DISCUSSION PAPER SERIES

DP15955

OLS Estimation of the Intra-Household Distribution of Expenditure

Valérie Lechene, Krishna Pendakur and Alexander Wolf

DEVELOPMENT ECONOMICS

CEPR

ISSN 0265-8003

OLS Estimation of the Intra-Household Distribution of Expenditure

Valérie Lechene, Krishna Pendakur and Alexander Wolf

Discussion Paper DP15955
Published 23 March 2021
Submitted 18 March 2021

Centre for Economic Policy Research
33 Great Sutton Street, London EC1V 0DX, UK
Tel: +44 (0)20 7183 8801
www.cepr.org

OLS Estimation of the Intra-Household Distribution of Expenditure

Abstract

We develop a method to estimate resource shares — the fraction of total household expenditure allocated to each household member — using OLS estimation of Engel curves. The method reframes the nonlinear model of Dunbar, Lewbel and Pendakur (2013), extended to allow single parent households, scale economies in assignable goods, complementarities between non assignable goods, and supplemented by a linear identification test. We apply the model to data from 12 countries, and investigate resource shares, gender gaps, and poverty. We reject equal sharing, and find large gender gaps in resource shares, and consequently in poverty rates, in some countries.

JEL Classification: D13, D63, I32

Keywords: collective households, resource shares, intra-household inequality, poverty

Valérie Lechene - v.lechene@ucl.ac.uk
University College London and CEPR

Krishna Pendakur - krishna_pendakur@sfu.ca
Simon Fraser University

Alexander Wolf - alexfwolf@gmail.com
Compass Lexecon

Acknowledgements
The authors acknowledge the financial support of the following agencies: the Insight Grant Program of the Social Sciences and Humanities Research Council of Canada (Pendakur); Economic and Social Research Council (Lechene); and, the World Bank. We also thank our friends and colleagues who gave us feedback on this work: Orazio Attanasio, Andrew Chesher, Geoffrey Dunbar, Dave Freeman, Arthur Lewbel, Kevin Milligan and Kevin Schnepel.

OLS Estimation of the Intra-Household Distribution of Expenditure

Valérie Lechene, Krishna Pendakur, Alex Wolf*

January 21, 2021

Abstract

We develop a method to estimate *resource shares* — the fraction of total household expenditure allocated to each household member — using OLS estimation of Engel curves.

The method reframes the nonlinear model of Dunbar, Lewbel and Pendakur (2013), extended to allow single parent households, scale economies in assignable goods, complementarities between non assignable goods, and supplemented by a linear identification test.

We apply the model to data from 12 countries, and investigate resource shares, gender gaps, and poverty. We reject equal sharing, and find large gender gaps in resource shares, and consequently in poverty rates, in some countries.

*JEL codes: D13; D63; I32.
Keywords: collective households; resource shares; intra-household inequality; poverty.
Acknowledgements: The authors acknowledge the financial support of the following agencies: the Insight Grant Program of the Social Sciences and Humanities Research Council of Canada (Pendakur); Economic and Social Research Council (Lechene); and, the World Bank. We also thank our friends and colleagues who gave us feedback on this work: Orazio Attanasio, Andrew Chesher, Geoffrey Dunbar, Dave Freeman, Arthur Lewbel, Kevin Milligan and Kevin Schnepel.

1 Introduction

Many aspects of well-being depend critically on the intra-household distribution of resources. For example, the Millenium Development Goals include the promotion of gender equality and the empowerment of women. Some of this has to do with women's access to resources within households, and achieving these goals requires measuring the consumption of women, as opposed to the consumption of households.

Similarly in labour, public and development economics, many important questions hinge on the intra-household distribution. For example, it is well-understood that poverty in childhood has long-term negative consequences (see, e.g., Campbell et al 2014), but much of the literature studying the consequences of childhood experiences is underpinned by the assumption that if a child is in a low-income household, then they have low consumption. But what if parents devote a greater fraction of resources to their children than to themselves? No question regarding the measurement of poverty, its persistence and correlation with long-term outcomes, can be answered properly if the tools we use preclude within-household inequality.

Allowing for the possibility of within-household inequality requires measuring how household expenditure is divided across the people inside households. Conceptually, this is a simple exercise consisting in measuring the expenditures of all the individuals in a household, and comparing them. But this exercise is frustrated by a lack of data on expenditure at the individual level. For instance, we may observe in the data that the household bought a bottle of milk, but we generally do not observe who drank it. Furthermore, there are goods with different degrees of shareability inside households, such as a common dwelling or shared means of transport, and ascribing a value to the services from the use of these goods to each individual is not straightforward. Our approach accomodates both the facts that people may have unequal access to resources and that some goods may be shared (to unknown degrees).

A useful tool to describe the within-household distribution of expenditure is the *resource share*, defined as the fraction of total expenditure allocated to each household member. If, in a given household, women have smaller resource shares than men then there is gender inequality in expenditure. Further, in that case, there is the possibility that although the household may have enough resources to keep all its members out of poverty, the women may have such poor access to resources that they are nonetheless poor.

The standard World Bank poverty measurement strategy assigns to each household member their per-capita share of household expenditure, and compares that to, e.g., the extreme poverty threshold of US$1.90 per day. But this is a matter of convenience and data availability, rather than a matter of principle: it has neither behavioural nor theoretical justification.

A strategy that respects the idea that consumption lives at the individual level would instead assign each person their share of household expenditure and compare that to US$1.90 per day.

Resource shares can help us understand a wide variety of phenomena. For example, Calvi (2019) estimates resource shares and poverty at the individual level in India and finds that women—especially older women—have lower resource shares than men. This then implies that older Indian women have much higher poverty rates than previously thought. Calvi shows that these higher poverty rates (driven by lower resource shares) among older women can explain the finding of Anderson and Ray (2010) that Indian women over the age of 45 have higher mortality rates than do Indian men (a phenomenon they call "missing women").

Many researchers have studied the consequences of unequal sharing within households using reduced form approaches. For example, Jayachandran and Pande (2017) provide evidence that Indian children further down the birth order are considerably more stunted, which they attribute to favoritism for first-born children. But does this favouritism run through a channel of greater access to household resources, that is, higher resource shares? If expenditure was observed, our econometrically simple method would illuminate this channel. Indeed, our method could provide reduced form researchers with an appealing structural methodology to supplement their analyses.

Our proposed method is in the lineage of Chiappori (1992)'s seminal contribution, which develops models of *collective households*, defined as households comprised of individual people who maximize utilities, and together reach the Pareto frontier. Using this general framework, Browning, Chiappori and Lewbel (2013) and Dunbar, Lewbel and Pendakur (2013, hereafter DLP) introduce structural models that allow us to use off-the-shelf data, of the sort collected routinely by statistical agencies, to reveal the resource shares of individual household members. DLP's model requires the use of nonlinear models that can be computationally difficult to estimate. The core of the computational difficulty lies in that resource shares must be between 0 and 1, and they enter the model nonlinearly, implying that bounded nonlinear estimation is required. However, many researchers in applied economics lean towards linear methods, in part because of their transparency. Therefore a simple linear model to estimate resource shares will significantly broaden our ability to measure within household inequality and individual poverty.

In this paper, we provide a linear reframing of the structural model of DLP yielding a theory-consistent and easy-to-implement model, requiring only the estimation of linear Engel curves for assignable goods. (An *Engel curve* relates the fraction of total household expenditure spent on a good to total household expenditure on all goods, at a fixed price vector; an *assignable goods* is one where we observe expenditure at the person level rather than at the household level, e.g. women's clothing.)

Our main methodological contribution is to show that, conditional on covariates, the model of DLP can be written as a linear reduced form wherein the structural parameters—the resource shares—are nonlinear functions of the reduced form coefficients. Furthermore, we extend the model of DLP to allow for complex household types, including those with multiple adult men and/or women and single parent households (see also: Calvi 2019, Dunbar, Lewbel and Pendakur 2019 and Calvi et al 2020), and to our knowledge, ours is the first paper to consider the latter households. We also extend the model of DLP to allow for assignable goods that have scale economies in consumption, and non-assignable goods that have complementarities (and scale economies) in consumption. Finally, we provide a test based on OLS regression that indicates whether the model is identified. Previous tests of this identifying restriction required nonlinear estimation methods; our linear reframing of DLP delivers an OLS test.

In this model, the *levels* of Engel curves pick up a mixture of differences across household members' resource shares and differences across household members' preferences. But the *slopes* of Engel curves with respect to the household budget are driven only by differences across household members' resource shares. So, resource shares are identified by the slopes of Engel curves: if the slope of a person's assignable good Engel curve is twice as large (in absolute value) than another person's, then their resource share is twice as large. Empirically, we run OLS regressions of observed household-level spending on assignable goods (as a fraction of total household expenditure) on observed household demographics, log total household expenditure and their interactions. Then, we compute estimated resource shares as functions of estimated regression coefficients.

We use the model to estimate resource shares and individual poverty rates (including women's poverty and children's poverty) with data from 12 countries, using household surveys from the World Bank LSMS data, and 1 other national survey from Bangladesh. We use person-level clothing expenditure as the assignable good. Clothing Engel curves pass the identification test for 5 of 12 countries, so we estimate estimate resource shares and person-level poverty for these countries.

We find that equal sharing—the implicit assumption underlying standard household-level poverty calculations—is rejected by the data, and we find evidence of gender gaps in resource shares and poverty rates in some countries. For example, we find estimated women's resource shares to be 5 and 4 percentage points lower than men's in Bangladesh and Iraq, respectively. This results in women's poverty rates that are 9 and 2 percentage points higher than men's in Bangladesh and Iraq, respectively.

Our data from Bangladesh have both person-level clothing expenditure and person-level food expenditure (including implicit expenditure on home-produced food). We find that using food data to identify resource shares delivers estimates that are similar to those generated from clothing data.

Given that we offer a simple and tractable methodology with low data requirements, we hope that researchers focused on intra-household inequality and its consequences will adapt their data collection strategies accordingly. In particular, field experimentalists and statistical agencies could add to their surveys questions on total household spending and person-level expenditure on at least 1 assignable good. Such data would be sufficient to estimate resource shares.

In section (2), we review the theoretical foundations of our work, and discuss identification of resource shares in DLP's nonlinear model. We then show our linear reduced form. In section (3), we present the data, and in section (4), estimated resource shares, gender gaps and poverty rates. We finish with a brief discussion of the implications of our work.

2 Theory

Consider for the moment a world where all goods are non-shareable (we will come to the case where some goods are shareable shortly). Let *expenditure* on a good be the quantity of the good times its price.[1] Dream data to measure the expenditure of individuals within households would look like Table 1a. Here, we directly observe the expenditure on each good by the man, woman and child in a nuclear household with 1 child. Suppose the poverty line is \$1.90 per person per day, which defines a household-level poverty line of \$2080. Since this household has total expenditure of only \$1850, standard poverty measure measurement (which assumes equal division within the household) would call all members of this household "poor".

However, with the dream data here, we observe the (unequal) expenditure level (and thus the resource share) of each person, and can compare individual expenditure levels to individual poverty thresholds. The column totals give individual total expenditure levels. The man's total expenditure on all goods is \$800, so his resource share (fraction of total household expenditure on all goods) is 43% (equals 800/1850). The individual poverty threshold is \$1.90 per day, equalling \$694, so the man is not poor. However, the woman's total expenditure is \$600, and her resource share is 32%. Since her expenditure falls below \$694, she is poor. Similarly, the child's total expenditure is \$450 and their resource share is 24%, and they are poor. Thus, the dream data reveal within-household inequality in resource

[1] In principle, we include home-produced quantities, so expenditures include the imputed value of home-production.

shares, and the fact the some members are poor while others aren't.

