (see Lechene, Pendakur and Wolf (2022) Appendix 6.3).

We note that the model is testable if there exist two assignable goods. The spirit of the test follows from the fact that if there are two assignable goods, they would each be sufficient to identify relative resource shares. The model would demand that the resource shares thus identified are the same regardless of which assignable good is used. A test of this kind is presented in Lechene et al (2022).

Alternatively the model could be testable if there existed time use data as well as expenditure data. Arduini (2025) explores what can be learned from time use data alone.

In this section, we have shown how to identify the relative resource shares in the specific case of PIGLOG preferences and without time or covariates. Identification is possible with a single-cross section of singles and households, and may be thought of as conditional on covariates. We discuss how to include time and covariates (cf. section (3.1) below, essentially by interacting all model parameters with time and observed covariates.

To situate our identification strategy within the literature, we provide a review of existing approaches to identification of resource shares in online Appendix (10). In online Appendix (14), we show that our strategy identifies relative resource shares that are compatible with a large class of models of consumption technologies and therefore of public or shareable goods.

3 Empirical implementation

In this section, we discuss the introduction of observed heterogeneity (household and person characteristics) in the model, as well as the introduction of time, in a context where we will be estimating the model with a set of repeated cross sections in (3.1). We then discuss how to restrict the model to allow for linear (Oaxaca-style) decompositions in (3.2). We finally discuss how we estimate the structural parameters from the estimated reduced forms in (3.3).

3.1 Covariates, Time and Approximation with Linear Reduced Forms

In the exposition above, with no covariates, the expression for the relative resource share is exact. However, if resource shares η_h^j and preference parameters β_h both depend on covariates z linearly, then equation (7) says that the reduced form coefficient b_h^j would be quadratic in z. More generally, it would be an unknown function of z. We deal with this by approximating the true dependence on z with a linear index in z, in the same way as Lechene, Pendakur and Wolf (2022). Since $b_h^j = \eta_h^j \delta^j \beta_h$, for b_h^j to be exactly linear in z, two of η_h^j, δ^j and β_h need to be independent of z and one of them needs to be linear in z.

Let z_i be a vector of covariates, varying across households, that affect resource shares η_h^j, the household consumption technology and therefore shadow prices $\widetilde{p}$ and preferences $\widetilde{w}_h^j$. Let t be time, which affects prices, resource shares, the household consumption technology and preferences. Like the cross-sectional work of Lechene, Pendakur and Wolf (2022), we linearly approximate the dependence of reduced form parameters on z. Our setting is repeated cross-sectional, so we additionally include time-varying coefficients on regressors to account for time t. Let ε_{ist} be an error term in the reduced form equation.

For the assignable good of a person of type j with characteristics z in time period t living in household of type h, we use the linear reduced form:

$$W_{iht}^j = a_{iht}^j + b_{iht}^j \ln y_i + \varepsilon_{iht}^j$$

where

$$a_{iht}^j = a_h^j(t) + a_h^{zj}(t)'z_i, \tag{10}$$

and

$$b_{iht}^j = b_h^j(t) + b_h^{zj}(t)'z_i. \tag{11}$$

The terms depending on t are time-varying coefficients, modeled empirically as smooth functions of time—splines for $a_j^j(t)$ and $b_j^j(t)$ linear or quadratic in time for $a_j^{zj}(t)$ and $b_j^{zj}(t)$. For example, for quadratic time-varying coefficients $a_j^{zj}(t)$, we include regressors z, tz, and $t^2 z$ and for quadratic time-varying coefficients $b_j^{zj}(t)$, we include regressors $z \ln y_{th}$, $tz \ln y_{th}$, and $t^2 z \ln y_{th}$.

Let $\widehat{b}_{iht}^j$ be an estimate of b_{iht}^j. An estimate of the relative resource share is given by the ratio

$$\widehat{R}_h = \frac{\widehat{b}_{ith}^f}{\widehat{b}_{ith}^f + \widehat{b}_{ith}^m \widehat{\delta}_{it}} \qquad \text{where} \qquad \widehat{\delta}_{it} = \frac{\widehat{b}_{ist}^f}{\widehat{b}_{ist}^m}. \tag{12}$$

Here, a key feature is that, since R_h is identified cross-sectionally for each z, we can identify how the resource share function itself changes over time (rather than just how changes in demographics z might drive change over time). Note that since $\widehat{b}_{ist}^j$ depend on covariates z and time t, the estimate $\widehat{\delta}_{it}$ generally depends on both z and t.

3.2 Model Restrictions and Oaxaca Decomposition

Since the estimated relative resource share $\widehat{R}_h$ in equation 12 is equal to a ratio of $\widehat{b}_{ith}^j$, each of which is linear in z_i, it is generally nonlinear in z_i and therefore not amenable to Oaxaca-style decompositions that break changes over time into a part driven by time t and a part driven by demographics z. To facilitate such decompositions, we impose two restrictions that, together,

make the denominator a function of time but not of demographics. With these restrictions, at each time t, the estimate of $\widehat{R}_h$ is linear in z_i. With this, we can use standard Oaxaca decompositions to break the overall change in relative resource shares over time into a part driven by changing demographics and a part driven by the changing function over time, analogous to decompositions of wage disparity into a part driven by characteristics and a part driven by different wage functions across groups.

First, we impose the restriction that all slope coefficients for singles are proportional to each other for males and females:

$$b_s^f(t) = db_s^m(t) \qquad \text{and} \qquad b_s^{zf}(t) = db_s^{zm}(t) \quad \text{for all } z \qquad (13)$$

for some scalar constant d. This implies that

$$\delta_{it} = \frac{b_{ist}^f}{b_{ist}^m} = \frac{b_s^f(t) + b_s^{zf}(t)'z_i}{b_s^m(t) + b_s^{zm}(t)'z_i} = \frac{db_s^m(t) + db_s^{zm}(t)'z_i}{b_s^m(t) + b_s^{zm}(t)'z_i} = d = \delta. \qquad (14)$$

The structural parameter $\delta_{it} = \delta$ (which itself equals the reduced form parameter d) is invariant over time t and demographics z_i. To estimate relative resource shares, all we need from singles' Engel curves is an estimate of δ.

Second, we impose the restriction on household Engel curves that the marginal effects of men's and women's covariates cancel each other out in the denominator of the resource share,

$$b_h^{zf}(t) + b_h^{zm}(t)\delta = 0 \qquad (15)$$

Given these two restrictions, we get the following expression for the women's relative resource share

$$R_{iht} = \frac{b_{iht}^f}{b_{iht}^f + b_{iht}^m \delta}.$$

We may therefore form an estimate of R in terms of estimated reduced form coefficients and an estimate $\widehat{\delta}$ of δ,

$$\widehat{R}_{iht} = \frac{\widehat{b}_h^f(t) + \widehat{b}_h^{zf}(t)'z_i}{\widehat{b}_h^f(t) + \widehat{b}_h^m(t)\widehat{\delta}}. \qquad (16)$$

The key feature here is that the denominator depends on time t but not demographics z.

Here, the average relative resource share in a time period t equals the resource share function

evaluated at average characteristics at time t, $\bar{z}_{ht}$:

$$E\left[R_{iht}|t\right] = E\left[\frac{\left(b_h^f(t) + b_h^{zf}(t)\right)z_{iht}}{b_h^f(t) + b_h^m(t)\delta}\bigg|t\right] = E\left[\kappa_h(t) + \lambda_h(t)'z_{iht}|t\right] = \kappa_h(t) + \lambda_h(t)'\bar{z}_{ht} \qquad (17)$$

where

$$\kappa_h(t) = \frac{\widehat{b}_h^f(t)}{\widehat{b}_h^f(t) + \widehat{b}_h^m(t)\widehat{\delta}} \qquad \text{and} \qquad \lambda_h(t) = \frac{\widehat{b}_h^{zf}(t)}{\widehat{b}_h^f(t) + \widehat{b}_h^m(t)\widehat{\delta}}$$

We can use Oaxaca (1973) decompositions to analyze changes over time by writing the overall change $R_{2019}(\bar{z}_{2019}) - R_{1978}(\bar{z}_{1978})$ as the sum of two pieces:

$$R_{2019}(\bar{z}_{2019}) - R_{1978}(\bar{z}_{1978}) = [R_{2019}(\bar{z}_{2019}) - R_{2019}(\bar{z}_{1978})] + [R_{2019}(\bar{z}_{1978}) - R_{1978}(\bar{z}_{1978})] \quad (18)$$

The first part, $[R_{2019}(\bar{z}_{2019}) - R_{2019}(\bar{z}_{1978})]$, gives the part of the overall change due to the change in average characteristics, evaluated at the relative resource share function of 2019. The second part, $[R_{2019}(\bar{z}_{1978}) - R_{1978}(\bar{z}_{1978})]$, gives the part of the overall change due to the change in the relative resource share function, evaluated at the average characteristics of 1978. We provide tables like this in the empirical work below in section 5.5. We note that this analysis is possible because our methodology identifies how the relative resource share function changes over time. The fact that we can identify how that function changes from one period to another is a key difference between this paper and the previous models that have been used to estimate resource shares over time (namely the models of Lise and Seitz (2011) and Bargain et al (2022)).

The same decomposition holds at each t. We can therefore similarly ask how much of the difference between women's relative resource share at time t and women's relative resource share of 1978 is due to the change in average characteristics, evaluated at the relative resource share function of t and how much is due to the change in the relative resource share function, evaluated at the average characteristics of 1978.

$$R_t(\bar{z}_t) - R_{1978}(\bar{z}_{1978}) = [R_t(\bar{z}_t) - R_t(\bar{z}_{1978})] + [R_t(\bar{z}_{1978}) - R_{1978}(\bar{z}_{1978})] \qquad (19)$$

We will present graphs with these types of comparisons in our empirical work.

We note that, with the 2 restrictions introduced above, in order to estimate relative resource shares, the only information needed from singles' Engel curves is the ratio δ, and the estimate of this ratio is statistically independent of all parameters in households' Engel curves. We therefore use a 2 step estimation procedure. We first estimate δ from singles, and then use household Engel curves and the estimate of δ to estimate women's relative resource shares.

Apart from allowing easy decomposition, linearity of the estimated relative share in

z has 3 other nice benefits worth noting. The first benefit is related to an identifying restriction of the model noted above. The relative resource share function is identified if the denominator is not zero. The denominator of the relative resource share (16) is equal to $b_h^f(t) + b_h^m(t)\delta$. It is a function of time *t* (and not characteristics *z*), and so we can graph the estimate of this object over time to easily check if the model is identified.

A second benefit relates to the sampling variability of the predicted values. When we come to measuring inequality, we will use predicted values of relative resource shares for each household. If these predicted values have a lot of spurious variability due to sampling variability of the estimated relative resource shares, it will inflate our estimates of inequality. If we don't impose the second restriction, the dependence of the denominator on characteristics *z* could lead to some observations where the denominator is very close to zero, yielding "wild" predicted values of resource shares for those households. These observations would have an undue influence on estimated inequality. In contrast, with the restrictions, once we confirm that the denominator is far from zero for all time periods, we can use the predicted values with confidence in our investigation of inequality.

A third benefit of linearity is that the marginal effects of *z* on relative resource share are constants in a given time period, and therefore easily graphed as functions of time. The marginal effect of *z* on the relative resource share is the time-varying (but nothing else-varying) function $\lambda(t)$. We show graphs of such marginal effects in our empirical work below.

We estimate models maintaining these 2 restrictions, which together result in a relative resource share function that is linear in the covariates in each time period. Empirically, we show in online Appendix (16) that the first restriction, given by equation (13), is not rejected in our data, and the second restriction given by equation (15) is strongly rejected in our data. The benefits of these restrictions we outlined above. The cost of imposing one rejected restriction is misspecification error. We maintain both restrictions in our main text analysis because it greatly facilitates transparently answering our basic questions about whether or not changing demographics are responsible for changing relative resource shares.

For the cautious reader, we provide estimates that relax the second restriction while maintaining the first in online Appendix (16). These estimates are less precise but still reveal the headline results of our main text analysis. In particular, we still see that the resource shares of childless women rose during the 1980s and 1990s, driven roughly half by an increase in the relative resource share function. And we see that equalization within households canceled out some of the rising inequality across households over this period.

3.3 Estimation

We assume that expenditures on the private assignable good W_h^j, for individuals of type j in households of type h total budgets y and relevant covariates z are observed for iid samples of both singles and households. Consequently, the reduced form parameters $a_h^j(t), a_h^{zj}(t), b_h^j(t), b_h^{zj}(t)$ are identified from linear regressions. For our baseline model, we use restricted cubic splines with 6 basis functions (including the constant), restricted so that boundary splines are linear, for $a_j^j(t), a_h^j(t), b_j^j(t), b_h^j(t)$ and linear time-varying coefficients for $a_j^{zj}(t), a_h^{zj}(t), b_j^{zj}(t), b_h^{zj}(t)$. In the online Appendix, we consider longer splines (9 basis functions) and quadratic time-varying coefficients. We use linear time varying coefficients in the main text, because it turned out *ex post* that none of the quadratic time terms are statistically significant and the qualitative results are the same with linear and with quadratic time varying coefficients. Similarly, a longer spline, while statistically significant, does not affect our qualitative story.

We recover estimates of our structural parameters from linear regressions of private assignable goods Engel curves on

$$(1, \text{splines}, \quad z, tz, \quad \ln y, \text{splines} * \ln y, \quad z \ln y, tz \ln y)$$

for each household type.

The model allows for time-varying preferences and time-varying resource share functions. Time variation includes the effect of prices, of the year of survey as well as any other relevant time-varying aspect of the environment.

In the online Appendix (16), we consider robustness by presenting estimates with 9 spline basis functions and quadratic time-varying coefficients. With the more complex splines, the estimates show the same patterns as in the main text, but are (predictably) somewhat less precise.

4 Data

We use the UK Living Costs and Food Survey (LCFS) (previously known as the Family Expenditure Survey (FES), the Expenditure and Food Survey (EFS) and the National Food Survey (NFS)), a repeated cross-section, from 1978 to 2019. This is a standard household budget survey, recording expenditure on assignable clothing for men, women and children, as well as the usual information required to estimate Engel curves. There are about $7,000$ households per wave of the survey, totaling 275,194 households in the years that we use.[5]

The sample we use for estimation of relative resource shares keeps households: a) in regions other than Northern Ireland (263,325 remain); b) where ages of children are available and children

---

[5]Although the survey starts in 1957, there is no information on the respondents' education prior to 1978.

Figure 1: Fraction of different types of households over time
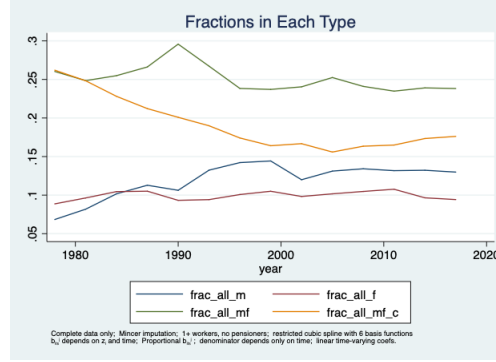
are aged less than 16 (250,657 remain); c) that are unmarried or where marriage is not a same sex marriage (250,181 remain); d) that have less than 5 children and less than 4 adult men and less than 4 adult women (248,663 remain); e) where head and spouse information are valid (240,808 remain); f) the average education leaving age of either adult men or adult women is less than 26 (236,976 remain); g) the average age of adult men and of adult women is less than 65 (169,014 remain); h) data on benefits income is present (168,981 remain); i) where at least one adult in the household is in work (aka "employed") (139,639 remain). Since the big restrictions are the exclusion of pensioners and of households not in work, we refer to this sample as *complete data; at least one member in work; no pensioners*.

Figure 1 shows the fraction of different types of households in the data over time. The proportion of single households as well as households of couples without children has slightly increased over time, whereas the proportion of households with couples with children has fallen. We present results for resource shares of couples with and without children separately. Singles are used for estimation as described above.

Our estimate of δ uses data from singles. Our estimates of the relative resource share R_h use data on people living in 3 household types: married couples with no children and no additional adults; married couples with at least 1 child but no additional adults; and, all other households with at least 1 man and 1 woman. We generally present figures for just the first 2 household types.

