Ordinary Least Squares Estimation of the Intrahousehold Distribution of Expenditure

Valérie Lechene

University College London

Krishna Pendakur

Simon Fraser University

Alex Wolf

Compass Lexicon

We provide a method to estimate resource shares—the fraction of total household expenditure allocated to each household member—using linear (e.g., ordinary least squares) estimation of Engel curves. The method is a linear reframing of the 2013 nonlinear model of Dunbar, Lewbel, and Pendakur, extended to allow single-parent and other complex households, scale economies in assignable goods, and complementarities between nonassignable goods and supplemented with a linear identification test. We apply the model to data from 12 countries and investigate resource shares, gender gaps, and poverty at the individual level. We reject equal sharing and find large gender gaps in resource shares, and consequently in poverty rates, in some countries.

I.   Introduction

Many aspects of well-being depend critically on individual-level expenditure and consumption. The United Nations' Millennium Development Goals include the promotion of gender equality and the empowerment of women, which partly have to do with women's access to resources within households. Many important questions in labor, public, and development economics also hinge on the intrahousehold distribution of resources. For example, part of the motivation for women's microcredit and child transfers delivered to women (rather than male household heads) is that these strategies may increase women's bargaining power within the household and therefore strengthen their claims on household expenditure.

Children's access to household resources is also of critical importance. Poverty in childhood is well understood to have 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 a child in a low-income household has low consumption, even though parents might devote a greater fraction of resources to their children than to themselves. Poverty—and its persistence and correlation with long-term outcomes—can be measured adequately only if we use tools that accommodate within-household inequality.

The standard poverty measurement strategy assigns to each household member a per capita share of household expenditure and compares that to a poverty threshold, for example, US$1.90 per day. This is a matter of convenience and data availability, rather than a matter of principle: it has neither behavioral basis nor theoretical justification. A strategy that respects the idea that poverty is experienced at the individual level would instead assign each person an individual-level consumption measure and compare that to $1.90 per day.

It is conceptually simple, but practically difficult, to measure the consumption of each individual in a household. 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 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 consumption level of these goods to each individual is not straightforward.

Resource shares, defined as the fraction of total expenditure allocated to each household member, describe the within-household distribution of expenditure. If, in a given household, women have smaller resource shares than men, there is gender inequality in expenditure. Resource

shares describe claims to total household expenditure, and they allow for the fact that some goods may be privately consumed but other goods may be shared to unknown degrees. In this paper, we show how to identify resource shares from Engel curves estimated with linear (e.g., ordinary least squares [OLS] or two-stage least squares [2SLS]) regressions on household-level expenditure data.

Resource shares can help us understand a wide variety of phenomena. Calvi (2020) 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"). Relatedly, women's resource shares can serve as a measure of women's empowerment in the consumption domain (complementing those of, e.g., Pulerwitz, Gortmaker, and DeJong 2000, Alkire et al. 2013, and Ewerling et al. 2017).

Many researchers have studied the consequences of unequal sharing within households, using reduced-form approaches. Jayachandran and Pande (2017) provide evidence that Indian children farther down the birth order are considerably more stunted, which they attribute to favoritism for firstborn children. But does this favoritism run through a channel of greater access to household resources, that is, higher resource shares?

Like household models going back to Becker (1981), our paper promotes the stance of methodological individualism. Our model 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 (see also Cherchye, De Rock, and Vermeulen 2007). Using this general framework, Browning, Chiappori, and Lewbel (2013) and Dunbar, Lewbel, and Pendakur (2013) introduce structural models (referred to below as the BCL and DLP models, respectively) 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. Both these works propose nonlinear structural models to estimate resource shares.

Nonlinear structural models can be computationally difficult to estimate and opaque in terms of their identifying variation. In the BCL and DLP models, a key computational difficulty is that resource shares must be between 0 and 1, and they enter the model nonlinearly, implying that bounded nonlinear estimation is required. Modern statistical software (e.g., R, Stata, and Matlab) allows users to write their own code to estimate nonlinear models, bound parameter spaces, and tell optimizers

684   JOURNAL OF POLITICAL ECONOMY

what to do when parameters hit boundaries. So the BCL and DLP models can be estimated with currently available data and computing resources.