Table 1a: Dream Data

	Man	Woman	Child	Total
Food	400	300	200	900
Clothing	50	75	25	150
Shelter	100	100	100	300
Transport	250	125	125	500
Total	800	600	450	1850

Table 1b: Real Data

	Man	Woman	Child	Total
Food				900
Clothing	50	75	25	150
Shelter				300
Transport				500
Total				1850

If data like those in Table 1a were widely available, poverty measurement at the person level would be straightforward. Cherchye, Demuynck, De Rock, and Vermeulen (2017) collected this type of data for the Netherlands, and used it to, among other things, estimate consumption inequality within households. Brown, Ravallion and van de Walle (2019) use this type of data to investigate individual-level poverty and food deprivation. Bargain, Lacroix and Tiberti (2019) use this type of data to validate the modeling assumptions of collective household models. To our knowledge, these are the only cases where individual-level expenditure data for all (or most) consumption categories is collected.

In this paper, we use the terms *shareable* and *non-shareable* to refer to the consumption technology of the household and to related scale economies in household consumption. A nonshareable good has the property that the quantities consumed by each person add up to the total quantity purchased by the household. For example, food may be nonshareable because food eaten by one member cannot be eaten by another, so if two members eat 1 unit each, the household must buy 2 units. In contrast, a shareable good has the property that the quantities of each person add up to more than the quantity purchased by the household. For example, if two people ride a motorcycle together, they both consume a motorcycle ride, but the household only has to purchase gasoline for one motorcycle ride. If the two people ride together only part of time, then this good is partially shared. If they ride the motorcycle together all the time, then it is fully shared.[2]

A key feature is that shareable goods feel cheaper to people living in households than they do to people living alone, but non-shareable goods feel just as expensive to people living in households as they do to people living alone. Thus in Table 1a, if we assume that food is non-shareable, then the $400 expenditure on food for the man buys the same quantity that it would if he were living alone. In contrast, if we assume that shelter were fully shareable, the

[2] *Public* goods are a particular type of shareable good. First, public goods are *fully shareable* in the sense that the household can attain a quantity level of q for each household member by spending pq on that good. (For a nonshareable good, it would have to spend Npq, where N is the number of members.) Second, for public goods, each member must consume *the same amount* (equal to q).

$100 expenditure on shelter for the man would buy a larger quantity of shelter than it would if he were living alone. In principle, that $100 expenditure for the man in a three-person household could buy three times the quantity attained by a $100 expenditure for a person living alone (see Appendix 1 for a full description of this).

Thus, the benefit of shareability of goods to people in households is that the shadow price of consumption of those goods in the household is lower than the market price of those goods. Regardless of how shareable different goods are, the resource share of each person has the same interpretation: it is the fraction of total household expenditure spent by each person. Because of sharing, this spending results in greater quantities consumed than if the person were living alone.

Real-world expenditure data tend to look more like Table 1b. In this type of data, we see *household-level* expenditure for all the goods and services comprising total expenditure, and we may see one or two goods at the person level (in this case clothing). Such data are widely available in rich countries, because they are collected by statistical agencies that estimate the rate of price inflation, and are increasingly common in developing countries, in part due to international research efforts like the 100+ datasets in the Living Standards Measurement Study (LSMS) of the World Bank. So, with real-world data, we face an incomplete data problem: we do not have full data on individual expenditure; instead, we have data on just 1 or 2 commodities collected at the individual level.

DLP aim to address this missing data problem via modeling the allocation problem of the household, and "backing out" the individual resource shares from these incomplete data. In particular, DLP uses information on individual-level spending on non-shareable assignable goods (clothing in the Tables) to infer the individual-level to total expenditure of each person (the bottom row of Table 1a). DLP *does not* infer the individual-level expenditure on any particular good, just the individual-level total expenditure on all goods. Importantly, the fact that that the man has less clothing expenditure than the woman does not imply that he has less total expenditure than her. Instead, the link between individual assignable goods expenditure and individual-level total expenditure is driven by response of the former to the total household budget. If the man's clothing expenditure responds more to the household budget than does the woman's, then he has a larger claim on household resources than she does.

If no goods are shareable, then there are no scale economies in household consumption; if some goods are at least somewhat shareable, then there are scale economies in household consumption. The methods proposed by DLP allow for unknown degrees of shareability for each consumption good, and therefore for unknown scale economies in consumption. But, they require that at least one good is assignable (observed at the individual level) and without scale economies. Their methods reveal resource shares whose interpretation does not depend on the degree of shareability of the non-assignable goods. One of our contributions is to show that the methods apply, regarless of whether the assignable good is shareable.

Browning, Chiappori and Lewbel (2013: BCL) provide a very general efficient collective household model with scale economies in consumption, preference heterogeneity across people, and possibly unequal distributions of household resources. DLP take that model and impose sufficient restrictions on it to make it implementable with real-world data via nonlinear estimation of household-level Engel curves for assignable goods. In the next subsection, we briefly sketch those models and show how we extend the model of DLP to allow for non-nuclear households and for more general scale economies and complementarities. Following this, in (2.2) we show a linear reframing of DLP that may be estimated by ordinary least squares. We then describe how to add covariates to the model (2.3). We then show the conditions under which the model is exact and the conditions under which it is an approximation (2.4). Finally, we provide an OLS test of identification (2.5).

2.1 An Efficient Collective Household Model

Collective households are households comprised of a collection of individuals. The individuals have utility functions; households are just environments in which individuals live. *Efficient* collective household models are those in which the individuals in the household are assumed to reach the (household) pareto frontier. Like in earlier results in general equilibrium theory, the assumption of pareto efficiency is very strong: it means that the household-level allocation problem is observationally equivalent to a decentralised, person-level, allocation problem. In this decentralised allocation, each household member demands a vector of consumption quantities given their preferences and a personal budget constraint, and the household purchases the sum of these demanded quantities (adjusted for shareability/economies of scale).[3]

The model thus has us picture the household as a machine that makes budget constraints for its members. Each person's budget constraint is characterised by a shadow budget and a shadow price vector. They are "shadow" budgets and prices because they govern each person's consumption demands but they are not observed and do not equal the observed household budget or market prices. Welfare analysis of people in living in households considers these shadow budget constraints.

Let $h = 1, ..., H$ index households. Let t index the types of individuals, in our case, m for

[3] Much effort has gone into testing the restriction of efficiency, e.g., Cherchye et al (2007), Attanasio and Lechene (2014) and Rangel and Duncan (2019) and many others. In this paper, we take efficiency as a maintained hypothesis.

adult male, f for adult female and c for children. Let household h consist of N_h^t individuals of each type t, and let $N_h = \sum_t N_h^t$ be the total number of individuals in household h. The types of individuals are in some sense defined by the data, as we will see below. Let y_h denote observed household income (budget). Each type of person gets a shadow budget, and these shadow budgets must add up to the full household budget.

The share of the household budget allocated to persons of type t in household h is called their *resource share*, denoted η_h^t. Resource shares sum to 1 in each household h so that $\sum_t \eta_h^t = 1$. They may in general depend on household budgets, prices, household and individual characteristics (including so-called "distribution factors"). Most importantly, they can vary across the types of individuals in the household, for example, men's and women's resource shares are not assumed to be equal.

Within types, we assume that resources are distributed equally (if there is one person for each type, then this is not restrictive).[4] For example, in a household with two children where the children's resource share is $\eta_h^c = 0.40$, we have that 40 per cent of the household budget is allocated to children, with 20 per cent going to each child. In general, the total shadow budget of all the people of a given type t in a household h is $\eta_h^t y_h$, and the shadow budget of each person of that type is $\eta_h^t y_h / N_h^t$.

DLP allowed for multiple children, but not multiple adults of a given gender. Here, we allow for multiple members of any type (see also Calvi 2019, Dunbar, Lewbel and Pendakur 2019 and Calvi et al 2020). This extension is trivial mathematically, but is vital to allowing the model to accomodate the complex multigenerational households observed in many developing countries.

Shadow prices for goods are the within-household prices of consumption. Let $\boldsymbol{p}$ denote the market price vector for goods and let $\widetilde{\boldsymbol{p}}$ denote the shadow price vector of goods. Shadow prices are the same for all household members. If they were not the same, then there would be gains from trade across household members, a violation of the assumption of efficiency.

Shadow prices are lower than market prices for shareable goods. The more shareable the good, the lower is its shadow price of consumption within the household.

Without complementarities, the shadow price is the value such that the sum of quantities consumed by each member, q^t, multiplied by the shadow price, equals the product of the quantity purchased by the household multiplied by the market price: for each good, $\widetilde{p}\,(q^m + q^w + q^c) = pq$, so that $\widetilde{p} = \frac{q}{q^m+q^w+q^c}p$. For goods without scale economies, $q = q^m + q^w + q^c$, and the shadow price is equal to the market price. Shadow prices may in

[4]In order to estimate the resource share of a type (e.g. children), in the linear reframing, we don't need to make any assumption on sharing within that type. However, for welfare analysis, some assumption must be made.

principle be as low as p/N_h for a good that is so shareable that each household member can consume the purchased quantity.

There are also intermediate cases. Consider a motorcycle, used every day by the man, riding together half the time with the woman and half the time with the child. Here, the total quantity of transport consumed by the household is the sum of the man's quantity, q^m, the woman's quantity q^w and the child's quantity, q^c, but also, $q^w = q^c = \frac{1}{2}q^m$. However, the quantity of transport purchased by the household is only q^m. Thus, the shadow price of transport $\widetilde{p}$ is such that the household's expenditure on transport, pq^m, is equal to $\widetilde{p}\,(q^m + q^w + q^c)$. Therefore the shadow price of transport for this household is $\widetilde{p} = \frac{1}{2}p$.

Economies of scale arise because of the nature of the goods and of the size of the household. If there are scale economies in household consumption, shadow prices are less than market prices, and individuals face lower prices within households than when they live alone. (Holding the budget constant, lower prices means higher utility.)

Note that shadow prices of shareable goods are not the same as Lindahl prices for public goods. Lindahl prices are *individualized* prices that rationalize a quantity choice that is the same quantity level across different household members: the Lindahl price for a household member equals their marginal willingness to pay for the good (see Lindahl 1919 and Chiappori 1992). Shadow prices in this paper, as in BCL and DLP, are *household level* prices that determine quantity choices for all household members, and those quantity levels may differ across household members: the shadow price for a good embodies the scale economies relevant for that good.

Resource shares and shadow prices are interesting because they completely characterise the budget constraint of each person in the household. To the extent that budget constraints provide welfare measures, resource shares and shadow prices allow us to peer inside the household to consider the distribution of welfare therein. See BCL and Pendakur (2018) for a discussion of how to use resource shares and shadow prices in welfare analysis.

In this work, we identify resource shares from household-level consumption choice data, but we do not try to identify shadow price vectors (which reveal scale economies in consumption due to sharing).[5] Resource shares are interesting even without knowledge of shadow prices. First, resource shares provide a measure of consumption within the household: higher resource shares mean higher consumption. Second, they speak to inequality within the household: if resource shares are very unequal, then there is a lot of inequality within the household. Third, resource shares may respond to policy variables in the context of poverty

[5]Our methodology estimates resource shares at a given price vector, without knowledge of prices. Since we don't observe market prices, we cannot estimate shadow prices. Other methodologies use observed price variation to identify shadow prices and thus scale economies (e.g., BCL and Pendakur 2018).

reduction. If we can find policy variables that shift resource shares upwards for disadvantaged individuals, then their poverty rates may decrease.

An *assignable good* is one where we can observe the quantity consumed by, or expenditure of, each (type of) individual. Such goods are very useful for identification of household models (see, e.g., Chiappori and Ekeland 2009).[6] We will assume that there is a single assignable good observed in the data. The main assignable good we use in this paper is clothing; we observe the clothing purchases for the 3 types of individuals, adult males, adult females and children.

Like DLP, we do not allow for externalities in the consumption of the assignable good, where one person's assignable good consumption affects another person's demand functions. As our main assignable good is clothing, ruling out externalities may be unpalatable (see Cherchye et al 2007 for a discussion of how to incorporate externalities into collective household models). Nonetheless, it is a maintained assumption in this work.