We allow for covariates z_i in the level (a_{iht}^j) and in the budget-response (b_{iht}^j) of the individual Engel curves. These covariates are: the average age of men in the household, the average age of women in the household, the average of school-leaving age of men in the household, the average of school-leaving age of women in the household, wage ratio (equal to the average women's hourly wage in the household divided by the average men's hourly wage in the household), the benefit share of household gross income, the number of men, women and children in the household, and an indicator that there is a child under 6 years old in the household.

We note that some of these variables are constants for some household types. For example, for

24

Figure 2: Sample mean of characteristics over time

(a) Wage ratio

(b) Share of household income from benefits

childless couples, the indicator that there is a child under 6 is always zero. However, this is not problematic because we use separate linear models for each of the 3 household types, and stack the equations in the regression.

We also allow for additional covariates in the level (a_h^j) of the Engel curves, with monthly and regional effects on spending (e.g., The North; December).

The wage ratio is equal to the ratio of the observed woman's to man's hourly in households where both hourly wages are observed. For all people who have zero or negative labour income or who are self-employed—and therefore do not have an hourly wage–we impute their wage using a standard Mincer regression with education and age regressors. As a robustness exercise, we also estimate models that replace the wage ratio with the market income ratio (see online Appendix (16)).

The left hand panel of Figure 2 shows the wage ratio over time for different types of households. In childless couples, the wage ratio went from about 75% to almost parity. In households with children, it went from about 65% to almost parity.

The right hand panel of Figure 2 illustrates the share of household income from benefits over time. The share of benefit income is small in childless households (between 1.5% and 4%) and larger in households with children (between 6% and 12%), since many benefits for households are linked to the presence of children. There is no constant trend in the evolution of this characteristic. It fluctuates over the period, due to policy changes affecting the relative generosity of benefits over time. For example, there were discretionary increases to level of benefits in the 1990s and early 2000s, whereas the 2010s saw cuts to benefits due to austerity measures (Belfield et al., 2017).

Figure 3 shows the mean age at which women and men left education over time. The mean age women left education has increased from less than 16 to close to 19 years old for women without children and from 16 to over 19 years old for mothers. The mean age at which men left education has increased from less than 16 to just above 18 for men in childless couples and from about 16 to

Figure 3: Sample mean of age left education over time

(a) Women

(b) Men

Figure 4: Budget responses
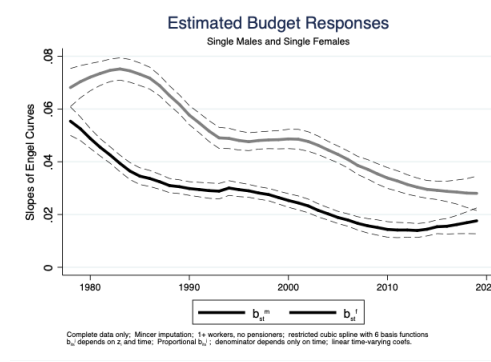
just below 19 for men in couples with children.

We often evaluate relative resource shares at the average value of covariates z_i in a given year t, which we denote as $\bar{z}_t$.

5   Results: Women's relative resource shares

We display most of the results in the form of graphs of objects of interest rather than in the form of tables of numbers. Generally, we provide 2 graphs for every object of interest, with one graph for each of the two most prevalent household compositions: childless couples (29% of the households in the sample) and couples with children (36% of the households in the sample). In most cases, we also include 90% confidence intervals, represented with dotted lines.

5.1   Preliminaries

We start by illustrating the estimated budget responses for clothing, in other words the estimated slopes of the Engel curves. Figure (4) gives the estimated budget responses, $b_{ist}^{j} = \widehat{b}_s^{j}(t) + \widehat{b}_s^{zj}(t)'z_i,$

for clothing budget shares for single males and single females without imposing the restriction (15). They are evaluated at the mean value of covariates for single males and single females, respectively. In the Figure, we see that for singles, estimated budget responses are positive throughout the period, meaning that clothing is a luxury for both single men and single women. However, the budget response is larger for women than for men throughout the period, meaning that clothing is more of a luxury for women than it is for men. This means that the SAP restriction of Dunbar, Lewbel and Pendakur (2013) does not hold for singles, because it requires these slopes to be identical across people (men and women) within a household type (singles). However, WSAP might hold.

Considering the restriction (15), the eyeball test suggests that the ratio between the budget responses of men and women might indeed be fixed over time. Formal tests in online Appendix (16) show that in these data we cannot reject the restriction that the ratio $\delta_{it} = b_{ist}^f / b_{ist}^m$ is fixed over time t and across household characteristics z_i (and so equals a constant δ). Consequently, in our main text results, we impose the restriction (14). We estimate δ by nonlinear GMM estimation of the model with linear budget shares on data for singles, under the (nonlinear) restriction that the slopes of the Engel curves for single men and single women are proportional to each other in every year. Our GMM estimate of δ is 1.80 with an estimated standard error of 0.068.

In online Appendix (16), we discuss the test of this restriction and in online Appendix (13.2), we discuss various ways to estimate δ. In particular, we offer a two-step linear estimator of δ which amounts to a minimum distance estimator. This estimate is 1.78 with a standard error of 0.26, which is very imprecise compared to our GMM estimate, so so we use the GMM estimate in our main text results.

In Figure (5) below, we present the estimated denominator in the expression for relative resource shares, equal to $\widehat{b}_h^f(t) + \widehat{b}_h^m(t)\widehat{\delta}$, as a function of time t for childless couples and for couples with children. An identifying restriction of the model is that this quantity be nonzero. In the case where the denominator equals zero, relative resource shares are not identified because observed behaviour would be consistent with any value of the relative resource share. In the Figure, we can see that the denominator of the expression for relative resource shares is statistically significantly positive throughout the period, though it does become small in the last decade.

Although the model is identified, the fact that denominators get small is consequential. As is visible in the graphs, the standard errors of the estimated denominators do not change much over time. Since relative resource shares are ratios, the sampling variability of the estimated relative resource share increases with the proportional variability of the denominator, and the proportional variability of the denominator is increasing over time. Consequently, the estimated sampling variability of estimated relative resource shares increases mechanically over time.

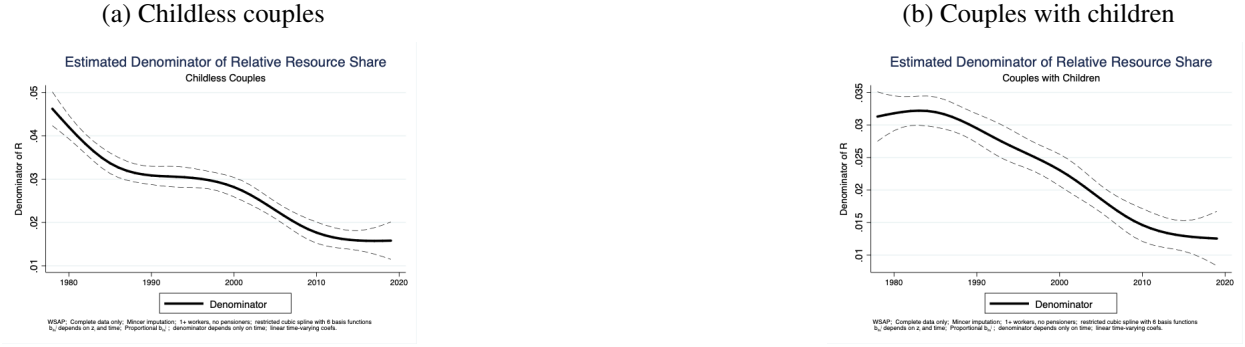Figure 5: Estimated denominator of relative resource shares

(a) Childless couples

(b) Couples with children

Figure 6: Estimated average relative resource shares

(a) Childless couples

(b) Couples with children

5.2 Women's average relative resource shares 1978-2019: $R_t(\bar{z}_t)$

Figure 6 shows the value of the women's estimated relative resource shares, evaluated in each year at the average value of household characteristics z for that year, denoted $R_t(\bar{z}_t)$. The left hand panel shows average resource shares for women in childless couples, and the right hand side panel for women in couples with children. Because estimated relative resource shares are linear in characteristics (under restriction (15)), this object is also the average of women's relative resource shares across households in a given year.

It is worth noting that the data is repeated cross-section, which means that we are describing the situation of different childless women and different mothers, in different time periods, rather than the evolution of the relative resource share of women as they go through the life-cycle. With such data, our model can inform of changes in household sharing through time for different cohorts and different family compositions, but not of changes in household sharing as time passes, or as household composition changes, for any cohort or group of women.

On the left-hand side graph, we see there seems to have been three distinct periods in the evolution of women's relative resource shares in childless couples. Between 1978 and the mid-1990s, the average woman's relative resource share in these households rose by about 11 percentage points.

28

This is a large change from substantial gender disparity to roughly gender parity. This change in the early decades is statistically significant: the confidence intervals in 1978 and the mid-90s don't overlap; and a formal test of equality of the resource share in 1978 and 1995 rejects at the 1% level. This initial period of increase in relative resource shares is followed by about a decade of slightly decreasing relative resource shares and by a subsequent increase in the point estimate since 2010, bringing the relative resource shares of childless women to above 55%. In the final period, from 2010 to 2019, the confidence intervals are consistent with either an increase, a stagnation or even a decrease in women's relative resource shares. The reason that the confidence intervals are wide at the end of the period is mechanical: the denominator of the relative resource shares decreases towards the end of the period, but its standard error remains roughly constant.

Overall, between 1978 and 2019, the increase in the relative resource share of women in childless couples is about 14 percentage points and marginally statistically significant. However, the increase of 11 percentage points from 1978 to 1995, where all estimates are precise, is strongly statistically significant.

On the right-hand side, we see that the women's relative resource share in couples with children did not have as much of a pronounced time trend. If anything, the women's relative resource share in these couples was stable before 2010, and stable or rising after 2010. Looking at the point-estimates, the increase after 2010 was substantial, showing an increase of about 11 percentage points in mothers's relative resource shares. However, as in the case of childless women, the increase after 2010 is not statistically significant. Despite the wide confidence intervals, we believe the increase in estimated relative resource shares at the end of the period might reflect an actual change in within household gender inequality, since we observe it for both household types. As with childless women, this increase brought them from substantial gender disparity towards near gender parity. But, as can be seen via the confidence intervals, we cannot reject the hypothesis that mothers's relative resource shares were the same in 2019 as they were in 1978.

The salient facts that emerge from the estimates of women's relative resource shares are: 1) different temporal patterns of change for mothers and for childless women; 2) for childless women, the period of increase in relative resource shares is early (1978-1995) and for mothers, it is late (2010-2019); 3) a 14 percentage point increase in childless women's relative resource shares from below parity (37%) to parity (51%) between 1978 and 2019; 4) an insignificant increase of around 10 percentage points in mother's relative resource shares, from below parity (40%) to parity (50%).

5.3   Distribution of Women's Relative Resource Shares $R_t(z_{it})$

In addition to documenting the average relative resource shares, we can also illustrate the distribution of women's relative resource shares over time. Figure 7 shows the distribution of women's

Figure 7: Distribution of women's relative resource shares

(a) Childless couples

(b) Couples with children

estimated relative resource shares in each year from 1978 to 2019, with a scatter plot showing the distribution of the resource shares (values corresponding to the left hand side axis), and a line showing the standard deviation of resource shares over time (values corresponding to the right hand side axis). Both for women in childless couples (left-hand graph) and mothers (right-hand graph), we see substantial variation in relative resource shares in any given year.

The central tendency in each year follows, by construction, the patterns observed in Figure 1, for the average values of the relative resource shares, since they are linear in characteristics z. That is, the means of these scatter plots are the lines shown in Figure 1. Both for women in childless couples and for mothers, the variation in resource shares increases over time, which will have an implication on person-level inequality (discussed in section (6)).

The scale of the increase in the value of the relative resource share over time (14 percentage points for childless women and 11 percentage points for mothers) is close to the 90-10 ratio of resource shares as they vary across demographics in any given time period.

The evolution of the distribution of relative resource shares addresses the question of whether the increase in average values of the resource share reflects a rising tide raising all boats, or whether only some women experience an increase. For childless women, there seems to be a rising tide in the first part of the period (1978 to 1995, when the distribution of resource shares moves up) followed by a period during which the distribution seems more constant, from 1995 to 2010. In the latter part of the period, from 2010 to 2019, the distribution fans out, so that it could be that resource shares increase more for women at the top of the distribution of resource shares (which is possible, although the increase in the resource shares over the end of the period is not significant).

It is worth noting that 99.9% of predictions are in $[0, 1]$, which constitutes an informal test of the model, since nothing in the econometric model constrains the relative resource shares to be in this range. For relative resource shares to be outside the $[0, 1]$ range, the signs for the slopes of men and women's Engel curves would have to be different. Empirically, we find that they have the same sign of slope (positive, as clothing is found to be a luxury) for almost all values of z.

5.4 Interpretation of the estimated women's relative resource shares

In the previous sections we have discussed the evolution of women's average relative resource shares and of the distribution of women's relative resource shares over time. We find that these resource shares increased over time both within childless couples and couples with children, although the increase was larger (14 percentage points) among childless couples than among couples with children (11 percentage points). The results were also temporally different for these two types of women – in particular only women in childless couples saw large increases in average resource shares between 1978 and 1995.

In order to interpret these results further, we first turn to think about the role of characteristics and time in driving changes in resource shares. There is no pre-determined way in which the evolution of characteristics and the evolution of relative resource shares (and thus the resource share function itself) should be linked, but rather this is in empirical question.

We saw in section (4) that, apart from the share of income coming from benefits, the characteristics that we might expect to increase the relative resource shares of women, such as education and the wage ratio (which can be considered to improve bargaining power within the household), have increased in level over the period. Indeed, it is tempting to attribute the increase in relative resource shares to the increase in level of the characteristics (such as education) of women, which is what Lise and Seitz (2011) and Bargain et al. (2023) find. Importantly, the methodologies used in those two papers could not help but do so: they assumed a time-invariant link between characteristics and resource shares.

However, we find that there are several periods in the data we examine, where these same characteristics continued to increase, but relative resource shares remained constant or even fell. This implies that the effect of these characteristics on resource shares has changed over time.

Given these patterns, in order to be able to evaluate the role of characteristics and the role of time in determining the value of the relative resource shares, we need to distinguish between the effect of characteristics and the effect of changes in the relative resource share function on the relative resource shares. In other words, we have so far documented the value of an object which is $R_t(z_t)$; we now need to examine how much of the change in the relative resource share is due to changes in the function $R_t()$ and how much is due to changes in z_t, the argument of the function. [6]

In addition to considering the role of changes in average characteristics and time in determining

---

[6]The increase in resource shares over time has already been documented in Bargain, Donni and Hentati (2023), who find a 5pp increase over 1978 to 2007, using the same data but a different model. They explain the increase in resource shares by increased incomes and increased female education. However, their model does not non parametrically identify a time trend, nor does it allow for changing returns to characteristics in the resource share function. They do *parametrically* identify a time trend, estimate it, and find it to be statistically insignificant.Consequently, all changes in resource shares have to be pinned on changing demographics. In other words, they estimate resource shares under the restriction that the function $R_t(z_t) = R(z_t)$. Additionally, they identify scale economies, using annual price variation, but at the cost of not identifying time-trends. Our model, on the other hand, does not identify scale economies, but

Figure 8: Counterfactual relative resource share with 1978 characteristics: $R_t\left(\overline{z_{1978}}\right)$

(a) Childless couples

(b) Couples with children

average resource shares, we also want to understand whether an increase in average characteristics and average resource shares in the first part of the period is due to some women having better characteristics and driving up the average shares, or to all women receiving larger shares of household expenditures. We presented some evidence to speak to this in (5.3).

We will use several techniques to attribute changes in characteristics and changes in the return to those characteristics inside marriage or cohabitation relationship (through resource shares) to the evolution of the resource shares on average and in distribution. Once we have done this, we will provide further discussion of the results in (5.7).

5.5 Counterfactuals and Oaxaca Decomposition

Having documented the evolution of relative resource shares over the period 1978 to 2019, the next question is to understand what drives this evolution. We see that relative resource shares rose for women in childless couples more than for women in couples with children. During this period, average characteristics of women, as well as how those characteristics are valued, changed, largely due to societal changes. We would like to disentangle the part of the evolution of relative resource shares which is driven by the changing characteristics of women from the part that is driven by changes in the resource share function itself. Before looking at Oaxaca decompositions, we show figures to illustrate the counterfactuals graphically. Figure 8 gives the counterfactual relative resource share corresponding to the 2nd term in the right-hand side of equation (19) for childless women and for mothers.