Unlike DLP, we allow for the possibility that the assignable good has economies or diseconomies of scale (an $\boldsymbol{A}$ matrix element not equal to 1). Consider the example of food waste in food preparation. Suppose that one portion of food is wasted whatever the number of portions are prepared. For example, to prepare 3 portions, the household must buy 4 portions (since one is wasted), so that 1/4 of food purchases are wasted. Then, $3\widetilde{p} = 4p$, so that $\widetilde{p} = \frac{4}{3}p$ for a household with 3 people. For a household with N_h people, $1/N_h + 1$ of food purchases are wasted, and $\widetilde{p} = \frac{N_h+1}{N_h}p$. These are diseconomies of scale, that decrease with household size. If consumed food quantities are observed at the individual level, then food could be an assignable good used to estimate resource shares, even in the presence of scale economies in food consumption. Allowing for this possibility does not change the DLP demand equations, and so does not affect estimation. We show this in the Appendix.

Generally speaking, the available data on assignable goods will define the typology of individuals. In many data sources, assignable spending on clothing is available for adult men, adult women and children. But it could be recorded by gender for children as well as for adults, in which case there would be 4 types of individuals, adults or children and males or females. This is why we specify the model in terms of types of individuals rather than in terms of individuals themselves. If we had an assignable good for each individual, we could estimate the resource share of each rather than estimating the resource shares by type.

BCL show that, given the model described above, household demands are related to individual demands in an intuitive way. Given the sharing in the household, the household purchases enough of each commodity (adjusted for sharing) so as to give each individual in the household exactly what they would have purchased had they faced their individual

[6] Chiappori and Ekeland refer to nonshareable assignable goods as *private* assignable goods.

shadow budget and the shadow price vector.

Much work on consumer demand estimation models *Engel curves*. The Engel curve of a good is the fraction of the overall budget (spent on all goods) commanded by that good (Engel 1857, 1895).[7] Engel curve functions hold prices constant at some vector, and evaluate the fraction of expenditure as a function of the total household budget (and possibly other demographic characteristics). DLP derive the implications of the BCL model on Engel curves for assignable goods.

Let $\boldsymbol{p}$ be the price vector, sorted so that its first element is the price for the assignable good. BCL consider models where the shadow price vector, $\widetilde{\boldsymbol{p}}$, is linear in the market price vector, with $\widetilde{\boldsymbol{p}} = \boldsymbol{A}\boldsymbol{p}$ where $\boldsymbol{A}$ is a square matrix (they consider more complex models as well). To see how $\boldsymbol{A}$ matters, suppose first that it is a diagonal matrix. In this case, nonshareable goods have a corresponding $\boldsymbol{A}$ element of 1, so that the shadow price equals the market price. Shareable goods have corresponding elements that are less than 1, so that the shadow price is less than the market price. The off-diagonal elements of $\boldsymbol{A}$ allow goods to be complements in household consumption. For example, if food and cooking fuel are complementary, the off-diagonal element corresponding to that combination would be positive.

Whereas BCL allow for a unrestricted $\boldsymbol{A}$ matrix (with any numbers in both diagonal and off-diagonal elements), DLP consider a restricted $\boldsymbol{A}$ matrix, where $\boldsymbol{A}$ is a diagonal matrix, with 1 in the element corresponding to the assignable good. This means that the assignable good must be nonshareable (have no scale economies), so that its shadow price equals its market price. It also means that there are no complementarities in the household consumption technology.

It turns out that DLP assumed too much.[8] We show in the appendix that one can derive the same demand equations as DLP with a weaker restriction on the scale economies as embodied in the matrix $\boldsymbol{A}$. Here, we require that $\boldsymbol{A}$ is a block-diagonal matrix $\boldsymbol{A}$ satisfying

$$\boldsymbol{A} = \begin{bmatrix} A_1 & 0 \\ 0 & \boldsymbol{A}_2 \end{bmatrix} \tag{1}$$

where A_1 is a scalar giving the price scale that multiplies the market price of the assignable good to get the shadow price of the assignable good. So, the assignable good can exhibit

[7]Engel curve functions are often called "budget share" functions, for obvious reasons. We use the phrase Engel curve rather than budget share so that it is not confused with "resource share".

[8]Specifically, DLP assumed too much for identification of resource shares given their similar across people (SAP) preference restriction. For that, the restriction (1) is sufficient for identification. But, for identification of resource shares given their similar across types (SAT) restriction, allowing $A_1 \neq 1$ doesn't make sense in the context of the model.

scale economies, because A_1 can equal values other than 1. The matrix $\boldsymbol{A}_2$, which governs the scale economies of (and complementaries between) all other goods, is unrestricted.

The restriction on the matrix $\boldsymbol{A}$ is that the off-diagonal blocks are 0, which rules out complementarities in consumption between the assignable good and all other goods. This means that other than knowing that the quantity of one good per person is observed and assignable, we don't need to know *a priori* about the degree of shareability or scale economies of any good (including the assignable good), or about complementarities across non-assignable goods.

The structure on the matrix $\boldsymbol{A}$ plays a sideshow role here. We do not try to estimate it in this paper. Tractable estimation of $\boldsymbol{A}$ is a job for future research (see, e.g., Calvi et al 2020). Instead, it defines the set of models for which our estimated resource shares are valid. The interpretation of the resource share is the same no matter what value $\boldsymbol{A}$ takes: the resource share is the fraction of the (observed) household budget enjoyed by a type of person, and spent at shadow prices $\boldsymbol{A}p$.

Let $\eta^t(y)$ be the resource share of person t when the household faces a fixed market price vector $\boldsymbol{p}$.[9] Let an assignable good (e.g., clothing) be observed for each type of person in a collective household. Let $w^t(y)$ be the *Engel curve function* for a person of type t for their assignable good. It gives the fraction of expenditure commanded by that good for a person of that type if they lived alone and faced the shadow price vector $\widetilde{\boldsymbol{p}}$ and a budget y. Since the only demander for this good is type t, the household Engel curve, $W^t(y)$, for assignable good t, evaluated at the market price vector $\boldsymbol{p}$, is given by:

$$W^t(y) = \eta^t\,(y)\,w^t\left(\eta^t\,(y)\,y/N^t\right). \tag{2}$$

The relationship (2) says that the household's Engel curves (at market prices, held fixed) for the assignable goods for $t = m, f, c$ are equal to the resource share of the relevant type times the unobserved Engel curve of a person of that type facing the shadow price vector and their shadow budget. We formally show this in Appendix 2.

In the linear reframing of the model that we develop (see section 2.2), we are able to incorporate complex households, including multi-generational, polygamous, and single-parent households. This extends DLP whose model is written for households with one adult male and one adult female and any strictly positive number of children ($N^m = N^f = 1$ and $N^c >= 1$).

Browning, Chiappori and Lewbel (2013) show that if we observed the functions $w^t(y)$ and the functions $W^t(y)$, then the resource shares $\eta^t(y)$ are identified. In general, this is possible if we observe Engel curves at many observed price vectors *and* assume that single

[9] The resource share may also depend on other covariates, including but not limited to those which affect preferences. Since the entire model can be conditioned on such variables, we suppress that dependence here.

individuals have the same preferences as individuals that live in collective households *and* that the Engel curves of single individuals are observable. In many settings, including most developing countries, at least one of these conditions is likely to be violated. For example, we do not observe children living alone and, in many countries, unmarried men and women live in collective households rather than on their own.

DLP provide sufficient restrictions on the model such that resource shares are identified from data on just Engel curve functions of collective households facing a single price vector. They do not assume that the Engel curves of single individuals are observed. They assume: a) resource shares do not depend on the household budget so that $\eta^t(y) = \eta^t$; b) individual Engel curve functions are given by the Almost Ideal demand system of Deaton and Muellbauer (1980), so that $w^t(y) = \alpha^t + \beta^t \ln y$; and c) preferences are similar—but not identical—across people, such that $\beta^t = \beta$.[10,11] Substituting these assumptions into (2) gives

$$W^t(y) = \eta^t \left(\alpha^t + \beta \left(\ln y + \ln \eta^t - \ln N^t \right) \right) \tag{3}$$

The assumption that resource shares do not depend on the household budget is strong. It implies, for example, that, all else equal, if a household gets richer the intra-household relative consumption distribution will not change.[12] Surprisingly, there is some empirical support for this restriction. Menon, Perali and Pendakur (2011) show that reported (stated preference) resource shares in Italian survey data do not vary much with household budgets. Cherchye, De Rock, Lewbel and Vermeulen (2015) use revealed preference methods to show that although resource shares do depend on variables like relative wages and education, they

[10] DLP define a property called "similar across people" (SAP) as being satisfied if the Engel curves for assignable goods are given by $w^t(y) = w^t(y/G^t) + g^t$ for some constants G^t and g^t. This condition is satisfied if preferences satisfy "shape-invariance" (see, e.g., Pendakur 1999 or Blundell, Chen and Kristensen 2014). It is also satisfied if cost functions satisfy "independence of base" (Lewbel 1989) or "equivalence-scale exactness" (Blackorby and Donaldson 1993). In addition to the assumptions of BCL, DLP assume that: $\widetilde{\boldsymbol{p}} = \boldsymbol{A}\boldsymbol{p}$, where $\boldsymbol{A}$ is diagonal with a 1 for the assignable good; resource shares do not depend on household budgets; and SAP holds. Given these assumptions, DLP show that resource shares are identified from the Engel curves of collective households at a single price vector. So, they do not require the Engel curves to be log-linear for identification. When applied to the log-linear Engel curves, SAP implies $\beta^t = \beta$. We (and they) use log-linear Engel curves to make estimation easier, not to achieve identification.

[11] DLP also define a restriction on preferences called "similar across types" (SAT) that is sufficient for identification of resource shares. We focus on their SAP restriction instead because: a) SAP is consistent with our more general A matrix, but SAT is not; and b) Given SAP, resource shares are identified with just one household type, e.g., just nuclear households with 1 child. However, we note that Bargain et al (2018) find that their (different from our) Bangladeshi data favour SAT over SAP.

[12] In fact, DLP require less than full independence, in two ways. First, they only require that resource shares are invariant to expenditure *over some range* of household expenditure. So, for example, if this invariance held only for the poorest households, we could still identify resource shares for the very poor, and consequently identify poverty at the individual level for this subpopulation. Second, the independence of resource shares from household expenditure is *conditional* on other observed covariates, which may include, for example, wealth.

do not vary much with household budgets.

The intuition for identification of resource shares in the above model is as follows. The observable budget semi-elasticity of household-level Engel curves for assignable goods, $\partial W^t(y)/\partial \ln y$, is equal to $\eta^t\beta$. Since η^t sum to 1, the sum of this semi-elasticity across types is β. Thus, we have that $\eta^t = (\partial W^t(y)/\partial \ln y) / (\Sigma_t \partial W^t(y)/\partial \ln y)$. That is, the relative magnitude of budget semi-elasticities determines resource shares. If the household's response to an increase in the budget is larger for men's clothing than for women's clothing, it is because the men's resource share is larger. Note that that it is budget responses, not levels, that identify resource shares. If women's clothing Engel curves were higher than men's, but men's had the larger budget response, then men would have the higher resource share.

A key feature of DLPs model is that identification requires $\beta \neq 0$, because if $\beta = 0$ clothing budget shares are homothetic and $\partial W^t(y)/\partial \ln y = 0$. In this case, because all slopes are zero, we cannot identify resource shares. In economic terms, this means that because we use budget responses to identify resource shares, the assignable goods must be either necessities (whose Engel curve declines with the budget) or luxuries (whose Engel curve increases with the budget).

The econometric model defined by equation (3) is nonlinear due to the fact that η^t multiplies β, and requires positive resource shares, due to the $\ln \eta^t$ term. So, estimation requires nonlinear optimization subject to bounding restrictions on parameters. This combination can be tricky to estimate, because nonlinear optimizers may have numerical issues if they search through a space with a negative value of η^t (and thus undefined $\ln \eta^t$). Further complications arise if one tries to condition the model on observed covariates, as we do below, because negative resource shares have to be avoided at all observed values of the covariates.

We now propose a reframing of DLP that uses ordinary least squares regression to identify resource shares. This strategy is easier to implement because negative resource shares don't cause the estimator to crash; instead, they are estimated, and provide evidence to the researcher that the model is not a good one for the data at hand.

2.2 OLS Estimation of Resource Shares

We now present a theory-consistent linear reframing of DLP. Consider first the case with no demographic covariates (the entire model can be written conditionally on covariates, which we do below). Rewrite equation (3) with a subscript $h = 1, ..., H$ indexing households, and an additive error term ε_h^t, as the following linear model:

$$W_h^t = a_h^t + b^t \ln y_h + \varepsilon_h^t \tag{4}$$

where

$$a_h^t = a_0^t + a_{\ln N}^t \ln N_h^t,$$

where

$$a_0^t = \eta^t \alpha^t + \eta^t \beta \ln \eta^t \qquad \text{and} \qquad a_{\ln N}^t = -\eta^t \beta,$$

and

$$b^t = \eta^t \beta. \tag{5}$$

Here, the troublesome $\ln \eta^t$ term is absorbed into the scalar parameter a_0^t, and so solves the problem of taking the log of a negative and crashing the estimator.[13] The model may be estimated by linear regression of the observed household-level assignable good expenditure share, `W_t`, on a constant, the log of the number of members of type t, `lnN_t`, and the log of the household budget, `lny`. E.g., for data on households with $t = m, f, c$, one could implement the linear seemingly unrelated regression system in Stata via:

- `sureg (W_m lnN_m lny) (W_f lnN_f lny) (W_c lnN_c lny)`

Rearranging equation (5), we have

$$\eta^t = b^t / \beta.$$

Denote the estimated regression coefficients from above as $\widehat{a}_0^t, \widehat{a}_{\ln N^t}^t$ and $\widehat{b}^t$. Since resource shares sum to 1, we can use $\sum_t \widehat{b}^t$ as an estimate of β, which implies that an estimate of the resource share of type t, η^t, is given by

$$\widehat{\eta}^t = \widehat{b}^t / \left(\sum_{t=1}^{T} \widehat{b}^t \right).$$

Notice that the estimated resource share does not depend on the estimate of the level term $\widehat{a}_0^t$. We identify resource shares in DLP solely via the relative slopes of Engel curves. Here, we see that if, for example, the household Engel curve for the men's assignable good has twice the slope (twice the value of b^t) as that for the women's assignable good, then the men have twice the resource share of the women.

One could implement this estimator for η^m in Stata via:

- `generate eta_m = [W_m]lny/([W_m]lny+[W_f]lny+[W_c]lny)`

In this model, $\beta \neq 0$ is an identifying restriction. If $\beta = 0$, then the estimated value of the denominator may be close to 0, yielding "crazy" estimates of resource shares. We use this fact to form the basis of our test of identification, described below.

[13] One could impose the restriction that $a_{\ln N}^t = -\eta^t \beta = -b^t$, but it is not necessary to do so. An easy way to do this would to replace $\ln y_h$ with $\ln y_h - \ln N_h^t$ (and drop $\ln N_h^t$).

2.3 Adding Covariates

The model above does not include any covariates, such as demographic preference shifters. Including them does not affect identification, but does require some additional notation. Let $\boldsymbol{z}$ be all variables that affect preferences and/or resource shares, including the numbers of household members of each type $\boldsymbol{N} = \{N^t\}$. Let $\widetilde{\boldsymbol{z}}$ be the subvector $\boldsymbol{z}$ of that excludes $\boldsymbol{N}$, so that $\boldsymbol{z} = [\boldsymbol{N} \quad \widetilde{\boldsymbol{z}}]$. Normalize $\widetilde{\boldsymbol{z}}$ so that $\widetilde{\boldsymbol{z}} = 0$ for some meaningful reference value of these characteristics, for example, the modal value of characteristics. Note that $\widetilde{\boldsymbol{z}}$ can be empty; in this case, only the numbers of household members of different types are preference shifters.

Assume that resource shares, η^t, and preference parameters, α^t and β, all depend on $\boldsymbol{z}$.[14] Substituting this into (3), and expanding out the terms, we have:

$$W^t(y, \boldsymbol{z}) = \eta^t(\boldsymbol{z})\alpha^t(\boldsymbol{z}) + \eta^t(\boldsymbol{z})\beta(\boldsymbol{z}) \ln y + \eta^t(\boldsymbol{z})\beta(\boldsymbol{z}) \ln \eta^t(\boldsymbol{z}) - \eta^t(\boldsymbol{z})\beta(\boldsymbol{z}) \ln N^{\mathrm{t}}. \tag{6}$$

This nonlinear structural model (6) has been implemented by several researchers on data from several countries (e.g., DLP in Malawi; Bargain, Donni and Kwenda 2014 in Cote D'Ivoire; Calvi 2014 in India; De Vreyer and Lambert 2016 in Senegal; Calvi et al 2020 in Bangladesh).[15]

As with equation (3), this model contains a term linear in the log of the resource share, $\ln \eta^t(\boldsymbol{z})$. If η^t is parameterised as a linear index (and especially if it contains an unbounded variable), then search algorithms trying to find the minimum/maximum of the sum of squares, likelihood function or GMM criterion function can stop before finding a solution. This is similar to the problem of the linear probability model giving predicted probabilities outside $[0, 1]$, but with the additional consquence that it may induce numerical problems in nonlinear solvers. For example, they may try to evaluate the function in a region of the parameter space where $\eta^t(\boldsymbol{z})$ is negative, yielding a missing value for $\ln \eta^t(\boldsymbol{z})$. Alternatively, η^t may be parameterised as a bounded function of z, but then the behaviour of the function near the boundaries may present problems for nonlinear solvers. An additional problem relevant to equation (6) comes from the fact that the term $\eta^t(\boldsymbol{z})\beta(\boldsymbol{z}) \ln y$ has quadratic interactions in $\boldsymbol{z}$ multiplying $\ln y$. These make it difficult to precisely identify the dependence of resource shares $\eta^t(\boldsymbol{z})$ on $\boldsymbol{z}$, because $\boldsymbol{z}$ affects both η^t and β.

[14] *Distribution factors* are variables that affect resource shares, η^t, but not preferences, α^t and β. We can think of these as elements of z, imposing the restriction that they do not affect α^t and β. These restrictions do not affect the linear estimator for resource shares that we propose below. So, in this paper, we don't separately track distribution factors from other observed covariates that affect preferences and/or resource shares.

[15] Menon, Perali and Piccoli (2018) provide a different almost linear formulation of Engel curves that could be approximated similarly to our development below.

/*Modern statistical software (especially, e.g., `R` and `Matlab`) allows users to write their own code to estimate nonlinear models, bound parameter spaces, and tell optimizers what to do when parameters hit boundaries. So, the nonlinear estimation issues outlined above are quite solvable with currently available computing resources. However, our objective is to provide OLS-based estimators for household models, so that more researchers see these models as practicable in their empirical settings.

Rewrite equation (6) with a subscript h on all observed variables, and an additive error term ε_h^t, as the following linear model:

$$W_h^t = a_h^t + b_h^t \ln y_h + \varepsilon_h^t, \tag{7}$$

where

$$a_h^t = \eta^t(\boldsymbol{z}_h)\alpha^t(\boldsymbol{z}_h) + \eta^t(\boldsymbol{z}_h)\beta(\boldsymbol{z}_h)\ln\eta^t(\boldsymbol{z}_h) - \eta^t(\boldsymbol{z}_h)\beta(\boldsymbol{z}_h)\ln N_h^{\mathrm{t}}, \tag{8}$$

and

$$b_h^t = \eta^t(\boldsymbol{z}_h)\beta(\boldsymbol{z}_h). \tag{9}$$

Here, a_h^t and b_h^t are functions of the vector of conditioning variables $\boldsymbol{z}_h$. Suppose that η^t, α^t and β are linear indices in $\boldsymbol{z}_h$. Then, a_h^t is a third-order function in $\boldsymbol{z}_h$, and b_h^t is quadratic in $\boldsymbol{z}_h$. Defining $\boldsymbol{Z}_h$ as the list of level and interaction terms up to the third order in $\boldsymbol{z}_h$, OLS regression of W_h^t on a constant, $\mathbf{Z}_h, \ln y$ and $\boldsymbol{Z}_h \cdot \ln y$ would suffice.

Alternatively, η^t, α^t and β could have unknown functional forms. In this case, one could let a_h^t and b_h^t be nonparametric functions of $\boldsymbol{z}_h$, and use standard semiparametric methods to estimate the model. One such approach would be to let a_h^t and b_h^t be multivariate polynomials over $\boldsymbol{z}_h$, with the degree of the polynomials increasing with the sample size.

2.4 Approximation

Unfortunately, neither of these approaches is practical with a high-dimensional conditioning vector $\boldsymbol{z}_h$. For example, with a constant and 9 conditioning variables in $\boldsymbol{z}_h$, third-order interactions requires 444 regressors.[16] So, we recommend approximating the model. Approximate the a_h^t term with

$$a_h^t = a_0^t + a_{\ln N^t}^t \ln N_h^t + \boldsymbol{a}_z^{t\prime}\boldsymbol{z}_h. \tag{10}$$

Similarly, we recommend approximating the slope term analogously to the level term as

$$b_h^t = b_0^t + \boldsymbol{b}_z^{t\prime}\boldsymbol{z}_h, \tag{11}$$

[16] $10^3 = 1000$ triples, deleting permutations, is 222 unique combinations, times 2 for the intercept and slope.

and let $\boldsymbol{b}_N^t$ and $\boldsymbol{b}_{\tilde{z}}^t$ refer to the relevant subvectors of $\boldsymbol{b}_{\boldsymbol{z}}^t$. From inspection of equation (9), it is easy to see that this approximation for b_h^t is exact if η^t is linear in $\boldsymbol{z}_h$ and β is independent of $\boldsymbol{z}_h$ (that is, if β is a constant).

This approximate model may be estimated via OLS regression of W_h^t on a constant, $\ln N_h^t$, $z_h, \ln y$ and $\boldsymbol{z}_h \cdot \ln y$. The estimated coefficients on $\ln y$ and $\boldsymbol{z}_h \cdot \ln y$ are estimates of b_0^t and b_z^t, respectively. These may be used to construct an estimate $\widehat{b}_h^t$ of b_h^t:

$$\widehat{b}_h^t = \widehat{b}_0^t + \widehat{\boldsymbol{b}}_z^{t\prime} \boldsymbol{z}_h.$$

Regardless of the specification of b_h^t, and regardless of whether not it is taken to be an approximation or exact (due to prior knowledge of the functional form of η^t and β), we can solve for resource shares. Since resource shares sum to 1, we can use $\sum_t \widehat{b}_h^t$ as an estimate of $\beta(\boldsymbol{z}_h)$, which implies that an estimate of the resource share of type t in a household with characteristics $\boldsymbol{z}_h$ is given by

$$\widehat{\eta}_h^t = \widehat{\eta}^t(\boldsymbol{z}_h) = \widehat{b}_h^t / \left(\sum_{t=1}^{T} \widehat{b}_h^t \right). \tag{12}$$

Engel curves may be estimated by equation-by-equation ordinary least squares (OLS) or with linear seemingly unrelated regression (SUR).[17] Resource shares may then be computed via (12).

Suppose we have a dataset on childless couples, so that $T = m, f$, $N^m = 1$, $N^f = 1$, $\ln N^m = 0$ and $\ln N^f = 0$. Let the data be log budgets `lny`, budget shares `W_t`, a covariate `z,` and the interaction (products) of log budgets and the covariate `lny_z`. Since N^t and $\ln N^t$ are constants, they are not included as regressors. The following two lines of Stata code implement our estimate of the man's resource share, as a function of the covariate `z`:

- `sureg (W_m z lny lny_z) (W_f z lny lny_z)`
- `generate eta_m = ([W_m]lny+z*[W_m]lny_z) / ([W_m]lny+z*[W_m]lny_z+[W_f]lny+z*[W_f]lny_z)`

Here, the first line estimates the model, and the second line delivers the resource share of the man in each household.