For women without children (the left hand side panel), the shape of the counterfactual is very similar to that of the factual, rising for the first two decades and then falling and rising again. However, the fall and rise over the 2nd two decades do not add up to as much increase as in the

---

allows for time trend and it is compatible with any scale economies where $Q_h = G_h\left(\left\{\widetilde{Q}_h^j\right\}_j\right) = A_t\left(\Sigma_j \widetilde{Q}_{hj}^j\right)$ for some time-varying matrix A_t.

Figure 9: Factual and counterfactual relative resource shares

(a) Childless couples

(b) Couples with children
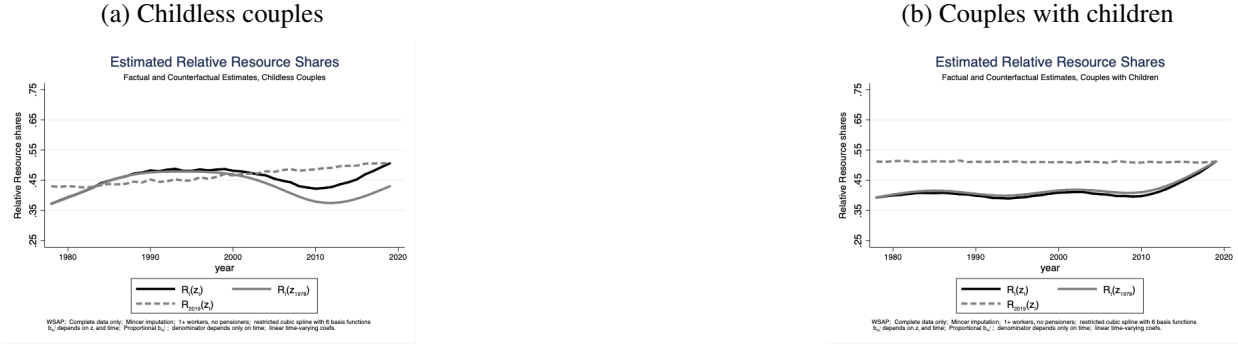
factual. The counterfactual relative resource shares (evaluated at z_{1978}) first increased by about 11 percentage points between 1978 and 1995. But, after 1995, the point-estimate of the counterfactual woman's resource share keeping characteristics fixed declined by 8 percentage points by 2010, and then increased again by 6 percentage points by 2019 (this increase is not statistically significant).

On the right-hand side of Figure 8, we see the estimated counterfactual women's relative resource share in couples with children. This estimate looks remarkably like the factual. If we take the point-estimates as revealing, it shows an increase of about 11 percentage points over the entire period. However, the increase in the counterfactual relative resource shares at the end of the period is not statistically significant.

It is worth noting that there is an increase in the point estimate of the counterfactual resource share function, both for mothers and childless women towards the end of the period (2010 to 2019), as we saw with the overall resource shares as well. Although the confidence intervals are wide and do not rule out other movements of the function, the fact that the point estimate is increasing for both types of women again indicates that the more recent increase in women's resource share function seems to be a real feature of the data.

The importance of the changing resource share function itself in driving changes in the resource share over time is particularly striking given that other empirical approaches for estimating resource shares rule out changes over time in this function. Although it is not possible with the data at hand and the tools of this approach to detail the drivers of change of the function, we can speculate as to what brought this about. One interpretation is that society is reevaluating the sharing inside households, for given characteristics. Another possibility is changes in assortative matching, as men and women being more closely matched in terms of their characteristics may change sharing inside household. Another possibility is that the resource share function is changing as a result of price changes and we cannot separate the effect of price changes from the effect of time. We further discuss possible explanations in (5.7).

Figure 9 shows these same results but comparing the factual resource share with the two coun-

terfactuals in one figure. Comparing the thick gray and thick black lines gives the first parenthetical expression in the Oaxaca decomposition (19) of the change in the factual from 1978 to year t. The difference between them gives the amount of change driven by the change in the relative resource share function. On the left-hand side, we see that for childless couples, before 1995, women's resource share was increasing mostly driven by the increasing resource share function. Between 1995 and 2010 the resource share function pushed the resource share down while characteristics were improving. Since 2010 both the resource share function and the changing characteristics have again increased women's resource share. On the right-hand side, we see that for couples with children, the resource share function did all the work over the whole period. The (statistically insignificant) increase in relative resource shares of mothers has taken place since 2010, driven by changes to the resource share function.

In addition to seeing the relative role of changing characteristics and the resource share function graphically, we can also express these changes in a table. Table 1 gives an Oaxaca decomposition corresponding to Figure 9, for the whole period, and for 3 subperiods of interest following the decomposition given by equation (18).

Table 1: Oaxaca Decompositions

type	time	from	to	$R_t(z_t)$	(se)	z_t	(se)	R_t	(se)
childless	all	1978	2019	0.134	0.083	0.076	0.047	0.058	0.091
couples	early	1978	1995	0.109	0.030	0.002	0.005	0.106	0.031
	mid	1995	2007	-0.038	0.036	0.027	0.010	-0.065	0.037
	late	2007	2019	0.063	0.093	0.020	0.012	0.044	0.093
couples	all	1978	2019	0.119	0.103	0.000	0.067	0.119	0.124
with	early	1978	1995	-0.001	0.041	-0.008	0.007	0.008	0.042
children	mid	1995	2007	0.006	0.043	-0.005	0.013	0.011	0.045
	late	2007	2019	0.114	0.113	-0.001	0.023	0.115	0.115

Because the standard error for the factual relative resource share is large in 2019 (for reasons discussed above), the Oaxaca decompositions for the entire period are not statistically significant. However, Figure 9 suggests that relative resource shares rose for childless couples over 1978 to 1995, were flat or declining for the next 12 years, and then possibly rose from 2007 to 2019. The Oaxaca decomposition for these subperiods shows that the increase of 10.9 percentage points over the early period was statistically significant and was driven essentially entirely by the evolution of the resource share function and not by changes in characteristics. Over the next two subperiods, the Oaxaca decomposition is not very precise, but does suggest an increase in resource shares in the late period, again driven mostly by the changing resource share. Turning to couples with children, we see that nearly the entire change of 11.9 percentage points occurred in the late period, almost all of it explained by the changing resource share function (although the increase in that period is

not statistically significant).

While characteristics for women in both types of households improve over the period, the resource shares evolved differently. Childless women's resource shares increased by more than the resource shares of mothers overall (13.4 vs 11.9 percentage points), and changes in characteristics and the resource share function each contributed about a half to the total change in resource shares for childless women, whereas for couples with children about all of the increase was driven by the resource share function.

Temporal patterns were also particularly different for these two types of couples. The biggest differences in how the resource shares evolved can be seen when we split the time period into an early period from 1978 to 1995, a middle period from 1995 to 2010, and late period from 2010 to 2019. In the early period, resource shares of childless women increased by 10.9 percentage points, mainly driven by the resource share function itself changing, whereas the resource share of women was constant. In the middle period, for childless women, improving characteristics increased resource shares by 2.7 percentage points, but the resource share function pushed the resource shares down so that there was an overall reduction of 3.4 percentage points over this period. In this period, the relative resource share for mothers was again unchanged. Finally, between 2007 and 2019 both childless women and women in couples with children saw large (although imprecisely estimated) increases in their resource share functions of 6.3 percentage points and 11.4 percentage points, respectively, mostly driven by the resource share function.

5.6   Contribution of individual characteristics

So far we have looked at the impact of all characteristics and their contribution to the resource share function at the same time. In this section we provide the estimate marginal effect on resource shares of the relative wages of women and men. We focus on this characteristic because we have a strong prior that it should have a positive effect on the relative resource share. We provide analysis for other characteristics in the online Appendix. We additionally investigate whether using the women's relative income share (rather than the relative wage), which is measured directly for all households (rather than imputed via a Mincer regression for nonworkers) in online Appendix (16).

The effect we find in Figure 10 is indeed consistent with our priors: we find a positive marginal effect of the wage ratio on women's relative resource shares (in terms of the point estimates).

In online Appendix (15), we provide similar figures for selected other characteristics. We find large positive effects of men's education on the relative resource share but essentially no effect of women's education on the relative resource share. We provide an interpretation of this phenomenon as resulting from matching in the marriage market in online Appendix (15). We also find large effects of benefits receipt on relative resource shares: amongst childless couples, those

Figure 10: Marginal effect of the wage ratio on women's relative resource shares

(a) Childless couples	(b) Couples with children

receiving benefits have lower relative resource shares.

5.7 Further discussion: Characteristics and cultural change in the evolution of women's relative resource shares

In previous sections, we reported on relative resource shares and their evolution and used Oaxaca decomposition to shed light on a number of patterns.

How should we then interpret the fact that the influence of characteristics on the resource shares and the resource shares function are different for both types of women, most prominently in the first period in the data (from 1978 to 1995)?

The resource share is a function of characteristics. In efficient collective household models (be they full, limited or no-commitment), level and dependence on characteristics of the resource share function is driven by the outside options of each person. Outside options are described by the value function of being single, which is defined as the expected lifetime utility of being single, accounting for the fact that single people can marry. Resource shares depend on characteristics because men's and women's value functions depend on characteristics.

We interpret the increase in the woman's resource share function for childless couples in the early period of the data (shown in Figure 6), as being driven by an increase in the value of being a single woman that occurred over the 1980s and 1990s. There are several reasons this may have occurred. First, the increase in reproductive freedom of women was a sustained process and seemed to plateau after 1990. Access to abortion in the UK was codified in law with the Abortion Act of 1967, and then extended with the Human Fertilisation and Embryology Act of 1990. There have been no further extensions of abortion access in the UK (outside Northern Ireland) since 1990 (Lowe and Page, 2022). Second, the gender wage gap in the UK declined very rapidly over 1971 to 2001, from 37% to 20%, but then much more slowly over the next 20 years, to 14% in 2018. Similarly, the full-time employment rate of women rose from 29% to 38% between 1971 and 2001,

but then much more slowly afterwards, rising to 40% by 2011 (Bryson et al 2020). Furthermore, the fraction of the population that was single rose very rapidly between the 1960s and 1990s, but rose more slowly thereafter. All these suggest that the value of being a childless single woman rose rapidly until the turn of the millennium, and rose more slowly thereafter.

In contrast to what we see for women in childless couples, the resource share function in couples with children did not rise in the early period we study. We think that the key difference between these groups is driven by the fact that generally, upon dissolution of marriage, children follow women. This means that the relevant value function for a woman in a couple with children is the value function for a single mother. Although we do believe that the value function for being a single mother rose steadily over this period, we do not believe it rose as much as the value function of being a childless single woman. If the value function for being a single mother didn't rise much–that is, if the expected lifetime utility of a single mother (including the utility gained from the option value of future marriage) did not increase much—then the resource share of women in couples with children would not be driven upwards like that of women in childless couples in the 1980s and 1990s.

6   Results: Gender inequality in expenditure

Relative resource shares provide a way to construct a measure of expenditure at the individual level from the observation of household expenditure: individual expenditure is equal to person's relative resource share times adult expenditure. For childless couples, adult expenditure equals household expenditure, so this is easily constructed. For households with children, we must bring in auxiliary information about children's consumption so that we can net out children's expenditure from household expenditure to recover adult expenditure.

We use OECD equivalence scales to estimate children's expenditures to recover adult expenditures as follows. Letting nc be the number of children and na be the number of adults, this equivalence scale assigns the share

$$\eta_h^c = \frac{nc * 0.3}{1 + (na - 1) * 0.5 + nc * 0.3}$$

to children in the household.

To construct the shadow budget of the woman in household type h using the relative resource share R_h, we multiply the household budget less children's expenditure by R_h:

$$\widetilde{y}_h^f = R_h (1 - \eta_h^c) y$$

$$\widetilde{y}_h^m = (1 - R_h)(1 - \eta_h^c) y$$

Individual expenditure inequality based the distribution of $\left\{\widetilde{y}_h^f, \widetilde{y}_h^m\right\}$ results from variation in household budgets y and from variation in relative resource shares R_h. By contrast, standard measures of inequality assume relative resource shares equal 0.5 for all households, and so are affected only by variation in household budgets.

In this section we use the estimated relative resource shares to measure individual income inequality. We start by looking at to what extent using the estimated relative resource shares – rather than assuming perfect income sharing – affects standard measures of individual inequality in total net income.

Analysis of inequality on the basis of shadow budgets leaves out the fact that people in large households take advantage of scale economies. To deal with this, we need to equivalise the incomes, that is, adjust the household income measures to account for differences in shadow prices that result from differences in household size and structure. We do this by applying the OECD equivalence scales to adjust household incomes before applying the resource shares to calculate individual shares of income. This equivalence scale is defined by $E_h = 1 + 0.5(na - 1) + 0.3nc$. So, in our analysis of inequality, we construct datasets with individual equivalent-incomes equal to $\widetilde{y}_h^j n_h / E_h$, where n_h is the number of household members.[7] Our estimates of inequality are of the population of individual adults. We do not construct measures of individual expenditure for children.

Trends in income inequality in the UK over time have been well-document. In particular, Belfield et al (2017) show that amongst households with workers in the UK, market income inequality rose quite substantially between the 1990s and 2010s, but net income inequality did not. They interpret this as revealing that the government undid the rise in market income inequality through the tax and transfer system. They calculate this with per capita income measures.

We show the 80-20 ratio for individual adults for: 1) scaled household market income divided by the number of household members (per capita measure of market income inequality, thick gray line); 2) scaled household net income divided by the number of household members (per capita measure of net income inequality, thick black line); and 3) scaled household net income allocated by estimated resource shares (individual net income inequality, thin black line). The thick lines in this figure show similar patterns to Belfield et al (2017). Between the mid 1990s and the mid 2010s (the period which Belfield et al (2017) focus on, 1995 to 2015), the 80-20 ratio of market income inequality increased by about 0.5. At the same time, for household net income we see a much smaller increase in inequality: the 80-20 ratio increased by about 0.1 In other words, we also find that government redistribution undid a lot of the the increase in market income inequality over

---

[7]We multiply by n_h because equivalence scales are usually applied to total household income, not a person's share of household income.

Figure 11: The 80/20 ratios of household market and net total income over time
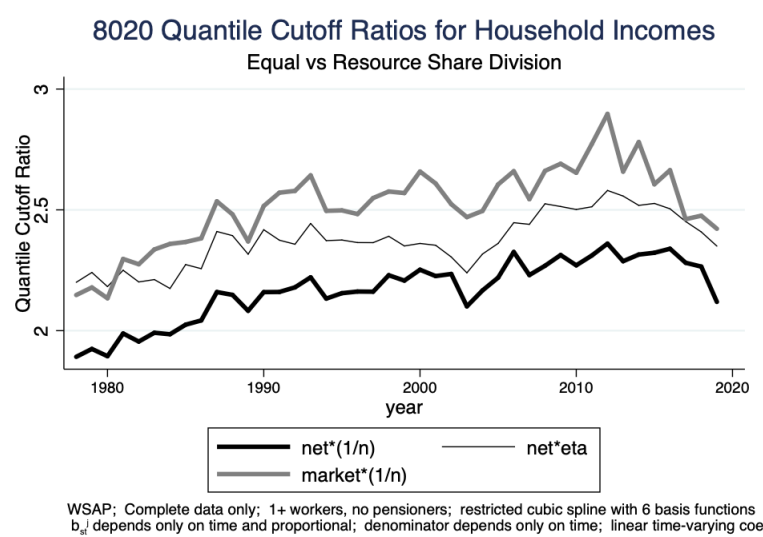

those two decades.

The thick lines in Figure 11 refer to traditional per capita measures of income inequality. The thin black line shows how changes to resource shares over this period affected income distribution, as economic changes *within* households can also affect the distribution of individual material well-being. Since our resource share estimates vary across households, they can induce inequality in individual-level income measures even for households with the same level of income. Consequently, the thin black lines (which allow for gender inequality) are above the thick black lines (which assume equal division), showing more inequality in every year.

We can also see that our individual-level income measures show a different pattern of the evolution of inequality over time. The 80-20 ratio using per capita measures of individual net income increased from about 1.9 in the late 1970s to about 2.2 in the early 2000s, an increase of about 0.3. On the other hand, the 80-20 ratio using relative resource shares to calculate individual income rose from about 2.2 in the late 1970s to about 2.4 in the early 2000s, an increase of about 0.2. In other words, before 2000, unequal-division net income inequality rose more slowly than equal-division net income inequality because resource shares were becoming *more equal*, as changes in sharing within households undid about a third of the increase in household level net income inequality. But after 2000 (or so), resource shares were flatter, and unequal-division inequality followed a similar path to equal-division inequality.

Figure 7 showed significant variation in resource shares across households in any given year. Given that these resource shares are point estimates given the observed covariates for each household, they contain estimation error, which might add spurious variation to our inequality estimates. One way to see if our main result is driven by this is to zero out all the variation in resource shares

Figure 12: The 80/20 ratio of household total net income using actual and average resource shares
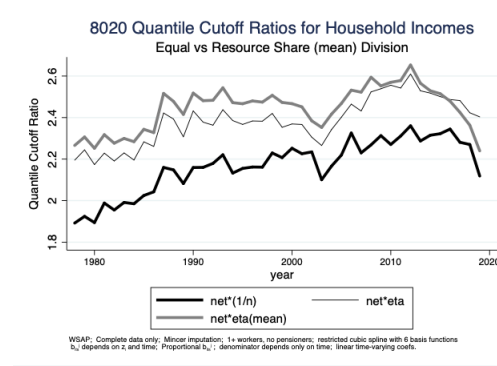
within year so that we can focus on the aggregate-level change in resource shares over time. Figure 12 gives an assessment of this type. Here, we show the thick black and thin black lines from Figure 11, replacing the thick gray line with the 80-20 ratio that would obtain if we assign the mean resource share in each year shown in figure (factual) to every collective household.

First, we can see that the thick gray line is higher than the thin gray line, meaning that some of the variation in resource shares within year was overall equalizing. Focusing on the time trend, the thick gray line shows the 80-20 ratio for net income evaluated at the person level and assuming share according to Figure 6 for every household in the relevant year. Here, we see an increase in the 80-20 ratio from about 2.3 to 2.5 between the late 1970s and the early 2000s, an increase of about 0.2. This is the same increase as we observe in the thin black line in the previous figure, where resource shares vary across households in a given year. This means that the pattern we observe in the first two decades—declining inequality within households canceling out rising inequality across households—is driven by the time trend of average resource shares shown in Figure 6.

So far in this section we have looked at all households together. We can also look at couples without children and with children to assess which types of households are driving the equalizing trend. Figure 13 shows the 80-20 ratios by household type, for childless couples on the left and couples with children on the right. These are analogous to figure 11, showing the per capita measures of market and total net income inequality, and the measure of unequal-division net income inequality. Recall that the difference in the time trend between the thick black line and thin black line describes the amount of the increase in net income inequality that is undone by changes in resource shares within households. For childless couples, we see the thick black line rising from an 80-20 ratio of about 1.8 in the late 1970s to about 2.2 in the early 2000s. In contrast, the thin black line is overall constant over this period at about 2.2 (with a small decline and small increase in the middle). This means that for childless couples, the increase in inequality in the last decades of the 20th century was completely undone by the equalization of resource shares within households.

For households with children, we do not see this pattern. This should not be surprising given

Figure 13: The 80/20 ratios of net household income for couples without and with children

(a) Childless couples

(b) Couples with children

that resource shares for women in these households were static over the first half of the period. In particular, we see the 80-20 ratio for net income using per-capita allocation rising from about 1.8 to 2.1, an increase of about 0.3. When we use resource shares to estimate person-level inequality (the thin black line), the 80-20 ratio rises from about 2.0 to about 2.3, a similar increase.

7 Conclusion

Measuring resource shares within households is important for measuring gender inequality, as well as welfare impacts of policies on individuals. In this paper we extend the theoretical framework for intra-household allocations to interpret differences in marginal responses to income in household Engel curves for private assignable goods (in our case clothing) as resource shares. Using a simple linear methodology, we measure how resources are shared between men and women within British households from 1978 to 2019. In particular, we estimate relative resource shares – the fraction of household resources that adult women command – and document how these evolve over four decades. We also use these estimated relative resource shares to estimate the evolution of individual-level consumption inequality over time.

We introduce several theoretical contributions to the literature on estimating resource shares using collective household models and private assignable goods. First, rather than trying to estimate all resource shares within a household (including children's), we introduce the concept of relative resource shares, which is defined as the fraction of total adult expenditure that women command. This notion of a relative resource means we have fewer parameters to estimate, which in turn allows us to use weaker identifying assumptions than previous literature. We call this a 'weakened' assumption for similarity of preferences across people. Where many previous models assume that men and women have similar preferences towards an assignable good (such as clothing), we relax this assumption so that we only require that men's preferences are similar whether they are in a couple or not, and the same applies for women.

41

Importantly, our approach also allows for the identification of relative resource shares from repeated cross sectional data with time varying resource share functions, which can be identified in each time period. We can thus assess whether changes in women's relative resource shares were driven by changes in their characteristics, or by how those characteristics were valued within households.

We also make several empirical contributions. Using our simple linear estimation methodology, we find that over the period of 1978 to 2019 the resource shares of childless women rose by up to 13 percentage points, while those of women with children rose by less, by about 12 percentage points. The timing of these gains differed by household type – among couples without children, women's shares rose over the 1980s and 1990s, whereas the change for mothers in couples was less pronounced until roughly 2010. This suggests that the outside options of single women, and thus their bargaining power in the relationship, shifted especially over the 1980s and 1990s, while this was not the case for mothers. Overall, the trend in women's relative resource shares was increasing and equalising over this period, rising from disparity to roughly parity with men.

We show that improvements in women's observed characteristics (e.g. higher wage ratios and more education) can explain about half of the increase in the relative resource shares for women without children, but the changing social and policy landscape – captured by changes in the resource share function itself – also played a major role. For couples with children, the increase was driven entirely by the resource share function.

Our estimated relative resource shares are used to estimate the evolution of individual-level consumption inequality over time. Knowing resource shares, we create measures of individual inequality that relax the usual assumption of equal sharing of resources within a household. We find that especially during the 1980s, greater equality within households offset a substantial fraction of rising inequality between households. Traditional household-level measures of inequality that assume equal sharing within households mask these kinds of differences in individual outcomes.

These contributions provide a framework and a tool for measuring allocation of resources within a household. By focusing specifically on relative resource shares under weaker identifying assumptions, and by allowing the resource share function to evolve with time, we significantly extend existing methods for measuring resource shares within households. Our empirical findings emphasise the importance of measuring resource shares within households and how they evolve over time, in order to better understand both changes in gender inequality and broader income inequality trends.

8 References

Akhal, A., 2019. The impact of recent government policies on early years provision. Education Policy Institute. https://epi. org. uk/wp-content/uploads/2019/01/Early-years-take-up_-EPI. pdf.

Angrist, Joshua D., and Jörn-Steffen Pischke. Mastering'metrics: The path from cause to effect. Princeton university press, 2014.

Apps, Patricia F., and Ray Rees. "Taxation and the Household." Journal of Public Economics 35, no. 3 (1988): 355-369.

Apps, Patricia, and Elizabeth Savage. "Labour supply, welfare rankings and the measurement of inequality." Journal of Public Economics 39, no. 3 (1989): 335-364.

Arduini, F., 2023. Estimating intra-household sharing from time-use data. lnstitute for Fiscal Studies Working paper, 23, p.23.

Bargain, Olivier, Olivier Donni, and Imen Hentati. "Resource sharing in households with children: A generalized model and empirical evidence from the UK." Journal of the European Economic Association 20, no. 6 (2022): 2468-2496.

Bargain, Olivier, Olivier Donni, and Prudence Kwenda. "Intrahousehold distribution and poverty: Evidence from Cote d'Ivoire." Journal of Development Economics 107 (2014): 262-276.

Bargain, Olivier, Guy Lacroix, and Luca Tiberti. "Validating the collective model of household consumption using direct evidence on sharing." Partnership for economic policy working paper 2018-06 (2018).

Becker, Gary S. "A Theory of the Allocation of Time." The economic journal 75, no. 299 (1965): 493-517.

Becker, Gary S. "Altruism in the Family and Selfishness in the Market Place." Economica 48, no. 189 (1981): 1-15.

Belfield, C., Blundell, R., Cribb, J., Hood, A. and Joyce, R., 2017. Two decades of income inequality in Britain: the role of wages, household earnings and redistribution. Economica, 84(334), pp.157-179.

Blackorby, Charles, and David Donaldson. "Adult-equivalence scales and the economic implementation of interpersonal comparisons of well-being." Social Choice and Welfare 10, no. 4 (1993): 335-361.

Blundell, Richard, Alan Duncan, and Krishna Pendakur. "Semiparametric estimation and consumer demand." Journal of applied econometrics 13, no. 5 (1998): 435-461.

Bourguignon, François, and Pierre-André Chiappori. "Collective models of household behavior." Developments in Agricultural Economics 10, no. C (1994): 205-216.

Brown, Caitlin, Rossella Calvi, and Jacob Penglase. "Sharing the pie: An analysis of undernutrition and individual consumption in Bangladesh." Journal of Public Economics 200 (2021):

104460.

Browning, Martin, and Pierre-André Chiappori. "Efficient intra-household allocations: A general characterization and empirical tests." Econometrica (1998): 1241-1278.

Browning, Martin, Francois Bourguignon, Pierre-André Chiappori, and Valerie Lechene. "Income and outcomes: A structural model of intrahousehold allocation." Journal of political Economy 102, no. 6 (1994): 1067-1096.

Browning, Martin, Pierre-Andre Chiappori, and Arthur Lewbel. "Estimating consumption economies of scale, adult equivalence scales, and household bargaining power." Review of Economic Studies 80, no. 4 (2013): 1267-1303.

Bryson, A., Joshi, H., Wielgoszewska, B. and Wilkinson, D., 2020. A short history of the gender wage gap in Britain. Oxford Review of Economic Policy, 36(4), pp.836-854.

Calvi, Rossella. "Why are older women missing in India? The age profile of bargaining power and poverty." Journal of Political Economy 128, no. 7 (2020): 2453-2501.

Cherchye, Laurens, Bram De Rock, and Frederic Vermeulen. "Opening the black box of intrahousehold decision making: Theory and nonparametric empirical tests of general collective consumption models." Journal of Political Economy 117, no. 6 (2009): 1074-1104.

Chiappori, Pierre-André. "Rational household labor supply." Econometrica: Journal of the Econometric Society (1988): 63-90.

Chiappori, Pierre-André. "Collective labor supply and welfare." Journal of political Economy 100, no. 3 (1992): 437-467.

Chiappori, Pierre-André, Richard W. Blundell, and Costas Meghir. Collective labour supply with children. No. 02/08. IFS working papers, 2002.

Chiappori, P-A., and Ivar Ekeland. "The microeconomics of efficient group behavior: Identification 1." Econometrica 77, no. 3 (2009): 763-799.

Chiappori, Pierre-Andre, and Maurizio Mazzocco. "Static and intertemporal household decisions." Journal of Economic Literature 55, no. 3 (2017): 985-1045.

Cherchye, Laurens, Bram De Rock, and Frederic Vermeulen. "The collective model of household consumption: a nonparametric characterization." Econometrica 75, no. 2 (2007): 553-574.

De Vreyer, Philippe, and Sylvie Lambert. "Inequality, poverty, and the intra-household allocation of consumption in Senegal." The World Bank Economic Review 35, no. 2 (2021): 414-435.

Deaton, Angus, and John Muellbauer. Economics and consumer behavior. Cambridge university press, 1980.

Donaldson, David, and Krishna Pendakur. "Equivalent-expenditure functions and expenditure-dependent equivalence scales." Journal of Public Economics 88, no. 1-2 (2004): 175-208.

Donaldson, David, and Krishna Pendakur. "The identification of fixed costs from consumer behavior." Journal of Business & Economic Statistics 24, no. 3 (2006): 255-265.

Dunbar, Geoffrey R., Arthur Lewbel, and Krishna Pendakur. "Children's resources in collective households: identification, estimation, and an application to child poverty in Malawi." American Economic Review 103, no. 1 (2013): 438-471.

Jorgenson, Dale W., and Daniel T. Slesnick. "Aggregate consumer behavior and household equivalence scales." Journal of Business & Economic Statistics 5, no. 2 (1987): 219-232.

Lechene, Valérie, Krishna Pendakur, and Alex Wolf. "Ordinary least squares estimation of the intrahousehold distribution of expenditure." Journal of Political Economy 130, no. 3 (2022): 681-731.

Lewbel, Arthur, and Krishna Pendakur. "Inefficient collective households: Cooperation and consumption." The Economic Journal 132, no. 645 (2022): 1882-1893.

Lise, Jeremy, and Shannon Seitz. "Consumption inequality and intra-household allocations." The Review of Economic Studies 78, no. 1 (2011): 328-355.

Lewis, J., 2013. The failure to expand childcare provision and to develop a comprehensive childcare policy in Britain during the 1960s and 1970s. Twentieth Century British History, 24(2), pp.249-274.

Lowe, P. and Page, S.J., 2022. Opposition to Abortion: A UK History. In Anti-Abortion Activism in the UK: Ultra-sacrificial Motherhood, Religion and Reproductive Rights in the Public Sphere (pp. 19-36). Emerald Publishing Limited.

McElroy, Marjorie B., and Mary Jean Horney. "Nash-bargained household decisions: Toward a generalization of the theory of demand." International economic review (1981): 333-349.

Menon, M., Pendakur, R. and Perali, F., 2015. All in the family: How do social capital and material wellbeing affect relational wellbeing?. Social Indicators Research, 124, pp.889-910.

Muellbauer, John. "The cost of living and taste and quality change." Journal of Economic Theory 10, no. 3 (1975): 269-283.

Oaxaca, Ronald. "Male-female wage differentials in urban labor markets." International economic review (1973): 693-709.

Pendakur, Krishna. "Semiparametric estimates and tests of base-independent equivalence scales." Journal of Econometrics 88, no. 1 (1999): 1-40.

Penglase, Jacob. "Consumption inequality among children: Evidence from child fostering in malawi." The Economic Journal 131, no. 634 (2021): 1000-1025.

Sokullu, S. and Valente, C., 2022. Individual consumption in collective households: Identification using repeated observations with an application to Progresa. Journal of Applied Econometrics, 37(2), pp.286-304.

9 Online Appendix: Notations

People types are adult males, m, adult females, f, or children, c (this means that children can potentially have utility functions). We use the subscript h to index the type of household an individual j might live in. Household types h are

- single, s

- households with: more than one household member, in which all the adults are female and any number of children, including none, fx

- households with: more than one household member, in which all the adults are male and any number of children, including none, mx

- childless male-female couples, mf

- male-female couples with children, mfc

- other household types with at least 1 man and 1 woman, mfx

If $h = s$, then this indicates a household where a person of type j lives alone. Only in the last 3 types of households might there be gender inequality in resource shares amongst the adults.

In a household of an adult couple with a child, the people might be $\{male, female\}$, or they might be $\{male, female, child\}$, depending on whether children are taken to be people (with utility functions) or taken to be attributes of the household. Households may have many attributes, denoted z, which we will explicitly pay attention to when we write the empirical model, but we will suppress z throughout the discussion of theory and identification. In our new work, children can be either attributes of the household *or* members with utility functions. For this part of the exposition, we will index the people in the household with $j = m, f$.

Resource shares $\eta_h^j(p,y)$ are in general functions of prices p and budgets y. Thus $\eta_{mfc}^f(p,y)$ gives the resource share of an adult female in a couple with children facing market prices p and a household budget y.

Let q_h be the observed quantity vector purchased at market prices by the household. Let $\widetilde{q}_h^j$ be the shadow demand of person j given the preferences they have if they live in household of type h. We may also write this as a function of price and budget arguments, $\widetilde{q}_h^j\left(\widetilde{p}_h^j, \widetilde{y}_h^j\right)$. Note that demand functions for singles, $\widetilde{q}_s^j\left(\widetilde{p}^j, \widetilde{y}^j\right)$ may not be a shadow quantity because it may be observable. We retain the tilda notation for $\widetilde{q}_s^j$ to keep things simpler.