From a practical standpoint, if the denominators in (12) had a lot of variation, or if they were close to zero, estimated resource shares might be somewhat wild. However, we can simplify the denominator by imposing the linear restrictions

[17] Although equation-by-equation OLS and SUR estimators are both consistent, they are not identical, because the regressor lists differ across equations ($\ln N_h^t$ only shows up in the regressor list for W_h^t). So, we recommend SUR.

$$\sum_t \boldsymbol{b}^t_{\tilde{z}} = \mathbf{0}. \tag{13}$$

implying that $\sum_t b^t_h = \sum_t \left(b^t_0 + b^t_{N^m} N^m_h + b^t_{N^w} N^w_h + b^t_{N^c} N^c_h\right)$.

Then, estimated resource shares are equal to

$$\widehat{\eta}^t(\boldsymbol{z}_h) = \frac{\widehat{b}^t_0 + \widehat{\boldsymbol{b}}^{t\prime}_z \boldsymbol{z}_h}{\sum_t \left(\widehat{b}^t_0 + \widehat{b}^t_{N^m} N^m_h + \widehat{b}^t_{N^w} N^w_h + \widehat{b}^t_{N^c} N^c_h\right)}.$$

Here, we expect $b^t_{N^t}$ to all have the same sign, and that the variation in the denominator would be tamped down. In our empirical work below, we impose this restriction.

This functional form for resource shares allows for the possibility that the resource shares of each type equals their per-capita share household members. In particular, if $b^t_0 = 0$ for all t, $\boldsymbol{b}^t_{\tilde{z}} = \mathbf{0}$ for all t, $b^t_{N^{t'}} = 0$ for all $t' \neq t$ and $b^t_{N^t} = \kappa$ for all t, then we get per-capita resource shares, $\eta^t(\boldsymbol{z}_h) = N^t\kappa / \sum_t N^t\kappa = N^t / \sum_t N^t$.[18] In our empirical work below, we test this restriction.

Let *composition* be a variable indicating whether or not different types of people are present in a household. In our estimation below, we consider 4 compositions of types: households with men, women and children; households with men and children only; households with women and children only; and households with men and women only. A pooled estimator would simply interact composition with all the regressors in the model ($\boldsymbol{z}$, $\ln y$ and $\boldsymbol{z} \cdot \ln y$); alternatively, one could estimate the model separately for each composition. In our empirical work, we do the latter. That is, to compute resource shares for people living in households with men, women and children, we run regressions on observations with at least 1 man, 1 woman and 1 child in each household. To compute resource shares for people living in households with just women and children, we run regressions on observations with no men, and at least 1 woman and 1 child. And analogously for the other 2 compositions. All test statistics, for example the Wald test of the per-capita model, are simply sums of chi-square test statistics across the samples for each composition.

2.5 A Linear Test of Model Identification

As noted above, if $\beta(\boldsymbol{z}_h) = 0$ then resource shares are not identified. In this case, the estimated value of the denominator may be close to 0, and the resulting estimated resource shares would be unreliable. If it were the case in the limit, then inference is polluted by

[18] One could additionally restrict $\sum_t \boldsymbol{b}^t_N = \mathbf{0}$, implying that $\widehat{\eta}^t(\boldsymbol{z}_h) = b^t_h / \left(\sum_t b^t_0\right)$. This further simplifies the denominator, but at the cost of not nesting the per-capita model.

weak identification problems (see Han and McCloskey 2019). Consequently, it is valuable to have a test of identification to tell us whether or not these methods will work at all. Previous papers (DLP; Dunbar, Lewbel and Pendakur 2019; Han and McCloskey 2019) have tested this identifying restriction, but their tests all involve estimating nonlinear models. Our linear reframing of DLP straightforwardly delivers an OLS-based test of whether or not this identifying restriction for resource shares is supported by the data.

Let the *overall assignable Engel curve* of the household be given by $W_h = \sum_t W_h^t$, and let $a_h = \sum_t a_h^t$, $b_h = \sum_t b_h^t$ and $\varepsilon_h = \sum_t \varepsilon_h^t$. If W_h^t is the fraction of the household spent on clothing for members of type t, then W_h is the fraction of the household budget spent on clothing in total for all members. Notice that a_h depends on $\ln \boldsymbol{N}_h$, the vector of logs of numbers of members. Then, our approximate model above implies

$$W_h = a_h + b_h \ln y_h + \varepsilon_h \tag{14}$$

and OLS regression of W_h on $1, \ln \boldsymbol{N}_h, \boldsymbol{z}_h, \ln y_h$ and $\boldsymbol{z}_h \ln y_h$ yields an estimate $\widehat{b}_h$ of $\beta(\boldsymbol{z}_h)$. We propose that an easy and useful test of identification for this model is to test whether overall assignable good Engel curves are statistically significantly upward or downward sloping, that is, test whether or not b_h is zero.

Below, we use two results from our overall assignable goods Engel curve regression to consider whether our methods should be applied to the data at hand. First, we use $E[\widehat{b}_h] = \widehat{b}_0 + \widehat{\boldsymbol{b}}'\overline{\boldsymbol{z}}_h$, where $\overline{\boldsymbol{z}}_h$ is the sample average of $\boldsymbol{z}_h$, as a test statistic. This is a test of the economic hypothesis that the overall assignable good Engel curve, evaluated at the mean value $\boldsymbol{z}_h$, is either a necessity or a luxury (is increasing or decreasing). If it is neither, then our strategy to estimate resource shares should not be used. Second, for every observation in the data, we test whether or not $\widehat{b}_h = \widehat{b}_0 + \widehat{\boldsymbol{b}}'\boldsymbol{z}_h$ is statistically significantly different from zero, and report the fraction of households for which it is statistically significant. Here, we think that a “large” fraction of households should have an estimated overall Engel curve that is either upward or downward sloping, where “large” is taken to be 75% of the sample (other cutoffs could be used).

3 Data

In most countries in the world, national statistical offices regularly collect household expenditure survey data. These data are used as input in national accounts, for the calculation of the GDP, to measure inflation, to analyse household spending patterns and behaviour, and to evaluate policy. Since the early 1980s, the World Bank has been providing assistance to national statistical offices in the design and implementation of household surveys through the Living Standards Measurement Study (LSMS). These data are standardised to some extent, and are the best tool available for cross country comparisons of poverty in low- and middle-income countries.

LSMS surveys exist for about 40 countries, and often several waves exist. There are in total 87 country-waves potentially available for the analysis of household consumption behaviour. We analyse the most recent waves from 12 countries for which LSMS data include clothing expenditure by type of individual (men, women and children), a measure of total expenditure for the household, and a minimal set of demographic variables (age, sex and education level of household members). We also include non-LSMS data from the Bangladesh Integrated Household Survey so that we can consider using food, both purchased and home produced, as the assignable good (see below). [19,20]

Table 2: Descriptive Statistics

Country	Total H	Single H	compositions	Our H	Nuclear H	budget	std dev
Albania	3599	239	mw, mwc	3279	612	11084	6477
Bangladesh	6434	219	mw, wc, mwc	6120	2122	6416	6268
Bulgaria	3018	801	mw, mwc	2099	412	13117	7954
Ethiopia	4717	503	mw, wc, mwc	3845	1481	3092	3645
Ghana	8687	1922	mw, mc, wc, mwc	6313	2195	5096	4835
Iraq	17513	288	mw, wc, mwc	14297	5487	26188	14287
Malawi	12271	1030	mw, wc, mwc	10873	5488	3189	3758
Nigeria	4600	349	mw, wc, mwc	3556	1013	6656	20322
Tajikistan	1503	54	mw, mwc	1275	192	10483	6250
Tanzania	3352	320	mw, wc, mwc	2677	1133	7219	5164
Timor Leste	4477	229	mw, wc, mwc	3788	1577	4954	4116
Uganda	3117	257	mw, wc, mwc	2468	1014	2462	2262
Bangladesh–Food	6434	219	mw, wc, mwc	5604	1916	6445	6287

Descriptive statistics for the sample of countries are in table 2. Altogether, these countries represent about 9% of the world population. Starting from the publicly available LSMS data (and the Bangladesh data), we exclude observations with missing data on clothing expenditures, total household expenditures or the age, sex and education level of household members. This yields sample sizes reported in column *Total H*. There is a wide range of sample sizes after this initial cleaning, from 1,503 households in Tajikistan to 17,513

[19] A variety of reasons makes the data from the other countries unusable. In some cases, no data on assignable goods is collected; in others, information on elements of non durable expenditure is missing.

[20] Code to go from publicly available online data to our working data files for each of the 12 countries, and code to estimate all tables, is available on request.

households in Iraq. Below we will pay attention to whether sample size matters to the feasibility of the method.

In column *Single H*, we report the number of households which are composed of a single adult man or woman. Since these households only have one individual, there is no sharing of resources, and they are not used in the estimation of resource shares, but they are included in the subsequent poverty analysis. It is worth noting that there are few singles, and that most households contain more than one type of person, highlighting the importance of modeling the within-household allocation of resources.[21]

For the estimation of the resource shares, we use all household compositions apart from households with a single type of individual, whatever their number (that is, we exclude households composed of men only, whatever their number, and similarly for women), and we allow for any number of individuals of each type. The possible compositions are mw, mwc, wc, and mc. These indicate that individuals of the type m for men, w for women and c for children, are present in the households, but it does not indicate how many individuals of each type there are. We exclude households belonging to a composition for which there are less than 100 observations (since estimation is done separately for each composition). The compositions remaining in the sample after this selection are indicated in column *compositions*, and column *Our H* gives the total number of observations of these compositions. This latter column shows that we are able to exploit most of the data.

Column *Nuclear H* shows the number of nuclear households in each country. In contrast to much of the previous work on resource shares, we are not limited to using only nuclear households. This shows that the selection to just nuclear households can be very restrictive indeed in some countries; nuclear households are less than 25% of all households in 6 of our 12 countries.

We then provide the mean and standard deviation in our sample (excluding singles) of the overall budget in (PPP) $US 2010. In some countries in our data, the average household budget is close to the World Bank poverty line of $US7.60 per day for a 4 person household (e.g., Ethiopia, Malawi and Uganda); in some countries, it is well above (e.g., Bangladesh, Iraq). In all countries, the standard deviation is of comparable order to the mean, which is desirable since identification rests on budget variation.

The bottom row gives summary statistics for the data on assignable food in Bangladesh. It is different from the clothing data because different observations have valid assignable food data versus clothing data. In the analysis below, we will compare estimated resource shares from assignable food with those from assignable clothing.

[21] For households with, e.g., multiple men but no women or children, the underlying model could be collective but it could only be estimated if there were an observed assignable good for each of the men.

4 Results

We estimate equations (7), (10) and (11) under the restrictions (13) via seemingly unrelated regression in Stata. Our observed vector of demographic variables $\boldsymbol{z}_h$ is comprised of: the numbers of men, women and children ($\boldsymbol{N}_h$); the average ages of men, women and children; the minimum age of the children; the average education levels of the men and women; and a dummy variable indicating that the household lives in an urban area.[22]

4.1 Test of Identification

The statistical significance of the slope of the Engel curve for the sum of household assignable goods provides a test of the applicability of the method. In Table 3, we give the mean and standard deviation of assignable goods budget shares (summed across household members), and the slope of the Engel curve evaluated at average characteristics, along with a t-test for its difference from zero. In the rightmost column, we give the fraction of observations whose estimated slope (conditional on their observed covariates) is statistically significantly different from zero.

Clothing is not a large budget share. Clothing represents between 1.7% and 7% of the budget. The standard deviation of clothing shares is high relative to the mean, so there is considerable dispersion in the distribution of clothing shares in each country.