Notation	Definition	Dimension
P	Market price of any good	vector
p	Market price of private assignable good of individual of type j	scalar
y	Household budget	scalar
Q	Household quantity demanded of any good, from centralised programme $O1$	vector
q^j	Household quantity demanded of private assignable good of individual of type j, from centralised programme $O1$	scalar
$F()$	Household demand function for any good, from centralised programme $O1$	vector
$f()$	Household demand function for private assignable good of individual of type j, from centralised programme $O1$	scalar
W_h^j	Budget share and Engel curve for any good, from centralised programme $O1$	vector
w_h^j	Budget share and Engel curve for a private assignable good of individual of type j, from centralised programme $O1$	scalar
$\widetilde{P}_h^j$	Shadow price of any good of individual of type j living in a household of type h	vector
$\widetilde{p}_h^j$	Shadow price of private assignable good of individual of type j	scalar
$\widetilde{y}_h^j$	Shadow budget of individual of type j when they live in a household of type h	scalar
$\widetilde{Q}_h^j$	Individual quantity demanded of any good, from decentralised programme $O2$, of individual of type j living in a household of type h	vector
$\widetilde{F}^j()$	Shadow demand function of individual of type j for any good, from decentralised programme $O2$	vector
$\widetilde{f}_h^j()$	Shadow demand function of individual of type j when they live in a household of type h, for a private assignable good of individual of type j, from decentralised programme $O2$	scalar
$\widetilde{w}_h^j$	Shadow budget share and Engel curve of individual of type j when they live in a household of type h, for a private assignable good of individual of type j, from decentralised programme $O2$.	scalar
G_h	Household consumption technology function	vector
η_h^j	Resource share of individual of type j when they live in a household of type h. With abuse of notation, can indicate either the level or the function.	scalar
R_h	Relative resource share of individual of type j when they live in a household of type h	

10 Online Appendix: Identification of Resource Shares: related strategies

The identification strategy presented in this paper relates to those provided in three papers: Browning, Chiappori and Lewbel (2013: BCL), Dunbar, Lewbel and Pendakur (2013: DLP) and Lechene, Pendakur and Wolf (2022: LPW). All three papers present ways to identify and estimate resource shares. The approach we use in this paper is an example of BCL, and is closely related to DLP and LPW.

10.1 Identification in BCL

BCL show generic identification of resource shares and shadow prices in any model with the very general household consumption technology $Q_h = G_h(\sum_{j \varepsilon h} \widetilde{Q}_h^j)$ for any function G_h. BCL also introduce a simpler model for consumption technology, based on a Barten consumption technology, wherein

$$G_h(q) = Aq$$

where q is the $K-$vector argument of G_h and A is a diagonal matrix.

For the entire $K-$vector of quantities, $Q_h = [Q_{1h}, ..., Q_{Kh}]$, for a household of type h, we have

$$Q_h = A_h(\sum_{j \varepsilon h} \widetilde{Q}_h^j) \tag{20}$$

where A_h is a diagonal matrix. BCL show that this relationship implies the shadow price vector $\widetilde{P}_h^j$ is the same vector for all household members, $\widetilde{P}_h^j = \widetilde{P}_h = A_h P$.

The diagonal elements of the matrix A define the shareability of each good. If the element of A corresponding to a good equals 1, then the good is nonshareable. In this case, if each of N_h household members demanded q units of a good, the household would need to purchase $N_h q$ units in the market to satisfy their demands. If it is less than 1, then the good is shareable, and the household would need to purchase less than $N_h q$ units of the good to satisfy their demands. If it is $1/N_h$ then it is fully shareable in the sense that if each household member demanded a quantity q, the household could satisfy their demands by purchasing q units in the market.

Recalling the definition of the shadow budget, $\widetilde{y}^j = \eta_h^j(P, y)y$, we substitute the shadow prices and budgets into the above equation to get the vector of household-level demands as

$$Q_h = A_h \left(\sum_{j \varepsilon h} \widetilde{Q}_h^j \left(A_h P, \eta_h^j(P,y)y \right) \right) \tag{21}$$

Quantities consumed by singles are Q_s^j. We may observe these quantities for no types, for some types (typically, quantities of private assignable goods are observed for single women and single men, but not for single children, as children tend not to live alone) or for all types. When a person j lives alone (in a household of type s for single), their resource share is 1 and no goods are shareable, so that $A_s = diag(1_K)$. Thus, for the types of singles for whom we observe quantities, we have

$$Q_s^j = \widetilde{Q}_s^j$$

and, since, for all j, $Q_s^j = F_s^j(P,y)$ and $\widetilde{Q}_s^j = \widetilde{F}_s^j\left(\widetilde{P}, \widetilde{y}_s^j\right)$ and for singles, $\widetilde{P} = P$ and $\widetilde{y}_s^j = y$, we have

$$Q_s^j = \widetilde{F}_s^j(P,y)$$

BCL show that resource shares η_h^j and scale economies A_h are identified from the $K-$vector of demands functions for all goods of all single-person households and collective households if: a) we assume the Barten model of BCL (described above); b) we assume that people's preferences are the same whether they live alone or live in households so that F_s^j and $\widetilde{F}_h^j$ are the same functions (F_h^j does not vary with h; and, c) we have sufficient market price variation for all goods.

This is the version of the model that they bring to the data. It is quite difficult to implement; only 3 papers have implemented this model (BCL; Pendakur 2017; and, Lewbel and Lin 2022). If $j = m, f$ so that the model has only adult males and adult females, one can plausibly obtain data on quantity demands for single men and women in some contexts. However, in some countries, single adults are not a common household type. Further, j includes children, observing the demands of single children presents an obstacle.

If shadow quantities were observed, they would not be shadow, but would rather be observed demands at the person-level within the household. Bargain et al (2020) make some progress in this kind of data environment. They have person-level consumption of all goods, and resource shares can therefore be observed directly in the data as the fraction of observed household consumption enjoyed by each person. For us, the identification challenge is to identify resource shares in the absence of person-level data on consumption of all goods.

10.2 Identification in DLP

DLP show identification of the model without data on singles and with children as welfare-relevant people. They use different restrictions on preferences and use the identifying power of a private

assignable good[8] to achieve identification of just the resource shares η_h^j (see also Chiappori and Ekeland 2009). A private good is not shareable, and so its element of A equals 1. DLP assume that there is one assignable good, the same good for each person, observed in the data. In their empirical work, as in ours, this good is clothing.

Let q_h^j be the observed scalar-valued quantity of the private assignable good purchased by the household for person j, and let $\widetilde{f}_h^j$ be the element of $\widetilde{F}_h^j$ corresponding to the private assignable good. Plugging into equation (20) gives

$$q_h^j = \widetilde{f}_h^j \left(A_h P, \eta_h^j (P,y) y \right). \tag{22}$$

The A_h premultiplying the parenthetical expression in (21) disappears because the private assignable good for person j is not shareable, so its element of A_h equals 1. The summation inside the parenthetical expression disappears because only person j wants to buy the private assignable good for person j.

Equation (22) can also be expressed in Engel curve form, where the *Engel curve* for a private assignable good relates the fraction of household expenditure commanded by that good to household income: $w_h^j = p q_h^j / y$, where p is the element of P corresponding to the private assignable good. The Engel curve for a household may be expressed in terms of the shadow Engel curve functions of individuals, $\widetilde{w}_h^j(p,y)$, which give the fraction of expenditure that would be commanded by the private assignable good if person j lived in household h and faced a budget constraint p,y. The above quantity demands may be expressed as household Engel curves for the assignable good of person j in household h, W_h^j:

$$W_h^j = \eta_h^j (p,y) \widetilde{w}_h^j \left(A_h p, \eta_h^j (p,y) y \right). \tag{23}$$

Preferences $\widetilde{w}_h^j$ depend on the household type h in which one lives. This equation says that the fraction of expenditure commanded by, e.g., men's clothing in a household of type h is equal to the men's resource share multiplied by the fraction of expenditure that would have been commanded by clothing by a man whose preferences are those of a man in household type h, facing prices $A_h P$ and with a budget of $\eta_h^j y$.

DLP show (in Appendix Theorems 1 and 2) that the resource share function is identified from private assignable goods demands of collective households if: a) we assume the Barten model of BCL; b) the resource share function does not depend on household expenditure y; and, c) the

---

[8]We defined assignable goods and private assignable goods in (2.2). Assignable goods are goods for which individual consumption is observed. For a private assignable good of individual j, the quantity demanded by the household, q^j is equal to the shadow quantity demanded by the individual, $\widetilde{q}^j$. We use lower case letters to denote private assignable goods as they have different properties from other goods. For all goods, including private assignable goods, we use upper case letters.

shadow Engel curves of individual satisfy one of two semiparametric preference restrictions, called "similar across people" (SAP) or "similar across types" (SAT).

The SAP restriction of DLP is that shadow Engel curve functions of individuals for the assignable good, $\widetilde{w}_h^j(P, y)$, satisfy the following semiparametric shape-invariance restriction:

$$\widetilde{w}_h^j(P, y) = d_h^j(P) + g_h\left(\frac{y}{G_h^j(P)}, P\right). \tag{24}$$

In their appendix, they show the class of utility functions that satisfy this restriction, which includes the semiparametric families of utility functions described in Lewbel (1989) (base-independent utilities), Blackorby and Donaldson (1993) (equivalence-scale exact utilities) and Pendakur (1999) (shape-invariant utilities). They call it a *semiparametric* restriction because the functions of prices, d_h^j and G_h^j, are unknown and unrestricted.

There have been a few applications of this methodology (Bargain, Donni and Kwenda 2014 in Cote d'Ivoire; De Vreyer and Lambert 2021 in Senegal; Bargain, Lacroix and Tiberti 2018 in Bangladesh; Brown, Calvi and Penglase 2019 in Bangladesh; Calvi 2019 in India; Penglase 2019 in Malawi), but the reliance on nonlinear econometrics is troublesome.

10.3   Identification in LPW

LPW offer a linear representation of the model of DLP, a strategy we also follow in this paper. They show that if we additionally assume that shadow Engel curves are PIGLOG (Muellbauer 1975), so that they are linear in the log of budgets, then the reduced form observed household Engel curves are linear, and so may be estimated by linear methods like OLS and 2SLS. Further, the resource share (an unobserved structural parameter of the model) may be recovered as a function of the reduced form coefficients.

Specifically, they show that resource shares are identified by reduced form coefficients in linear regression models of private assignable-goods budget shares of collective households if: a) we assume the Barten model of BCL; b) shadow Engel curves for the private assignable good are linear in the log-budget

$$\widetilde{w}_h^j(P, y) = \alpha_h^j(P) + \beta_h^j(P)\ln y \tag{25}$$

which satisfies similarity across people (SAP) if $\beta_h^j(P) = \beta_h(P)$; and c) resource shares do not depend on the budget so that $\eta_h^j(P, y) = \eta_h^j(P)$.

Together, these restrictions imply the following linear reduced form for household Engel curves at a fixed price regime P:

$$W_h^j = a_h^j + b_h^j \ln y_h$$

where

$$a_h^j = \eta_h^j \alpha_h^j + \eta_h^j \beta_h \ln \eta_h^j \qquad \text{and} \qquad b_h^j = \eta_h^j \beta_h.$$

Here, the structural parameters are functions of the fixed price vector P where, with some abuse of notation, $\eta_h^j = \eta_h^j(P)$, $\alpha_h^j = \alpha_h^j(AP)$ and $\beta_h = \beta_h(AP)$.

The structural parameters—the resource shares—are recovered as functions of the reduced form parameters:

$$\eta_h^j = \frac{b_h^j}{\sum_j b_h^j}.$$

The methodology of LPW relies on the SAP restriction which with PIGLOG Engel curves implies that the slope of men's Engel curves for the assignable good equals the slope of women's Engel curves for that good. For example, if the assignable good is clothing, and men's shadow Engel curves $\widetilde{w}_h^m(p,y)$ have that clothing is a luxury with a budget semi-elasticity (β_h^m) of 0.1, then women's shadow Engel curves $\widetilde{w}_h^f(p,y)$ must show the same budget semi-elasticity of 0.1. In the data we use, we observe the Engel curves of single men and single women and have direct evidence that the slopes of clothing Engel curves, β_s^m and β_s^f, are different for men and women who live alone. The identifying assumption we use is weaker than SAP in exactly this direction, specifically allowing men's and women's shadow Engel curves to have different slopes. However, like DLP, we provide a linear reduced form where functions of reduced-form coefficients identify the structural parameters of interest.

11 Online Appendix: Sketch of identification result in this paper: Gender inequality given weakened SAP (WSAP)

In the present paper, we provide an identification result that uses a weaker similarity restriction than SAP. Our identification result identifies less: instead of identifying the resource shares (defined by equation (3)) of men, women and children in the household, we identify the relative resource shares (defined by equation (4)) of women. That is, we identify gender inequality in resource shares, and do so with a weaker identifying preference restriction than DLP. However, we use more data than do DLP: in addition to data on household demands for assignable goods, we use data on singles' demands for those same goods.

Theorem 1 of Dunbar, Lewbel and Pendakur (2013) shows identification of resource shares η_h^j given SAP if Engel curves of all types of people in a particular household type h are observed. We show (in Appendix (12), corollary 1) that their theorem identifies the relative resource share given similarity across people (SAP) if just the Engel curves of men and women are observed.

In Theorem 1 (in Appendix (12)) we show that one can use a weaker preference restriction than

SAP of DLP, which we call weakened similarity across people (WSAP). We combine WSAP with the assumption that we observe the demands of singles to identify the relative resource share.

Let weakened similarity across people (WSAP or "weakened SAP") be satisfied if the shadow Engel curve functions of individuals for the assignable good, $\widetilde{w}_h^j(P,y)$, satisfy the following semi-parametric generalized shape-invariance restriction:

$$\widetilde{w}_h^j(P,y) = d_h^j(P) + \gamma_h \left(\left(\frac{y}{\Gamma_h^j(P)} \right)^{\delta^j}, P \right) \tag{26}$$

for some functions d_h^j, Γ_h^j and γ_h and scalars δ^j, with the normalization $\delta^m = 1$ (or $\delta^f = 1$). Note that SAP in equation (24) is satisfied here if $\delta^j = \delta$ doesn't vary across people. If δ^j is different for men and women, then WSAP is a weaker restriction than SAP. In Lemma 1 in Appendix (12), we show the class of utility functions that implies WSAP.

Our main new theoretical result is our Theorem 1 in Appendix (12). There, we show that the relative resource share of women in household type h, R_h, is identified from the assignable goods Engel curves of single men and single women and those of men and women in household type h if: a) the model of BCL holds; b) WSAP is satisfied; and c) resource shares don't depend on household expenditure.

The sketch of the proof for our Theorem 1 is as follows (details are in Appendix (12)). The budget share of singles for the private assignable good (e.g., for clothing), w_s^j, are observed for $j = m, f$ and are in the semiparametric class defined by WSAP:

$$w_s^j = \widetilde{f}_s^j(P,y) = d_s^j(P) + \gamma_s \left(\left(\frac{y}{\Gamma_s^j(P)} \right)^{\delta^j}, P \right).$$

These demand functions are a case of generalized equivalence-scale exactness described by Donaldson and Pendakur (2004). That paper covers unitary households only, but since single-member households are unitary, those results apply here. They show that the functions $\gamma_s(.,P)$, $d_s^j(P)$, $\Gamma_s^j(P)$ and the parameter δ^j are identified from the Engel curves of singles if preferences are not PIGLOG, and that δ^j is identified from the Engel curves of singles if preferences are PIGLOG but not homothetic. The identification of δ^j is up to the normalization that $\delta^j = 1$ for some j, and in our case we normalise $\delta^m = 1$.[9] Thus, from the behaviour of singles, we identify the parameter

---

[9]We can see the intuition for this result by considering singles' Engel curves in terms of logged budgets:

$$W_s^j = d_j^j(P) + \widetilde{\gamma}_s \left(\delta^j \left(\ln y - \ln \Gamma_j^j(P) \right), P \right)$$

where $\widetilde{\gamma}_s$ is γ_s defined over the logged version of its first argument. Here, Engel curves have a nonparametric shape

---

53

δ^f.

Then, given we have identified δ^f (and normalized $\delta^m = 1$), we can plug δ^j into (26) by replacing y with $\widehat{y} = y^{\delta^j}$ and $\Gamma_h^j(P)$ with $\widehat{\Gamma}_h^j(P) = \left(\Gamma_h^j(P)\right)^{\delta^j}$. With these replacements, these budget share functions over $P, \widehat{y}$ satisfy SAP. Finally, we recall that (from Corollary 1 of DLP Theorem 1), the relative resource share R_h is identified if: BCL holds; SAP holds; both men's and women's Engel assignable good curves are observed in household type h; and, resource shares do not depend on the budget y. Consequently, if: BCL holds; WSAP is satisfied; we observe single men's and single women's assignable goods demands and the assignable goods demands for men and women in households; and resource shares are independent of the household budget, then the relative resource share R_h is identified. Essentially, by observing singles, we observe δ^j. With δ^j in hand, identification from the Engel curves of collective households proceeds via DLP.

11.1 Remarks

We have four remarks about our identification result.

First, our identification theorem is semiparametric in that the functions d_s^j, γ_s and Γ_s^j are unknown and nonparametric functions. d_s^j and Γ_s^j are functions of prices P and γ_s is a function of prices P and the transformed budget y.

Second, because identification is shown for each price vector P, we can estimate R_h using cross-sectional data facing fixed prices (we could alternatively identify using data where every household faced a different price vector, but that data configuration is atypical in this literature). In data where the only observed price variation is over time (as in our setting), this means that we can separately identify the relative resource share in each time period/price regime. As a consequence, the relative resource share function can be arbitrarily flexible over time. This is in contrast to other identification strategies where identification is conditional on a fixed resource share function over all time periods/price regimes (e.g., Lise and Seitz (2011) and Bargain et al (2022)). On the other hand, if time and prices have no independent variation, then our methodology cannot separately identify the effect of the passage of time from the effect of price changes on the relative resource share function. For example, in our UK data, we do not observe price variation across households within a time period. Thus, we cannot identify the dependence of the relative resource share function on prices in any given time period using our UK data.Third, the vector of scale economies A_h is unknown and need not be estimated. The only restriction on A_h is that the value of the

---

$\widetilde{\gamma}_s(\ln v)$ over the log-budget that is shared by both men and women (there is no superscript j). At a given price vector P, the shapes of the Engel curves of men and women by 3 parametric shifters: a) $\ln \Gamma_s^j(P)$ translates the shape in the log-budget (horizontally); b) δ^j scales the shape in the log-budget (horizontally); and c) $d_s^j(P)$ is a vertical translation. Consequently, δ^j is identified under WSAP from the behaviour of singles: it is the vertical scaling of Engel curves over the log-budgets that makes (horizontally and vertically translated) Engel curves the identical for men and women.

element corresponding to the assignable good (in our case, clothing) is 1. Consequently, resource shares estimated using these identifying restrictions are compatible with any degree of shareability (ranging from fully shareable to not shareable at all) for any good, except for the assignable good which is assumed to be non-shareable. We discuss this in greater detail in Appendix (14).

Fourth, the meaning of the relative resource share (it is the fraction of adult expenditures consumed by adult women) does not depend on whether or not children are present, nor on whether or not children are members indexed by j or simply attributes of households or indexes of household type. Consequently, the researcher does not have to take a stand on "whether or not children are people". The estimate of, and meaning of, the relative resource share does not depend on this aspect of the model.

12 Online Appendix: Proof of identification under WSAP

In this appendix, we show that:

1. DLP Theorem 1 implies that the *relative* resource share is identified from the assignable goods Engel curves of men and women in a given household type if SAP is satisfied and resource shares are independent of the household budget y. Engel curves of children are not needed.

2. Donaldson and Pendakur (2004) show that, given the normalization $\delta^m = 1$, the scalar δ^f is identified if we observe person-level Engel curves as is the case, e.g., for singles.

3. Weakened SAP deviates from SAP only by the scalar δ^f. For known δ^f, application of DLP Theorem 1 shows that weakened SAP is sufficient to identify resource shares.

12.1 Relative Resource Shares are identified from just men's and women's Engel curves

The last two lines of the proof of Theorem 1 of DLP show that the limiting derivatives of Engel curves with respect to the household budget (in their notation, $\widetilde{\rho}_{ks}$) are sufficient statistics for the resource shares of all household members. Let θ_h^j be analogous to their $\widetilde{\rho}_{ks}$. Our θ_h^j is the observable λ-order partial derivative of the Engel curve for the assignable good of person j in a household type h, $w_h^j(y)$, with respect to the logged household budget $\ln y$. Their $\widetilde{\rho}_{ks}$ is the observable λ-order partial derivative of the Engel curve for the assignable good of person k in a household type s with respect to the logged household budget $\ln y$.

The last line of the proof of their Theorem 1 may therefore be restated in our notation as: Given DLP assumptions A1,A2,A3 and A4,

$$\eta_h^j = \theta_h^j / \left(\sum_{l \varepsilon h} \theta_h^l \right)$$

is identified if all θ_h^j for $j \varepsilon h$ are identified.

The relative resource share equals

$$R_h = \frac{\eta_h^f}{\eta_h^f + \eta_h^m} = \frac{\theta_h^f}{\theta_h^f + \theta_h^m}.$$

Note that θ_h^j is identified if $W_h^j(y)$ is observed and sufficiently differentiable.

Corollary to Theorem 1 of DLP: Given DLP assumptions A1,A2,A3 and A4, R_h is identified if $W_h^j(y)$ for $j = m, f$ are observed and sufficiently differentiable. ∎

12.2   Given WSAP, Singles' Engel Curves Identify δ

Let p be the scalar-valued price of the private assignable good; let $\widehat{p}$ be the $K - 1$ vector of the prices of all other goods. Let $P = [p \ \widehat{p}]$

Let the indirect utility of person j in household type h if they have individual budget y, V_h^j, be said to satisfy weakened similarity across people (WSAP, "weakened SAP") iff

$$V_h^j(P,y) = \psi_h^j \left(v_h \left(P, \left(\frac{y}{\Gamma_h^j(P)} \right)^{\delta^j} \right), \widehat{p} \right) \qquad (27)$$

where Γ_h^j is a function of prices, δ^j is a scalar, v_h is a function of prices and a transformed budget $\left(\frac{y}{\Gamma_h^j(P)} \right)^{\delta^j}$ and ψ_h^j is a function of v_h and non-assignable prices $\widehat{p}$. Normalize $\delta^j = 1$ and $\Gamma_h^j(P) = 1$ for one reference type of person. In the main text, we normalize $\delta^m = \Gamma_h^m(P) = 1$ for type m (men).

Denote $y_h^j(P,y) = \left(\frac{y}{\Gamma_h^j(P)} \right)^{\delta^j}$. We normalize for person m that $\Gamma_h^m(P) = \delta^m = 1$, implying $y_h^m = y_h^m(P,y) = y$.

We now derive demands just for singles, so $h = s$.

Lemma 1: Application of Roy's Identity to indirect utility given by (27) yields Engel

56

curves for the assignable good which satisfy WSAP given by equation (26):

$$-\frac{\frac{\partial V_s^j(P,y)}{\partial \ln p}}{\frac{\partial V_s^j(P,y)}{\partial \ln y}} = -\frac{\frac{\partial \psi_s^j(v,\widehat{p})}{\partial v}\left[\frac{\partial v_s(P,y_s^j)}{\partial \ln p} + \frac{\partial v_s(P,y_s^j)}{\partial \ln y_s^j}\left(-\frac{\partial \Gamma_s^j(P)}{\partial \ln p}\right)\right]}{\frac{\partial \psi_s^j(v,\widehat{p})}{\partial v_s}\frac{\partial v(P,y_s^j)}{\partial \ln y_s^j}\delta^j} = -\frac{1}{\delta^j}\frac{\left[\frac{\partial v_s(P,y_s^j)}{\partial \ln p}\right]}{\frac{\partial v(P,y_s^j)}{\partial \ln y_s^j}} + \frac{\partial \ln G_s^j(P)}{\partial \ln p}$$

or,

$$w_s^j(P,y) = \frac{1}{\delta^j}w_s^f\left(P,\left(\frac{y}{G_s^j(P)}\right)^{\delta^j}\right) + \frac{\partial \ln G_s^j(P)}{\partial \ln p} \tag{28}$$

∎

Remark: Engel curves of this form are a case of generalized equivalence-scale exactness, studied by Donaldson and Pendakur (2004) for unitary households. (They allow for δ^j to depend on prices P; in our case, δ^j does not depend on P). They show that if Engel curves for a single good are observed for at least two types including the reference type m, then:

- the parameters δ^j and functions Γ_h^j are identified if preferences are not PIGLOG

- the parameters δ^j are identified if preferences are not homothetic.

Remark: the Range Condition of Donaldson and Pendakur (2006) is satisfied if the functions Γ_h^j in (27) differ across j for a given h.

Lemma 2: Given the Range Condition of Donaldson and Pendakur (2006) and the assumption that preferences are not homothetic, δ^j are identified via Theorem 1 of Donaldson and Pendakur (2006) if Engel curves for the assignable goods of singles, $w_s^j(y)$ for $j = m, f$, are observed . ∎

Corollary 2: To connect with the main text analysis, if we observe $w_s^j(y)$ for $j = m, f$, δ is identified.

12.3 Given δ^j, Relative Resource Shares are Identified

Assumption A3b: for $j = m, f$ let indirect utilities be given by (27), implying that Engel curves are given by (28).

Theorem 1: Given DLP assumptions A1,A2,A4 and Assumption A3b above, the relative resource share R_h is identified if Engel curves for the assignable goods of single men and women and men and women in households, $w_s^j(y)$ and $w_h^j(y)$ for $j = m, f$, are observed and sufficiently differentiable.

Proof: If $\delta^j = 1$ for all j, we have SAP of Dunbar, Lewbel and Pendakur (2013). Corollary 1 of DLP Theorem 1, above, shows that in this case, the relative resource share is identified from the Engel curves of adult men and women in collective households.

If the scalar δ^j were known, we can replace y with y^{δ^j} and utilities of the form (27) would fit in the definition of SAP given Assumption A3 of the Appendix of Dunbar, Lewbel and Pendakur (2013). Consequently, for known δj, application of Theorem 1 of DLP and corollary 1 of DLP Theorem 1 (above) imply identification of the relative resource shares.

Suppose we observe assignable good Engel curves for people m, f in a household of type h and for singles. Then, δ^j is identified from singles, and we can plug y^{δ^j} for household Engel curves, and we have identification of R_h via Theorem 1 of Dunbar, Lewbel and Pendakur (2013) and corollary 1 of DLP Theorem 1 (above). ∎

Remark: Theorem 1 differs from Corollary 1 only if δ^j differs across j. Assumption A3b differs from DLP Assumption A3 only if δ^j differs across j.

13 Online Appendix: Empirical model given WSAP and PIGLOG

A parametric example may be helpful here (this is the parametric structure that we will take to the data). Let $V_h^j(p, y)$ be the indirect utility function of person j if they live in household h. Suppose that indirect utilities are in the PIGLOG class of Muellbauer (1975):

$$V_h^j(p, y) = \frac{\ln y - A_h^j(p)}{B_h^j(p)}$$

which implies, via Roy's Identity, shadow Engel curve functions of individuals for the assignable good, $\widetilde{w}_h^j(p, y)$,

$$\widetilde{w}_h^j(p, y) = \alpha_h^j(p) + \beta_h^j(p) \ln y \tag{29}$$

where

$$\beta_h^j(p) = \frac{\partial \ln B_h^j(p)}{\partial p_{assignable}}$$

and

$$\alpha_h^j(p) = \frac{\partial \ln A_h^j(p)}{\partial p_{assignable}} - \beta_h^j(p) A_h^j(p).$$

These demands[10] satisfy WSAP if and only if

$$\ln B_h^j(p) = \delta^j \ln B_h(p). \tag{30}$$

Letting

$$\beta_h(p) = \frac{\partial \ln B_h(p)}{\partial p_{assignable}},$$

we have that

$$\beta_h^j(p) = \delta^j \beta_h(p).$$

When applied to the case of linear Engel curves (which themselves are implied by PIGLOG indirect utility), the semiparametric restriction of Weakened SAP implies that the slopes of Engel curves are multiplicative in a person component δ^j and a household component $\beta_h(p)$. Arduini (2024) invokes this restriction in a reduced form setting, calling it "similarity in ratios across types". However, whereas Arduini provides a fully parametric result implemented with Cobb-Douglas utilities, our identification result is semiparametric and provides the class of utility functions consistent with the model. Sokullu and Valente (2021) also provide an identification result based on a multiplicative decomposition of slopes of Engel curves, but theirs decomposes over time periods rather than household types and person types.

The key difference between SAP and Weakened SAP is the following. The term $\beta_h^j(p)$ is the budget semi-elasticity of the assignable goods Engel curve for a person j if they live in household h. It is not directly observable. In SAP of DLP, this semi-elasticity must be the same for all household members (because $\delta^j = 1$), meaning that if the assignable good is, e.g., a luxury, it must be the same degree of luxury for both men and women. In contrast, with weakened SAP, it can be different degrees of luxury for men versus women.

The WSAP-restricted PIGLOG has demands of singles given by

$$W_s^j = \alpha_s^j(p) + \beta_s^j(p) \ln y = \alpha_j^j(p) + \delta^j \beta_s(p) \ln y. \tag{31}$$

Since we have not assume any error terms or other random variables, we have that for any (p, y), W_s^j

---

[10] In terms of the class of demand functions covered by the WSAP restriction (26) above, this PIGLOG structure implies

$$d_j^j(p) = \alpha_h^j(p)$$

is unrestricted and $g_h(v, p)$ has the form

$$g_h\left(\left(\frac{y}{G_h^j(p)}\right)^{\delta^j}, p\right) = \beta_h(p) \ln\left(y^{\delta^j}\right).$$

takes on a single value. Consequently, at the level of the population, since W_s^j, p and y are observed, we may assume that the partial derivatives of W_s^j with respect to $\ln y$, denoted $\partial W_s^j(p, \ln y)/\partial \ln y$, are identified. The observable log-budget-response of these Engel curves are related to the structural parameters via

$$\frac{\partial W_s^m(p, \ln y)}{\partial \ln y} = \delta^m \beta_s(p)$$

and

$$\frac{\partial W_s^f(p, \ln y)}{\partial \ln y} = \beta_s(p)$$

(because we normalized $\delta^f = 1$). So, we identify the structural parameter δ^m as the ratio of log-budget responses of singles' Engel curves:

$$\delta^m = \frac{\partial W_s^m(p, \ln y)}{\partial \ln y} \bigg/ \frac{\partial W_s^f(p, \ln y)}{\partial \ln y}$$

Substituting (29) under the restriction (30) into (23) gives Engel curves for collective households

$$W_h^j = \eta_h^j(p) \left[\alpha_h^j(A_h p) + \delta^j \beta_h(A_h p) \left(\ln y + \ln \eta_h^j(p) \right) \right]$$

$$= \left[\eta_h^j(p) \alpha_h^j(A_h p) + \eta_h^j(p) \delta^j \beta_h(A_h p) \ln \eta_h^j(p) \right] + \left[\eta_h^j(p) \delta^j \beta_h(A_h p) \right] \ln y \qquad (32)$$

Assuming that the partial derivatives of these are identified, the log-budget-response of these household Engel curves are related to the structural parameters via

$$\frac{\partial W_h^m(p, \ln y)}{\partial \ln y} = \eta_h^m(p) \delta^m \beta_h(A_h p)$$

and

$$\frac{\partial W_h^f(p, \ln y)}{\partial \ln y} = \eta_h^f(p) \beta_h(A_h p)$$

Finally, given identified δ^m, we identify the relative resource share R_h as

$$\frac{\delta^m \frac{\partial W_h^f(p, \ln y)}{\partial \ln y}}{\frac{\partial W_h^m(p, \ln y)}{\partial \ln y} + \delta^m \frac{\partial W_h^f(p, \ln y)}{\partial \ln y}} = \frac{\delta^m \eta_h^f(p) \beta_h(A_h p)}{\eta_h^m(p) \delta^m \beta_h(A_h p) + \delta^m \eta_h^f(p) \beta_h(A_h p)} = \frac{\eta_h^f(p)}{\eta_h^m(p) + \eta_h^f(p)} = R_h(p).$$

Note that R_h is identified at every p, so that identification is at the Engel curve level. That is, the relative resource share at a given price vector is identified from data on budgets y_i and assignable goods Engel curves W_i^j where households face that price vector.