Table 3: Test of Identification

country	sample N	budget share	std dev	slope at $\overline{\boldsymbol{z}}$	t-test of slope	% of sample significant
Albania	3279	0.041	0.042	0.014	4.7	84
Bangladesh	6120	0.039	0.021	-0.016	-21.4	100
Bulgaria	2099	0.036	0.040	0.014	5.2	90
Ethiopia	3845	0.072	0.064	-0.011	-3.5	63
Ghana	6313	0.048	0.040	-0.002	-1.0	63
Iraq	14297	0.07	0.047	0.021	14.8	99
Malawi	10873	0.025	0.036	0.009	10.0	98
Nigeria	3556	0.017	0.023	-0.002	-2.0	51
Tajikistan	1275	0.058	0.050	0.008	1.9	12
Tanzania	2677	0.044	0.058	-0.002	-0.9	14
Timor Leste	3788	0.022	0.020	-0.002	-1.8	24
Uganda	2468	0.055	0.052	-0.004	-1.1	5
Bangladesh–Food	5604	0.568	0.150	-0.120	-17.2	100

[22] The Bangladeshi data are not drawn from a nationally representative sample frame; rather these data are representative of rural households only. So, we do not include the urban dummy in the demographic shifter list for Bangladeshi estimates.

Clothing is found to be a luxury in Albania, Bulgaria, Iraq, and Malawi and a necessity in Bangladesh, Ethiopia, Nigeria.

The slopes of the clothing Engel curves are not statistically significantly different from zero in Ghana, Tajikistan, Tanzania, Timor Leste and Uganda. Since the formula for resource shares uses this slope as a denominator, for these countries, the model may not be identified.

We also report the percentage of the sample for which the slope is significant. For our method to work, this needs to be high enough, so that we further eliminate Ethiopia and Nigeria because less than 75% of observations in those countries have predicted Engel curve slopes that are statistically significantly different from zero. This leaves us with 5 countries which pass the test, hence for which the model is identified and resource shares can be estimated.

For Bangladesh, we also have assignable data on food consumption. We have fewer observations (5604) on food than clothing (6120) because there is some non-response in the daily food diary data. Food budget shares are much larger than clothing budget shares: whereas clothing accounts for only 3 per cent of total household consumption, fully 56 per cent of household consumption in our Bangladesh sample is food.

A long history of demand analysis, dating back to Engel (1890), has shown that food is a necessity whose Engel curve is therefore downward sloping. The Bangladeshi data reflect this with a strongly declining food Engel curve, whose estimated slope with respect to the log of household expenditure is -0.12, with a t-test of -17, and 100 per cent of the sample with significant slopes. Food Engel curves are therefore different from clothing Engel curves in two important ways: 1) Food budget shares are large while clothing budget shares are small; and 2) Food Engel curves slope downwards while clothing Engel curves sometimes slope upwards, sometimes slope downwards, and are sometimes flat. Both of these differences suggest that food is a preferable candidate for our methods. Consequently, in our analysis below, we pay special attention to the difference—or lack thereof—between estimates of Bangladeshi resource shares based on clothing versus food Engel curves.

4.2 Resource shares

Estimated per-person resource shares, η_h^t/N_h^t, of men, women and children, are shown in Table 4, for the countries whose data pass our test of identification. We report both the resource shares estimated at the mean of observed covariates, $\overline{\boldsymbol{z}}$, and the mean of the resource shares evaluated at all $\boldsymbol{z}_h$. For the former, we give the standard error and for the latter, the standard deviation.

In Albania, the estimated men's and women's per-person resource shares at the average

$\boldsymbol{z}_h$ are 28 per cent and 24 per cent, respectively, with small standard errors, of 3 per cent. Because resource shares are nonlinear functions of estimated OLS regression coefficients, the estimate of resource shares at average $\boldsymbol{z}_h$ does not equal the average of estimated resource shares over all $\boldsymbol{z}_h$. However, they are similar: the sample averages of the resource shares are 28 and 25 per cent, respectively, for men and women. Variation in estimated resource shares is driven by variation in observed covariates $\boldsymbol{z}_h$. The standard deviation of these estimated resource shares are 37 and 34 per cent, indicating quite a lot of heterogeneity in resource shares driven by the sample variation in observed covariates.

Table 4: Predicted Resource Shares, Selected Countries

Country	sample size	Evaluated at $\overline{\boldsymbol{z}}$ men est *std err*	women est *std err*	children est *std err*	Evaluated at all $\boldsymbol{z}_h$ men mean *std dev*	women mean *std dev*	children mean *std dev*	η outside [0,1]	per cap test Wald, df *p-value*
Albania	3279	0.282	0.247	0.134	0.28	0.256	0.126	0.062	45, 35
		0.032	*0.033*	*0.030*	*0.369*	*0.340*	*0.166*		*0.129*
Bangladesh	6120	0.312	0.286	0.120	0.311	0.284	0.122	0	387, 41
		0.011	*0.014*	*0.010*	*0.114*	*0.118*	*0.059*		*0.000*
Bulgaria	2099	0.304	0.372	0.188	0.292	0.387	0.173	0.079	49, 35
		0.038	*0.041*	*0.061*	*0.14*	*0.218*	*0.214*		*0.058*
Iraq	14297	0.268	0.237	0.041	0.267	0.236	0.042	0.011	530, 45
		0.009	*0.011*	*0.006*	*0.133*	*0.131*	*0.069*		*0.000*
Malawi	10873	0.312	0.274	0.124	0.31	0.267	0.127	0.015	267, 45
		0.028	*0.03*	*0.011*	*0.179*	*0.154*	*0.089*		*0.000*
Bangladesh	5604	0.300	0.241	0.173	0.304	0.235	0.174	0.028	232, 41
Food		*0.013*	*0.015*	*0.010*	*0.116*	*0.108*	*0.078*		*0.000*

The rightmost column of table 4 gives the fraction of resource shares which fall outside of the [0,1] interval. The largest is Bulgaria (clothing) with 7%. The consequence of estimating our model on data where Engel curves are not very steep, that is, where the assignable good is neither very strongly a necessity or a luxury, could be large fraction of estimated shares outside $[0, 1]$ and implausible estimates for resource shares.[23]

According to the point estimates, men get a larger share of household resources than women in all countries, except Bulgaria. Children get between 12% and 18% everywhere, except in Iraq where they get about 4% of resources each.

A standard resource share in current use by the World Bank and other agencies is the per-capita share of household members, that is, $\eta_h^t = N_h^t / \sum_s N_h^t$. This would assign each

[23]Estimates for all countries, even those where data do not pass the test of identification, are available on request.

person their per-capita share of household consumption. Given our model, this obtains if $\boldsymbol{b}^t_{\tilde{z}} = \mathbf{0}$ for all t and $b^t_{n^{t'}} = 0$ for all $t' \neq t$ and $b^t_{n^t} = \kappa$ for all t. The Wald test statistic for this hypothesis and its associated degrees of freedom are presented in the rightmost column of Table 4, with p-values in *italics* below.

Table 4 shows lots of inequality across household members, so it should not be surprising that the per-capita model is not supported by these estimates in most countries. The per-capita model is rejected in data from Iraq, Malawi and Bangladesh (for both clothing and food), but it is not rejected in Albania or Bulgaria. Notably, these latter two countries have the smallest samples by a factor of about 2. This suggests to us that rather large sample sizes are needed to estimate these models.[24]

We will show below that the failure of the per-capita model implies the existence of both gender gaps in consumption and gendered poverty that has been missed by previous investigations.

4.3 Gender gaps

In Table 4, we see some evidence that women get smaller per-person resource shares than men. However, those estimates include all types of households, including those that don't have an adult man or those that don't have an adult women. To construct an estimated gender gap that refers strictly to within-household inequality, we present in Table 5 estimates on the subset of households that include both adult men and adult women. In the leftmost columns, we present the mean and standard deviation of estimated resource shares evaluated at all values of the covariates. In the right-hand columns, we present estimated resource shares, and their standard errors, for men and women evaluated at the average value of observed covariates. The difference between these two per-person resource shares is our gender-gap estimate, provided with standard errors, and 1, 2 or 3 stars to indicate statistical significance at the 10, 5 and 1 per cent level.

Here, we see that the evidence given in Table 4 that women have a greater share of household resources than men in Bulgaria is not a statistically significant finding. Because the estimates of men's and women's resource shares covary, the estimated 6.8 percentage point gender gap has a large standard error of 7.0 percentage points, even though the estimated resource shares of men and women have standard errors of only around 4 percentage points. Consequently, the difference between them—the gender gap—is statistically indistinguishable from zero.

[24] We note that one can pool multiple waves of data for a given country, just by including a year dummy as an additional element of $\boldsymbol{z}$. To do this, the model should include the additional restriction that the function β is price-independent (as in Deaton and Muellbauer 1980, but not as in Muellbauer 1974,1975).

Table 5, Estimated Resource Shares and Gender Gaps, Selected Countries
Households with Both Men and Women Present

		Evaluated at all $\boldsymbol{z}_h$		Evaluated at $\overline{\boldsymbol{z}}$		Gender Gap at $\overline{\boldsymbol{z}}$	
	sample H	men	women	men	women		
		mean	mean	est	est	est	sig
		std dev	*std dev*	*std err*	*std err*	*std err*	
Albania	3279	0.333	0.287	0.282	0.247	0.035	
		0.298	*0.259*	*0.032*	*0.033*	*0.059*	
Bangladesh	5427	0.348	0.302	0.312	0.267	0.045	**
		0.111	*0.090*	*0.011*	*0.011*	*0.020*	
Bulgaria	2099	0.312	0.447	0.304	0.372	-0.068	
		0.150	*0.212*	*0.038*	*0.041*	*0.070*	
Iraq	14040	0.333	0.286	0.268	0.233	0.035	***
		0.139	*0.121*	*0.009*	*0.009*	*0.017*	
Malawi	9490	0.362	0.279	0.312	0.253	0.059	
		0.167	*0.135*	*0.028*	*0.029*	*0.054*	
Bangladesh	4958	0.334	0.240	0.300	0.224	0.076	***
Food		*0.119*	*0.096*	*0.013*	*0.013*	*0.024*	

The point estimates of the gender gap in Albania and Malawi are positive (3.5 and 5.9 percentage points, respectively), but are statistically insignificantly different from zero. In fact, we only see a statistically significant gender gap in Bangladesh and Iraq, and both of these show larger resource shares for men. The Iraqi data suggest a gender gap of 3.5 percentage points. In the Bangladeshi data, the estimated gender gap from assignable clothing data is larger, about 4.5 percentage points, and from the assignable food data, about 7.6 percentage points. The similarity between the estimates coming from clothing data and food data is striking (and, they are not statistically significantly different from each other).

4.4 Individual poverty

A standard poverty line used by the World Bank and other international organizations concerned with poverty is a fixed number of US dollars per day (using PPP adjusted values) per person. We call this the *per capita* approach. The threshold typically differs between richer and poorer countries: we use the World Bank poverty lines[25] of $1.90 a day (low income poverty line) for Malawi, $3.20 a day (lower middle-income poverty line) for Bangladesh and $5.50 a day (upper middle-income poverty line) for Albania, Bulgaria and Iraq.

The per capita approach does not account for scale economies in consumption (see Appendix 2 for more on this). Because our measure of resource shares simply divides the pie differently than the per-capita approach, our poverty rates are directly comparable to the

[25] Posted at https://www.worldbank.org/en/topic/poverty/overview

per-capita poverty rates. Our poverty rates differ from, and are larger than, the per capita rates for two reasons: first, our approach allows for inequality within households, so that some members may be poor even if the household budget exceeds the per-capita threshold; and, second, our approach allows for heterogeneity across households, and this heterogeneity may push some members of some households below the poverty line.