If we observe data at many price vectors, the estimated slopes of Engel curves would differ

at each price vector. If prices vary over time, as they do in the data we use below, we would be able to identify the relative resource share relevant to each time period—even if we didn't observe prices—by stratifying the data by time and estimating separately in each period. However, we would not be able to identify the effect of prices on resource shares separately from the effect of time on resource shares.

13.1 Linear reduced form

Now, similar to Lechene, Pendakur and Wolf (2022), we show that there are linear reduced forms corresponding to singles' Engel curves (31) and household Engel curves (32) whose coefficients have a simple mapping to the resource shares. Recall that a subscript s indicates a person living alone as a single.

The first step is to consider data where we hold prices constant, and have many households facing a single price regime. Because prices p enter every structural function as an argument, this is without loss of generality. Next, because real data do not fit the model exactly, we add an error term ε to each equation. This error term may be interpreted as measurement error or preference heterogeneity (see Lewbel and Pendakur 2022), but not as specification error. The linearity of PIGLOG demands is essential to our linear reduced form.

As we are holding prices constant, with some abuse of notation, let $\alpha_s^j = \alpha_s^j(p)$, $\alpha_h^j = \alpha_h^j(A_h p)$, $\beta_s = \beta_s(p)$, $\beta_h = \beta_h(A_h p)$ and $\eta_h^j = \eta_h^j(p)$ (where η_h^j doesn't depend on y by assumption).

Substituting all this into (31) and (32) under the restriction (30), adding a subscript i for households $i = 1,...,N$ facing a common price vector and adding an error term, we have for singles a linear model (with coefficients in roman face):

$$w_{is}^j = a_s^j + b_s^j \ln y_i + \varepsilon_{is}^j \tag{33}$$

where

$$a_s^j = \alpha_s^j \qquad \text{and} \qquad b_s^j = \delta^j \beta_s \tag{34}$$

and for households:

$$w_{ih}^j = a_h^j + b_h^j \ln y_i + \varepsilon_{ih}^j \tag{35}$$

where

$$a_h^j = \eta_h^j \alpha_h^j + \eta_h^j \beta^j \ln \eta_h^j \qquad \text{and} \qquad b_h^j = \eta_h^j \delta^j \beta_h. \tag{36}$$

Note that the singles' reduced form is obtained from the household reduced form by substituting in $\eta_s^j = 1$.

Manipulation of (34) and (36) reveals that relative resource share is given by the following

function of reduced form slope parameters[11]

$$R_h = \frac{\eta_h^f}{\eta_h^m + \eta_h^f} = \frac{b_h^f b_s^m}{b_h^m b_s^f + b_h^f b_s^m}$$

or, equivalently,

$$R_h = \frac{b_h^f}{b_h^m \delta^f + b_h^f} \qquad \text{where} \qquad \delta^f = \frac{b_s^f}{b_s^m}.$$

If we observe assignable goods demands for both singles and households, the reduced form coefficients $a_s^j, b_s^j, a_h^j, b_h^j$ are all identified via linear estimation techniques like OLS and 2SLS. For example, we could regress the fraction of household expenditure commanded by for assignable goods of person j, W_{ih}^j, on a constant and the log-budget, $\ln y_i$, to recover a_h^j and b_h^j. Let $\widehat{b}_s^j$ and $\widehat{b}_h^j$ be estimates of b_s^j and b_h^j, respectively. Consequently,

$$\widehat{R}_h = \frac{\widehat{b}_h^f}{\widehat{b}_h^f \widehat{\delta}^f + \widehat{b}_h^m} \qquad \text{where} \qquad \widehat{\delta}^f = \frac{\widehat{b}_s^f}{\widehat{b}_s^m} \qquad (37)$$

Like LPW, this estimate depends only on estimated slopes of Engel curves; the levels of Engel curves embodied in the reduced form parameters a_s^j and a_h^j do not affect the estimate. Like LPW, this estimate may suffer from weak identification if the denominator is close to zero (much more on this below). But, unlike LPW, the estimated relative resource share does not impose the assumption that men and women have the same preference parameter $\beta_h^j = \beta_h$ governing the budget responses of Engel curves. That is, this identification strategy allows for the possibility that clothing is, e.g., more of a luxury for women than for men.

13.2 Weighted Average Estimator of δ

Suppose that for singles, we have

$$W_{ist}^j = a_s^{j\prime}(1 \quad z_i \quad \tau(t)) + \left(b_s^{j\prime} \tau(t) \quad b_s^{jz\prime} z_i \right) \ln y$$

where $\tau(t)$ is an L-vector of basis functions in time, including a constant term, and where $b_s^j = \left[b_{j0}, ..., b_{jL} \right]$ with b_{j0} being the coefficient on $\ln y$.

Under the restriction $b^{f\prime} \tau(t) = \delta b^{m\prime} \tau(t)$ and $b^{fz\prime} \tau(t) = \delta b^{mz\prime} \tau(t)$, we have

$$b^f = \delta b^m$$

---

[11] $\dfrac{b_h^f b_s^m}{b_h^m b_s^f + b_h^f b_s^m} = \dfrac{b_h^f \delta^m \beta_s}{b_h^m \beta_s + b_h^f \delta^m \beta_s} = \dfrac{\eta_h^f \beta_h \delta^m}{\eta_h^m \delta^m \beta_h + \eta_h^f \beta_h \delta^m} = \dfrac{\eta_h^f}{\eta_h^m + \eta_h^f}$

and

$$\delta = \frac{b^{fl}}{b^{ml}}$$

for all $l = 0, ..., L$ and

$$\delta = \frac{b^{fz}}{b^{mz}}$$

for all demographics in z.

Suppose we have estimates

$$\widehat{\delta}_l = \frac{\widehat{b}^{fl}}{\widehat{b}^{ml}}.$$

Any of these provides a consistent estimate of δ. Similarly for demographics z. We can test the hypothesis that all these estimated ratios are the same.

Suppose we additionally have estimated variances $\widehat{V}\left(\widehat{\delta}_l\right)$, then the inverse-variance weighted average

$$\widehat{\delta} = \frac{\sum_l \frac{\widehat{\delta}_l}{\widehat{V}\left(\widehat{\delta}_l\right)}}{\sum_l \frac{1}{\widehat{V}\left(\widehat{\delta}_l\right)}}$$

is a consistent estimator for δ that uses information about the relative variances of $\widehat{\delta}_l$ to provide a more efficient estimate than any single one of them.

Empirically, in our model with 6 basis functions, we get

$$\widehat{\delta} = 1.78 \qquad (.26)$$

which is similar to, but much less precise than, the GMM estimate used in the main text. The Wald test statistic of the hypothesis that the values δ_l are identical for all coefficients (6 time terms and 3 demographics for singles: age, education and benefit income share) is distributed as a χ_7^2 and has a sample value of 1.95 with a p-value of 0.75. So, as in the main text, we don't reject the hypothesis that male and female slopes are proportional to each other with a fixed-over-time constant of proportionality δ.

13.3 Total Effect of z

The marginal effect of z on the relative resource share is the vector $\lambda(t)$, which is a vector-function of reduced form parameters. Suppose we are interested in the element corresponding to marginal effect of woman's education, $\lambda^{womeduc}(t)$. This gives the effect of increasing a woman's education, holding all other covariates—such as the education of the spouse and income ratio—constant. However, this may not be an interesting thought experiment because those variables are not held

constant if a woman's education is higher. Indeed, a large part of the return to education is in match attained in the marriage market. Consequently, it may be interesting to consider the total (rather than marginal) effect of a covariate on the relative resource share. Using the omitted variables bias formula (e.g., the multivariate version presented in Angrist and Pischke 2014), this is equal to

$$\frac{dR_h}{dz_0} = \lambda_0(t) + \sum_{k \neq 0} \rho_k(t)\lambda_k(t)$$

where z_0 is the element of z we are interested in, and $\rho_k(t)$ is the regression coefficient from a population-level regression of z_k on z_0 and the summation excludes element 0. Substituting sample regression coefficients for $\rho(t)$ and estimates of marginal effects for $\lambda_k(t)$ provides an estimate of the total derivative expressed in terms of $\rho(t)$ and reduced form parameters.

This can be generalized to hold other variables constant. For example,

$$\frac{dR_h}{dz_0}|z_1 = \lambda_0(t) + \sum_{k \neq 0,1} \rho_k(t)\lambda_k(t)$$

where $\rho_k(t)$ is the regression coefficient on z_0 from a population-level regression of z_k on z_0 and z_1, and the summation excludes element 0. More concretely, if we want to know the total effect of the woman's education on the resource share, holding her age constant, we regress z_k for $k = 2, ..., K$ on the woman's education and age and take the coefficient on her education as an estimate of $\rho_k(t)$.

14 Online Appendix: Public/private goods approach vs Shareable/Non shareable

In this appendix, we discuss the relationship between scale economies, shareability and publicness of goods.

Although we do not estimate scale economies, we can be precise about the model of scale economies that we allow for. For *nonshareable* goods (aka: private goods), demand for the nonshareable goods $q_{nonshare}$ is given by:

$$\text{nonshareable: } q_{nonshare} = \widetilde{q}^m_{nonshare} + \widetilde{q}^f_{nonshare}$$

because if the male demands 1 unit and the female demands 2 units, the only way for the household to satisfy both demands for this good is to go out to the market and purchase 3 units.

In contrast, for shareable goods (including both public goods and shareable goods), the demand

for the shareable goods q_{share} is given by

$$\text{shareable: } q_{share} < \widetilde{q}^m_{share} + \widetilde{q}^f_{share}$$

because if the male demands 1 unit and the female demands 2 units, the household can satisfy both demands for this good by purchasing, e.g., only 2.5 units in the market.

Public goods are a special type of shareable good where

$$\text{shareable and public: } q_{share} = \frac{\widetilde{q}^m_{share} + \widetilde{q}^f_{share}}{2},$$

and satisfying the additional restriction that

$$\widetilde{q}^m_{share} = \widetilde{q}^f_{share}.$$

Some papers that seek to identify collective household models divide goods a priori into purely public (where $q_{share} = \widetilde{q}^m_{share} = \widetilde{q}^f_{share}$) and purely private goods (where $q_{nonshare} = \widetilde{q}^m_{nonshare} + \widetilde{q}^f_{nonshare}$). We do not use this strategy. Instead, we use the model of Browning, Chiappori and Lewbel (2013), described below, which has

$$q_k = a_k \left(\widetilde{q}^m_k + \widetilde{q}^f_k \right) \tag{38}$$

for each good $k = 1, ..., K$, where $\frac{1}{2} \le a_k \le 1$. It thus accommodates goods that are not shareable ($a_k = 1$) and goods that are shareable ($a_k < 1$), and has the advantage that the exact degree shareability of all goods need not be known in advance.

In this paper, we describe tools to identify and estimate resource shares η^j (or, functions of them), but not to identify shadow prices. However, the resource shares we identify will be consistent with a large class of models of shadow prices, and therefore for a large class of models of scale economies in household consumption. We will work with models where $\Sigma_j \widetilde{y}^j = y$, and therefore resources shares sum to 1: $\Sigma_j \eta^j = 1$. Resource shares are shares of expenditure, spent at shadow prices.

Figure 14: Marginal effect of the benefit share of gross income on women's relative resource shares

(a) Childless couples

(b) Couples with children

15 Online Appendix: Marginal Effects

15.1 Individual characteristics: Marginal effect of benefit share of household income

Similarly to figure 10 depicting the marginal effect of the wage ratio on the relative resource share, in figure 14, the right hand side $y-$axis shows the average benefit share at each date, and the left hand side $y-$axis shows the marginal effect of the share of benefit out of gross income on the relative resource share. These figures suggest that the marginal effect of the benefit share of household income might be negative for all women and is slightly rising for mothers (although not statistically significantly).

15.2 Individual characteristics: Effect of education on women's relative resource shares

We start by reporting the effect of men's education on women's relative resource shares.

Figure 15 gives the marginal effect of men's education on the women's relative resource shares. For childless women, for two decades, men's education had no effect on the relative resource shares (from 1978 to 2000), then the effect became positive. For mothers, we can see the same pattern, with a longer period without an effect of the men's education on the relative resource share before the effect becomes positive at the end of the period.

We next turn to consider the effect of women's education on women's relative resource shares.

The two panels of figure 16 show, on the right hand side vertical axis, the mean age at which women have left education, which is the thick dashed line. Age of leaving education has increased from less than 16 to close to 19 years old for women without children and from 16 to over 19 years old for mothers. The thick black line in these graphs is the marginal effect of the education

Figure 15: Marginal effect of men's education on women's relative resource shares

(a) Childless couples

(b) Couples with children

Figure 16: Marginal effect of women's education on women's relative resource shares

(a) Childless couples

(b) Couples with children

of women on the relative resource share, with the scale indicated on the left hand side vertical axis. For women without children, the effect is zero. For mothers, not only is it not significantly different from zero, as evidenced by the confidence intervals, but the point estimate of the marginal effect of education on the relative resource shares is in fact negative for the entire period.

Altogether, women's education has either no effect or a negative effect on their relative resource shares. In an environment with assortative matching, we might expect, for a fixed distribution of education, that more education leads to a higher bargaining power, and therefore a higher resource share. We might also expect that, as all women become more educated on average, the returns to education decline. This is not what we observe. It could be that the marginal effect of education is not what we need to measure. Indeed, the graphs are showing what happens if you increase a woman's education in a household, holding everything else constant. But when education increases, other things change: for instance, more education means a higher salary and a match with a man with different characteristics. Therefore, it might be that, as more educated women match with more educated men, because of assortative matching, in order to measure the effect of women's education on their resource shares, we should be looking at the total effect of education rather than at the marginal effect.

Figure 17: Total effect of women's education on women's relative resource shares

(a) Childless couples

(b) Couples with children
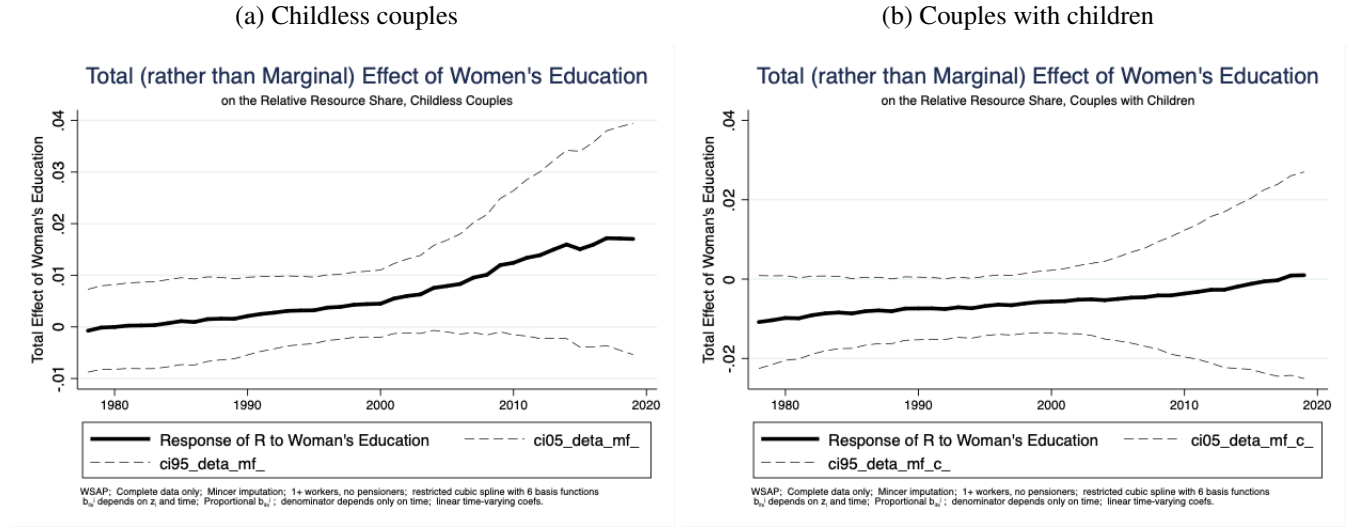
A simple way to get at the total effect of women's education on their relative resource share is to exploit the omitted variables bias formula to see what full load of education is when we assume that all other characteristics (such as mens' education and age and household benefits) are in fact driven by matching on education. We describe in the Appendix an estimator for this total effect of woman's education on the relative resource share, which sums its direct effect (the marginal effect described above) and indirect effects through the correlation of, e.g., men's education with women's education, and the return (on resource shares) to those correlated variables. Because the relative resource share is linear in covariates in each t, this is involves yet another nonlinear transformation of linear regression coefficients.

Figure 17 shows that the total effect of women's education on their relative resource shares is increasing (in the point estimates) over time for both childless women and mothers. It is positive and increasing for childless women and goes from negative to positive for mothers. However, if we look at the confidence intervals, the effect is still not significantly different from zero for either mothers or childless women.

16 Online Appendix: Robustness to Alternative Specifications

16.1 Do the restrictions that deliver R linear in z matter?

We impose 2 restrictions for convenience.

Restriction (14) says that, for singles, *slopes are proportional*. That is, the budget responses of single women are proportional to the budget responses of single men, with a factor of proportional-

ity that does not vary with demographics z or time t. We see from Figure 4 in the main text that the slopes of the Engel curves of the assignable good clothing for men and women are different from each other, which means that the assumption of similarity across people (SAP) does not seem to hold here. That assumption, which is often used for identification, has that the slopes of the Engel curves for the assignable good are the same across people for a given type of household, which in this case would translate into the slope of the Engel curve for men being equal to the slope of the Engle curve for women. In contrast, in our data, the slopes are not equal and women's Engel curves have larger slopes throughout the period. The slopes decline over time for both men and women indicating that clothing is becoming less of a luxury over time.

In the figure (in the main text), the slopes look roughly proportional, but are not exactly proportional (e.g., the ratio of slopes is smaller in 2019 than in 2000). We can estimate the value of δ under the restriction that the ratio doesn't change over time via GMM estimation of the conditional moment condition

$$E\left[W_{ist}^j - a_s^{j0} - a_s^{jz\prime}z_i - a_s^{j,splines}*splines_t - b_s^{j0}*\ln y_i - b_s^{j,splines}*splines_t*\ln y_i - b_s^{jz\prime}z_i | z, \ln y\right] = 0$$

under the restriction

$$b_s^{f0} = \delta b_s^{m0}$$

$$b_s^{f,splines} = \delta b_s^{m,splines}$$

$$b_s^{fz} = \delta b_s^{mz}$$

using instruments $1, z, splines, \ln y, z*\ln y, splines*\ln y$. The estimated value of δ is 1.80 with an estimated standard error of 0.068. [12]

The unrestricted model (shown in the Figure) has different splines with 6 basis functions each (including the constant) and 3 demographic characteristics (age, education, benefit share) for single men and single women, hence 18 parameters defining slopes of Engel curves. The restricted model has 10 parameters, because it only has one set of spline and demographic coefficients and one proportional shifter δ. Thus, there are 8 overidentifying restrictions. In the absence of these restrictions, the model is exactly identified. Consequently, we can use the J-test of overidentifying restrictions to test the proportionality restrictions. This χ^2 test statistic has 5 degrees of freedom and has a sample value of 12.9 with a p-value of 0.114. The hypothesis that the Engel curves of single men are proportional to those of single women is therefore not rejected.

Restriction (15) implies that demographics cancel out of the denominator in the expression

---

[12]Estimation of the parameter δ could be done via OLS, cf Appendix (13.2). The GMM estimator described here is more precise than the OLS based estimate, which is why we use it.

for the relative resource share, in other words that the *denominator of the resource share function depends only on time*. While the previous two restrictions relate to the Engel curves of singles, this is a restriction on the Engel curves of collective households. This is a linear restriction with 48 degrees of freedom. The sample value of the χ^2 test statistic for this test is 267 with a p-value of 0.000 (the 5% critical value is 69). Imposition of this restriction therefore does imply structure not supported by the data. However, the benefits of the restriction in terms of analytical simplicity and interpretation of the results are substantial, as illustrated in (5.5), since the restriction enables us to use Oaxaca decompositions. Furthermore, we show below that estimates that relax this restriction have the same patterns as the estimates obtained imposing it, but are less precise. Therefore, we retain this restriction for our main text analysis, but provide estimates that relax it below.

In the absence of the second restriction, the expression for the relative resource share is

$$\eta_{th}^f = \frac{b_{th}^f \delta}{b_{th}^m + b_{th}^f \delta} = \frac{\left(b_h^f(t) + b_h^{zf}(t)'z_h\right)\delta}{b_h^m(t) + b_h^{zm}(t)'z_h + \left(b_h^f(t) + b_h^{zf}(t)'z_h\right)\delta}.$$

The denominator of this expression may be close to zero for some values of *z*. To avoid weakly identified estimates, we drop observations of relative resource shares where the denominator is less than 0.003 in absolute value. The average value of the denominator in the entire sample is 0.03, so this cutoff is one-tenth as large. This drops 0.1% of observations in the entire sample.

In the expression above, the average factual relative resource share is not equal to the above expression evaluated at the average value of *z*. Thus, here we present average factual relative resource shares, which are comparable to the Figures presented in the main text. However, instead of giving standard errors for this average (which is cumbersome to compute analytically), we provide the 5th and 95th percentiles of the empirical distribution.

The first two figures below give the average and quantiles of the factual relative resource share distribution maintaining both restrictions. Here, the mean is identical to that shown in the main text. However, instead of giving a 90% confidence interval for the estimate of the mean, here we give the 5th and 95th percentiles of the distribution of relative resource shares evaluated over all the covariates in each year. The second two figures below give the average and quantiles of the factual relative resource share distribution maintaining only the first restriction (since it was not rejected by the data).

Mean and Q5 and Q95 of Estimated Relative Resource Shares
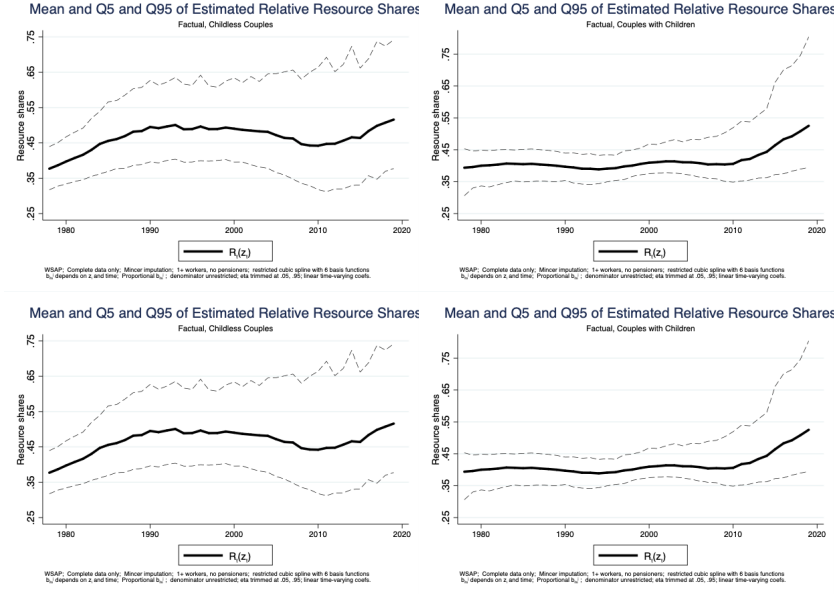
The 'eyeball check' suggests that the second convenience restriction—while rejected by the data—still leads to results consistent with the big picture about the evolution of mean resource shares over time.

However, it gives a different picture of the evolution of the dependence of resource shares on covariates over time. If we impose the restriction that the denominator of the expression for the relative resource share is invariant to demographics, the dependence of resource shares on demographics is driven entirely by the numerator. In this case, we observe that demographic variation drives about the same amount of variation in relative resource shares across the whole period.

In contrast, if we relax that restriction, then demographics drive more variation in relative resource shares at the end of the period than at the beginning of the period. Part of this increase in variation is spurious, due to to identification getting close to weak. Recall that the denominator declines over time. When we allow the denominator to depend on covariates, the risk of poorly identified estimates—and spuriously large estimated relative resource shares—rises towards the end of the period.

Our headline result concerned the increase in the average relative resource share. While it is cumbersome to do inference on the average relative resource share for this model, we can do inference for the relative resource share function evaluated at average characteristics. Below, we provide the factual estimates for childless couples and for couples with children for the model where the denominator is unrestricted. Here, we see that the statistically significant increase in relative resource shares for women in childless couples over the first 20 years that we saw in the main text is also evident.

Estimated Relative Resource Shares — Factual, Childless Couples / Factual, Couples with Children
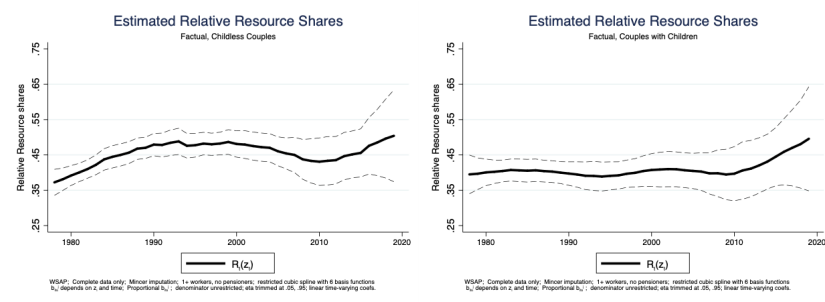
Below, we show the counterfactual, with demographics *z* held constant at their 1978 values. Here, we see that the statistically significant increase in the resource share function observed for childless women in the main text is again evident when we relax the denominator restrictions.

Estimated Relative Resource Shares — Counterfactual with fixed z_{1978}, Childless Couples
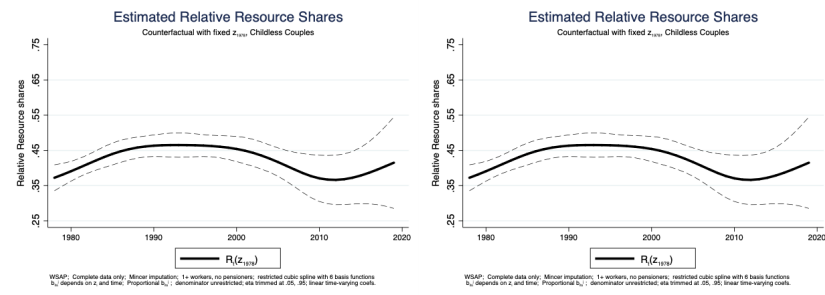
## 16.2 Does the Complexity of the Model with Respect to Time Matter?

The time-dependence in our baseline model is rather limited. We have a restricted cubic spline with 6 basis functions for the time trends and we have linear time-varying coefficients. Here, we consider relaxing this aspect of the baseline model, using a restricted cubic spline with 9 basis functions for the time trends and quadratic time-varying coefficients.

Below, we reproduce the estimated factual relative resource shares, evaluated at mean characteristics in each year, from the baseline and compare it to the estimated factual relative resource share with more complex time effects. Here, we impose both restrictions described above.

Estimated Relative Resource Shares — Factual, Childless Couples / Factual, Couples with Children
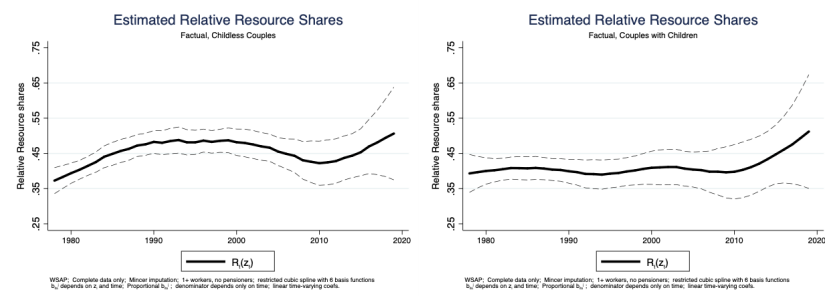
Now, we consider a more complex model with respect to time, where there are 9 basis functions (rather than 6 functions) in the spline and where we use quadratic (instead of linear) varying coefficients.

Estimated Relative Resource Shares
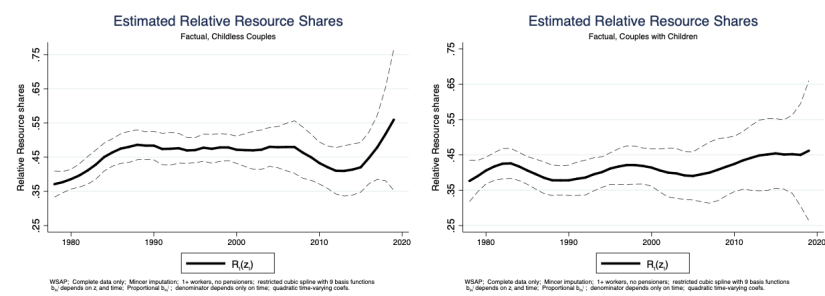
For childless couples, the more complex model shown below suggests that the increase in relative resource shares occurred earlier than indicated by the smoother model shown in the main text. In the main text, we talk about an increase in relative resource shares of about 13 percentage points over 1978 to 1992; but with the more complex model we see the same increase concentrated over 1978 to 1988. Otherwise, the two estimates of factual relative resource shares show the same magnitudes and timings of changes. Similarly, for mothers, we see relative resource shares that are very flat over 1978 to 2005 and then possibly rising after that, regardless of the complexity of the time effects.

Allowing for quadratic varying coefficients might also effect the estimated marginal effects. Below, we reproduce the estimated marginal effects of the wage ratio on women's relative resource shares, evaluated at mean characteristics in each year, from the baseline and compare it to the estimates with more complex time effects.

MFX of wagr on Relative Resource Share
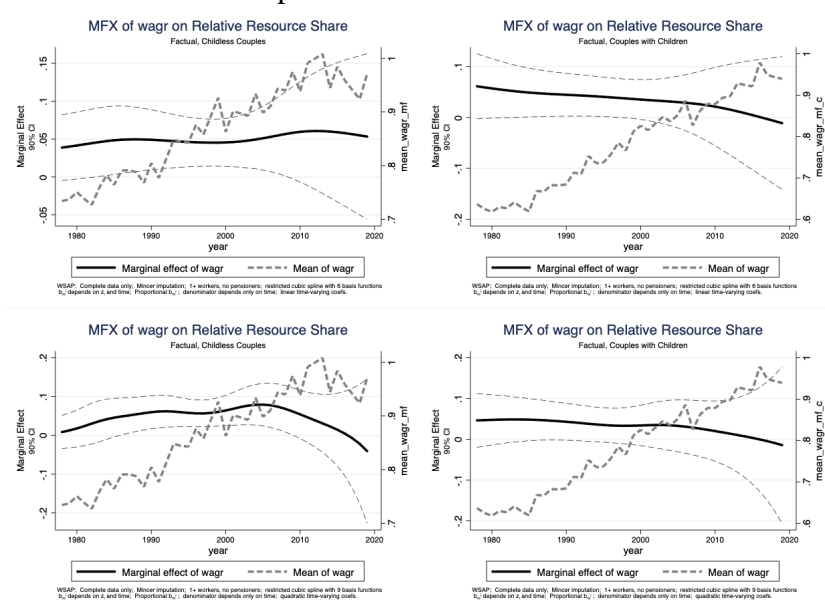
Here we see that for this covariate, allowing for quadratic varying coefficients does not affect out assessment of the marginal effects over time. It turns out that not a single quadratic term for a covariate is statistically significant. In the reduced form, this amounts to all the $z * \ln y * t^2$ terms being individually insignificant. Additionally, the estimated quadratic terms are small enough for

Figure 18: Marginal effect of women's share of market income on women's relative resource shares

(a) Childless couples

(b) Couples with children

all covariates that the graphs analogous to those shown above show that allowing quadratic varying coefficients does not affect the interpretation described in the main text for any covariate.

16.3 Robustness: Marginal effect of income share rather than wage ratio

Here we focus on the marginal effect of women's income share on women's resource share. We would expect women's resource shares to be increasing in their share of income, as we expect a positive relationship between income shares and bargaining power within the household. The effect we find in Figure 18 is indeed consistent with our *a priori:* we find a positive marginal effect of women's income share on women's relative resource shares(in terms of the point estimates).