Table 6, Estimated Poverty Rates, Selected Countries

country	per-capita	men	women	children	all people
	est	est	est	est	est
	std err	*std err*	*std err*	*std err*	*std err*
Albania	0.319	0.282	0.392	0.471	0.369
	0.008	*0.051*	*0.054*	*0.082*	*0.021*
Bangladesh	0.464	0.252	0.338	0.631	0.421
	0.006	*0.026*	*0.022*	*0.034*	*0.008*
Bulgaria	0.080	0.185	0.157	0.416	0.206
	0.005	*0.058*	*0.035*	*0.11*	*0.031*
Iraq	0.070	0.013	0.030	0.744	0.330
	0.002	*0.005*	*0.010*	*0.042*	*0.017*
Malawi	0.629	0.469	0.586	0.727	0.627
	0.004	*0.025*	*0.035*	*0.03*	*0.009*
Bangladesh	0.453	0.283	0.489	0.334	0.373
Food	*0.006*	*0.026*	*0.027*	*0.036*	*0.009*

In Table 6, we measure poverty using the per-capita approach and our measures of resource shares. In the leftmost column, we compute for each person in the household their per-capita budget, y_h/N_h, compare this to the poverty line, and report the poverty rate. In the middle three columns, we compute for each man, woman and child in the dataset, $y_h\eta_h^t/N_h^t$, compare this to the poverty line, and report the poverty rate. Like DLP, and we use a poverty line 40% lower for children. In the final column, we report the overall poverty rate, at the person-level and using our resource shares, for the entire sample. Note that for these estimates, we include single-member households, where $N_h^t = \eta_h^t = 1$, and households with just one type of person (e.g., a two-man household), where each of the N_h^t people is assigned y_h/N_h^t. We provide asymptotic standard errors, computed via the bootstrap.[26]

The key message from Table 6 is that the variation across types in resource shares that we observed in Tables 4 and 5 translates directly into variation in estimated poverty rates across types. The point-estimates of the gender gap in resource shares are largest for Bangladesh, Malawi and Iraq. In these countries, we see higher women's poverty than men's poverty. In

[26] We bootstrap the standard errors (rather than using the delta method) because poverty rates are a discontinuous function of the estimated resource shares, which are themselves nonlinear functions of estimated OLS regression coefficients. For an alternative, see Wouterson and Ham (2013).

Bangladesh, women are 8 percentage points more likely to be poor than men; in Malawi, they are 12 percentage points more likely to be poor than men.

5 Discussion

We provide evidence of substantial within-household consumption inequality. This suggests that current standard practice for poverty measurement in developing countries—asking whether or not per-capita household income falls below a threshold—can be misleading. This current practice ignores within-household inequality, and so mischaracterises poverty rates. For example, if a household has income slightly above the poverty line, then by the per-capita method we would call it non-poor, but even a small amount of within-household inequality will result in some of its members being poor. Further, within-household inequality may be biased against certain groups. Among the 5 countries for which we estimate resource shares, we see statistically significant gender gaps in resource shares that favour men over women in two countries and we see no statistically significant evidence of gender gaps that favour women. Further, these gaps in resource shares result in gender gaps in poverty rates.

If within-household inequality is real, and affects the incidence of poverty among men, women and children, then its accurate measurement is of paramount importance. Our work suggests that statistical agencies, and the World Bank programs they work with, should focus more data gathering effort on assignable goods. There are two strategies available here. First, resources could be directed to gathering assignable person-level consumption flows for all categories of goods and services (aka: dream data in Table 1). With these data, we would not need a structural model such as ours to estimate resource shares—we could measure them directly. Second, resources could be directed to gathering assignable consumption flows for 1 or 2 categories of goods and services that can be measured well and which represent a large fraction of total household expenditure. With these data, we could estimate resource shares using our structural model (or any household model that bases identification on assignable goods, see, e.g., Bargain et al 2020).

This recommendation applies similarly to field experimentalists where a potential outcome variable is individual poverty or consumption. If information on total household consumption is already being gathered, this may only require adding a few questions to a questionnaire.

Our estimates of resource shares, gender gaps and poverty rates for Bangladesh come from two different assignable goods. We use clothing, which is roughly 4 per cent of the household budget, and food, which is roughly 56 per cent of the household budget. Clothing has a venerable history as an assignable good used in this literature (e.g., survey of Donni and Molina 2018; Calvi 2019; etc). However, the use of clothing is due to its availability in public-use datasets, not to its superiority in other ways.

We find in our work that using food data as an assignable good to identify resource shares delivers estimates that are very similar to those generated from clothing data. But, food data have five advantages over clothing data. First, food is more plausibly assignable than is clothing. Clothing can be handed down from member to member, but the same food cannot be eaten by two members. Second, for food, we often collect data on both quantity and expenditure, whereas for clothing, we usually only know expenditure. There may thus be more unobserved price heterogeneity in clothing than in food. Third, food budget shares are known to be downward sloping (e.g., Engel 1857, 1895), and therefore satisfy the identifying restriction of our model. Fourth, clothing is much more durable than food. Consequently, observed clothing expenditure may not equal clothing consumption, due to infrequency of purchase. Fifth, food shares are typically much larger than clothing shares. This is not a gain in terms of the model in any formal sense, but it does seem like a worthwhile auxilliary feature. All together, this suggests that statistical agencies and the World Bank should focus significant data gathering resources on the collection of person-level food consumption.

6 Conclusions

We show how to estimate the resource share of each person in a collective household via simple linear regressions of assignable goods Engel curves. This may be implemented with off-the-shelf consumer expenditure micro-data, such as that collected through the World Bank's Living Standards Measurement Study. We apply the model to data from 12 countries, and investigate resource shares, gender gaps and individual poverty. We find that equal sharing—the implicit assumption underlying household-level poverty calculations—is rejected. We also find evidence of large gender gaps in resource shares, and consequently in poverty rates, in a few countries.

7 References

Anderson, S. and Ray, D., 2010. Missing women: age and disease. The Review of Economic Studies, 77(4), pp.1262-1300.

Bargain, O., Donni, O. and Kwenda, P., 2014. Intrahousehold distribution and poverty: Evidence from Cote d'Ivoire. Journal of Development Economics, 107, pp.262-276.

Bargain, O., Lacroix, G. and Tiberti, L., 2020. Validating the collective model of household consumption using direct evidence on sharing. Partnership for Economic Policy Working

Paper, (2018-06).

Blundell, R., Chen, X. and Kristensen, D., 2007. Semi-nonparametric IV estimation of shape-invariant Engel curves. Econometrica, 75(6), pp.1613-1669.

Brown, C. M. Ravallion and D. van de Walle. 2019. Are poor Individuals Mainly Found in Poor households?, forthcoming, the Review of Economics and Statistics

Browning, M., Chiappori, P.A. and Lewbel, A., 2013. Estimating consumption economies of scale, adult equivalence scales, and household bargaining power. Review of Economic Studies, 80(4), pp.1267-1303.

Bütikofer, A., 2011. Sibling sex composition and cost of children. Unpublished manuscript, Department of Economics, Norwegian School of Economics, Bergen, Norway.

Calvi, R., 2019. Why are older women missing in india? the age profile of bargaining power and poverty. forthcoming, Journal of Political Economy.

Calvi, R., Lewbel, A. and Tommasi, D., 2017. LATE with Mismeasured or Misspecified Treatment: An Application to Women's Empowerment in India.

Calvi, R., Penglase, J., Tommasi, D. and Wolf, A. 2020. The More the Poorer? Resource Sharing and Scale Economies in Large Families. Unpublished working paper.

Campbell, F., Conti, G., Heckman, J.J., Moon, S.H., Pinto, R., Pungello, E. and Pan, Y., 2014. Early childhood investments substantially boost adult health. Science, 343(6178), pp.1478-1485.

Cherchye, L., De Rock, B. and Vermeulen, F., 2007. The collective model of household consumption: a nonparametric characterization. Econometrica, 75(2), pp.553-574.

Cherchye, L., De Rock, B., Lewbel, A. and Vermeulen, F., 2015. Sharing rule identification for general collective consumption models. Econometrica, 83(5), pp.2001-2041.

Cherchye, L., Demuynck, T., De Rock, B. and Vermeulen, F., 2017. Household consumption when the marriage is stable. American Economic Review, 107(6), pp.1507-34.

Chiappori, P.A., 1992. Collective labor supply and welfare. Journal of political Economy, 100(3), pp.437-467.

Chiappori, P.A. and Ekeland, I., 2009. The Microeconomics of Efficient Group Behavior: Identification 1. Econometrica, 77(3), pp.763-799.

Chiappori, P.A. and Meghir, C., 2015. Intrahousehold inequality. In Handbook of income distribution (Vol. 2, pp. 1369-1418). Elsevier.

Deaton, A. and Muellbauer, J., 1980. Economics and consumer behavior. Cambridge university press.

Donni, O. and Molina, J.A., 2018. Household Collective Models: Three Decades of Theoretical Contributions and Empirical Evidence.

Doss, C., 2013. Intrahousehold bargaining and resource allocation in developing countries. The World Bank Research Observer, 28(1), pp.52-78.

Dunbar, G.R., Lewbel, A. and Pendakur, K., 2013. Children's resources in collective households: identification, estimation, and an application to child poverty in Malawi. American Economic Review, 103(1), pp.438-71.

Engel, E (1857), Die Productions- und Consumtionsverhaltnisse des Königreichs Sachsen" in Zeitschrift des Statistischen Bureaus des Königlich-Sächsischen, Ministerium des Innern, No. 8 u. 9, 1-54. Reprinted as in Engel E. (1895) "Die Lebenskosten Belgischer Arbeiter Familien früher und jetzt," Bulletin de l'institut international de statistique, tome IX, premiere livraison, Rome.

Han, S. and McCloskey, A., 2019. Estimation and inference with a (nearly) singular Jacobian. Quantitative Economics, 10(3), pp.1019-1068.

Jayachandran, S. and Pande, R., 2017. Why are Indian children so short? The role of birth order and son preference. American Economic Review, 107(9), pp. 2600-2629.

Lindahl, Erik (1958) [1919], "Just taxation – A positive solution", in Musgrave, R. A.; Peacock, A. T. (eds.), Classics in the Theory of Public Finance, London: Macmillan.

Menon, Martina, Krishna Pendakur and Federico Perali, 2012, "On the Expenditure-Dependence of Resource Shares", Economics Letters 117 (2012), pp. 739-742.

Menon, M., Perali, F. and Piccoli, L., 2018. Collective consumption: An application to the passive drinking effect. Review of Economics of the Household, 16(1), pp.143-169.

Pendakur, K. 1999. "Estimates and Tests of Base-Independent Equivalence Scales", Journal of Econometrics 88(1), pp 1-40.

Pendakur, K. 2018. "Welfare Analysis When People Are Different", Canadian Journal of Economics, May, 2018.

Rangel, M. and Thomas, D., 2019. Decision-Making in Complex Households (No. w26511). National Bureau of Economic Research.

De Vreyer, P. and Lambert, S., 2016. Intra-household Inequalities and Poverty in Senegal. mimeo, Paris School of Economics.

Woutersen, T. and Ham, J.C., 2013. Calculating confidence intervals for continuous and discontinuous functions of parameters (No. CWP23/13). cemmap working paper.

8 Appendix

8.1 Appendix 1: Dream Data with Scale Economies

Suppose that there are scale economies as in Browning, Chiappori and Lewbel (2013: BCL), with shadow prices that are linear in market prices $\boldsymbol{p}$, where shadow prices equal $\boldsymbol{Ap}$ for some diagonal matrix $\boldsymbol{A}$. ($\boldsymbol{A}$ can depend on observed household characteristics like household size, but we suppress that dependence here.) An element of $\boldsymbol{A}$ says how shareable a good is. If it is 1, then the good is not shareable; if it is less than 1, the good is shareable, and if it is more than 1, the good has diseconomies of scale. The essence of this model of scale economies is that if individuals demand the vector of quantities $\boldsymbol{q}^t$, the household can satisfy all these demands with a market purchase of the quantity vector $\boldsymbol{A}\sum_t \boldsymbol{q}^t$. For example, in a 3 member household, the value of the element of $\boldsymbol{A}$ corresponding to a perfectly shareable good might be 1/3. This means that the household could deliver a quantity q to each of the 3 people in the household with a market purchase of only q.

Table A1a gives dream data about the quantities consumed by each person. In this world, we observe for each person in the household the quantity of the good that they personally got to consume. Normalize market prices to 1 for all goods, so that we can think of consumed quantities as measured in dollars. We additionally observe the total expenditure of the household on each good.

Table A1a: Dream Data: Quantities

	Quantities			Expend.
	Man	Woman	Child	Total
Food	400	300	200	900
Clothing	50	75	25	150
Shelter	300	300	300	300
Other	500	250	250	500
Total				1850

Table A1b: Dream Data, Expenditure

	Expenditure			
	Man	Woman	Child	Total
Food	400	300	200	900
Clothing	50	75	25	150
Shelter	100	100	100	300
Other	250	125	125	500
Total	800	600	450	1850

For nonshareable goods (food and clothing in this example), the total expenditure of the household is simply the sum of the individual quantity levels (prices are normalized to 1). However, for goods that are shared, this is not the case. In this example, shelter is considered to be a fully shared good. Here, we have that each member reported that they personally consumed \$300 worth of shelter. But, because shelter is fully shared, the household only had to purchase \$300 of housing to accomplish this. This means that the household purchased only 1/3 of the total housing consumption of the 3 members. It is as if the household was able to scale its housing spending up by a factor of 3, and then each member bought housing as a private good out of this scaled purchase. Consequently, we identify the matrix $\boldsymbol{A}$ from these data: the element of $\boldsymbol{A}$ corresponding to shelter is 1/3, because the household only needs to buy 1/3 of the total consumed quantities of all the members.

Goods do not have to be either fully shared or non-shareable in the BCL model; they can be partly shared. Suppose that "other" is transportation, and that transportation costs are for riding a motorcycle. The individual-level quantities in Table A1a are the individual-level numbers of km ridden and the household purchased quantity would be the total number of kilometers shown on the odometer. The sum of the former would exceed the latter, because sometimes people ride together. Suppose the man is the only member who knows how to drive a motorcycle. If the man rode 250km with the woman and 250km with the child, then their consumed quantities would be as in Table A1a, with 1000 person-km driven. But, the motorcycle would only have travelled 500km, so the household would have purchased only 500km. Here, the element of $\boldsymbol{A}$ corresponding to transport (other) would be 1/2.

In Table A1b, we turn individual-level quantities into individual level expenditures by multiplying quantity by price. Since market prices $\boldsymbol{p}$ are normalized to 1, within household prices given by $\boldsymbol{Ap}$, this means we multiply by the diagonal matrix $\boldsymbol{A}$. Since nonshareable goods have an element of $\boldsymbol{A}$ equal to 1, for the nonshareable goods of food and clothing, the rows of Table A1b are identical to the rows of Table A1a. The elements of $\boldsymbol{A}$ for shelter and other, respectively, are 1/3 and 1/2. So, for shelter, we multiply by 1/3 and for other, we multiply by 1/2. This yields Table A1b which gives the expenditure of each person on each good. These can be summed down columns to yield the total expenditure of each person, and these person-level total expenditures add up to household-level total expenditure in the bottom right corner.

Scale economies in the BCL model are thus driven by the matrix $\boldsymbol{A}$ which scales prices. We like scale economies because we like low prices. The value of scale economies is just the cost of living index corresponding to the difference between facing a price vector $\boldsymbol{p}$ and facing a price vector $\boldsymbol{Ap}$. BCL show how to identify resource shares and the matrix $\boldsymbol{A}$ from knowledge of individual demand vector functions for all goods and household demand vector functions from all goods (as functions of prices and budgets).

DLP do not attempt to identify the matrix $\boldsymbol{A}$. Instead, they show how to identify just the resource shares from knowledge of just household Engel curve functions (without price variation) for assignable goods, where the assignable goods are assumed to be non-shareable. The model of DLP does not make any assumptions about how shareable the non-assignable goods are. In terms of the matrix A, DLP assume that the single element of A corresponding to the non-shareable assignable good equals 1, and make no assumptions about the other elements of A.

Although the model of DLP is not affected by whether or not scale economies are assumed to exist, the characteristics of the dream data are affected by this assumption. In particular, if we want to identify scale economies *as well as* resource shares directly from data, then such data must provide (at least) the individual-level experienced quantities of each good as well as household level expenditure on these goods.

The matrix A governs scale economies and is relevant to poverty calculations. The standard tool used to estimate poverty in developing countries is to compare per-capita income to a poverty threshold of US$1.90 per day. The assumption on scale economies underlying this strategy is that there are no scale economies. If we took scale economies seriously in the measurement of poverty, we would scale up household consumption by the matrix A to give an estimate of the total consumption of all people in the household. If we then take within-household inequality seriously in the measurement of poverty, we would multiply this scaled household consumption by the resource share of each person, and compare this quantity to the poverty threshold of US$1.90 per day. This paper deals with only the latter issue. Simple estimation tools to recover scale economy parameters in household models remain an important issue for future research.

8.2 Appendix 2: Weaker Restriction on A

Denote the household quantity demand vector as $\boldsymbol{Q}(\boldsymbol{p}, y)$ and individual demand functions as $\boldsymbol{q}^t(\boldsymbol{p}, y)$. The BCL model where $\widetilde{\boldsymbol{p}} = \boldsymbol{A}\boldsymbol{p}$ (for unrestricted $\boldsymbol{A}$) implies that, given the sharing in the household, the household purchases enough of each commodity so as to give each individual in the household exactly what they would have purchased given their shadow budget constraint $(\boldsymbol{A}\boldsymbol{p}, \eta^t(\boldsymbol{p}, y)y/N^t)$. Satisfying these demands requires only purchase of $\boldsymbol{A}$ times that vector of summed quantities. This gives equation (6) from BCL (slightly rewritten):

$$\boldsymbol{Q}(\boldsymbol{p}, y) = \boldsymbol{A} \sum \boldsymbol{q}^t(\boldsymbol{A}\boldsymbol{p}, \eta^t(\boldsymbol{p}, y)y/N^t).$$

Note that in BCL, $N^t = 1$.

We reproduce equation (1) from the current paper, which states a restriction on the matrix $\boldsymbol{A}$ as follows:

$$\boldsymbol{A} = \begin{bmatrix} A_1 & 0 \\ 0 & \boldsymbol{A}_2 \end{bmatrix}$$

Given this restriction, for the assignable goods of each person of type t , the household purchases A_1 times the quantity that they would have purchased. Let $Q^t(y) = Q^t(\boldsymbol{p}, y)$ be the household quantity demand for the assignable good of person t at the fixed price vector $\boldsymbol{p}$. Let $q^t(y)$ be the quantity demand of person t for the assignable good, evaluated at the fixed price vector $\boldsymbol{A}\boldsymbol{p}$ (this is the shadow price vector, not the market price vector). Substituting in, we have

$$Q^t(y) = A_1 N^t q^t(\eta^t(y)y/N^t).$$

Here, A_1 accounts for scale economies relevant to the assignable good.

To move from quantity demands (at fixed prices) to Engel curve functions, let p_1 ($\widetilde{p}_1$) be the first element of $\boldsymbol{p}$ ($\widetilde{\boldsymbol{p}}$), which gives the price (shadow price) of the assignable good, and notice from equation (1) that $\widetilde{p}_1 = A_1 p_1$. Dividing both sides by the y and multiplying by p_1 yields household Engel curves for the assignable good of type t:

$$W^t(y) = p_1 Q^t(y)/y = p_1 A_1 N^t q^t \left(\eta^t\left(y\right) y/N^t\right)/y.$$

Next, we substitute in the Engel curve function for a person of type t : $w^t(y) = A_1 p_1 q^t(y)/y$. Here, we multiply quantities by shadow prices $A_1 p_1$ because q^t is evaluated at shadow prices. Substituting in the shadow budget $\eta^t\left(y\right) y/N^t$ of each person of type t for y, we have,

$$w^t(\eta^t\left(y\right) y/N^t) = N^t A_1 p_1 q^t \left(\eta^t\left(y\right) y/N^t\right)/\eta^t\left(y\right) y.$$

Substituting the expression for $w^t(\eta^t\left(y\right) y/N^t)$ into the expression for $W^t(y)$ yields equation (2), reproduced here:

$$W^t(y) = \eta^t\left(y\right) w^t \left(\eta^t\left(y\right) y/N^t\right).$$

This is equation (3) of DLP (for $N^m = N^f = 1$). DLP assumed "too much" about $\boldsymbol{A}$ in the sense that while they assumed

$$\boldsymbol{A} = \begin{bmatrix} 1 & 0 \\ 0 & \boldsymbol{A}_2 \end{bmatrix}$$

with $\boldsymbol{A}_2$ diagonal, the weaker restriction (1) yields the same Engel curve functions (2) for assignable goods. Our restriction (1) is less restrictive than DLP's restriction in economically meaningful ways: a) our restriction allows for assignable goods with scale economies or diseconomies; and, b) our restriction allows for unrestricted and unknown complementarities between non-assignable goods. However, like DLP, our restriction requires that there be no complementarities between assignable and non-assignable goods.

8.3 Appendix 3: Accounting for Scale Economies in Poverty Measurement

The standard estimate of the poverty rate does not account for scale economies in consumption. In this paper, we do not estimate scale economies. Instead, we consider an off-the-shelf adjustment for scale economies.[27] The per-capita approach assigns y_h/N_h to each household member. The OECD uses an alternative approach, wherein each household member is assigned $y_h/\sqrt{N_h}$, to account for the fact that members of large households can access scale economies. We can think of the OECD approach as first inflating household expenditure by $\sqrt{N_h}$, and then dividing equally among household members, assigning $\sqrt{N_h}y_h/N_h = y_h/\sqrt{N_h}$ to each member. In Table A3, we pursue this approach. The analogous approach using our resource shares first inflates the household budget by $\sqrt{N_h}$ to account for scale economies in consumption, and then multiplies by resource shares to assign a consumption level $\sqrt{N_h}y_h\eta_h^t$ to each member.

Table A3, Estimated Poverty Rates, with Scale Economies					
country	per-capita	men	women	children	all people
	est	est	est	est	est
	std err	*std err*	*std err*	*std err*	*std err*
Albania	0.016	0.069	0.066	0.111	0.078
	0.002	*0.046*	*0.046*	*0.082*	*0.033*
Bangladesh	0.048	0.007	0.034	0.112	0.055
	0.003	*0.002*	*0.007*	*0.040*	*0.015*
Bulgaria	0.022	0.058	0.052	0.316	0.092
	0.003	*0.048*	*0.034*	*0.119*	*0.033*
Iraq	0.000	0.000	0.002	0.130	0.056
	0.000	*0.000*	*0.002*	*0.070*	*0.030*
Malawi	0.238	0.192	0.228	0.298	0.253
	0.004	*0.029*	*0.040*	*0.047*	*0.022*
Bangladesh	0.044	0.012	0.052	0.063	0.045
Food	*0.002*	*0.007*	*0.018*	*0.016*	*0.007*

The poverty rates shown in Table A3 are much lower than those in Table 6, because large households are assumed to enjoy substantial scale economies that raise the consumption of their members. Indeed, in Iraq, no household in the sample had total expenditures lower than $\sqrt{N_h}$ times the poverty threshold of US$5.50 per person per day, resulting in an estimated poverty rate of 0. However, because estimated resource shares of some members may be much smaller than $1/N_h$ for some households, we see that our estimated poverty rates for women and children in Iraq are positive.

The big-picture finding from Table 6 in the main text is unchanged by this accounting for scale economies. In those countries where the point-estimate of gender disparity in resource shares is positive, the point-estimates of poverty rates are higher for women than for men.

[27] Tractable estimation of scale economies in household consumption remains a task for future research (see Calvi et al 2020 for a promising approach.