However, in the years since the collective-household model first gained traction (circa 1991) and since the publication of the BCL and DLP models, these approaches have not made their way to informing policy. Despite the availability of suitable data and substantial interest in within-household inequality from both practitioners and academics, household models like these have not yet entered widespread use.

Over the past 2 decades, much of empirical economics has moved away from complex empirical models and toward linear methodologies such as OLS and 2SLS. These linear methods are widely understood, simple to implement, and computationally cheap and have a unique solution. For these reasons, linear methodologies are viewed by some as more transparent than more complex methodologies. Indeed, a commonly held view is that if you cannot see an empirical result in a linear regression, then it probably is not real. We believe that the lack of a simple and transparent empirical methodology is the reason that structural models identifying resource shares have not been used widely in policy work, studies of gender disparities, and poverty estimation.

Our proposed estimator is a linear reframing of the DLP model. It is simple, transparent, and implementable with real-world data. We hope that this paper will jump-start the endeavor of measuring within-household economic inequality and rescue the collective-household model from being a purely academic exercise. It will open the door to policy makers and practitioners using these methods and models to illuminate gender-based and other inequalities within the household and to more adequately formulate programs that are delivered at the household level. Our OLS-based estimator will lead more researchers to see these household models as practical in their empirical settings and therefore foster the application of theory to data.

Our linear reframing of the DLP structural model requires only the estimation of linear Engel curves for one assignable good for each person. 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 good* is a good where we observe expenditure at the person level rather than at the household level, for example, women's clothing.

Our methodological contribution is to show that, conditional on covariates, the DLP model can be written as a linear reduced form wherein the resource shares are functions of the reduced-form coefficients. Here, it is easy to see the variation that identifies resource shares: it is the relative size of the budget responses of household-level assignable-goods Engel curves.

As well as developing a theory-consistent OLS route to estimation of resource shares, we extend the DLP model to allow for complex household

types, including those with multiple adult men and/or women and single-parent households. This is particularly important, since nuclear households are far from the norm, in particular in developing countries. We also extend the model to allow for assignable goods that have scale economies in consumption and nonassignable 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.

Our empirical contribution is to 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 11 household surveys from the World Bank Living Standards Measurement Study (LSMS) data and one national survey from Bangladesh. We use person-level clothing expenditure as the assignable good. Clothing Engel curves pass the identification test for five of the 12 countries, and we 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 4 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, lending credibility to our methods. We provide arguments to support the use of food data to estimate resource shares and therefore suggest that national consumption surveys should collect data on individual food intakes.

Given that we offer a simple and tractable methodology with low data requirements, we hope that researchers interested in intrahousehold 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 one assignable good. Such data would be sufficient to estimate resource shares.

In section II, we review the theoretical foundations of our work. In section II.A, we present the specific BCL and DLP models that underpin our work. We show our new theoretical work in sections II.B–II.E. We present the data in section III and estimated resource shares, gender gaps, and poverty rates in section IV. We finish with a brief discussion of the implications of our work.

II.   Theory

A.   Prologue

Before reviewing the formal theory, it is useful to consider why direct measurement of individual consumption is difficult. Consider for the moment a world where all goods are private (we will allow for nonprivate goods shortly). Let *consumption* refer to quantities consumed and *expenditure* to products of quantities by prices. This distinction is critical when we come to the case where not all goods are private.

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 one child.

Suppose that the poverty line is $1.90 per person per day, which defines a household-level poverty line of $2,080 per annum. Since this household has total expenditure of only $1,850, standard poverty measurement (which assumes equal division within the household) would call all members of this household "poor."

However, with these dream data, we observe the (unequal) expenditure level of each person, and we can compare individual expenditure levels to individual poverty thresholds. Individual levels of total expenditure are given by the column totals. The man's total expenditure on all goods is $800. Since the individual poverty threshold is $1.90 per day, corresponding to $694 per annum, the man is not poor. However, the woman's total expenditure is $600, which falls below $694, and she is poor. Similarly, since the child's total expenditure is $450, the child is also poor.

TABLE 1
DREAM DATA AND REAL DATA ON EXPENDITURES

	Man	Woman	Child	Total
	A. Dream Data on Expenditures			
Food	400	300	200	900
Clothing	50	75	25	150
Shelter	100	100	100	300
Transport	250	125	125	500
Total	800	600	450	1,850
Resource shares (%)	43	32	24	100
	B. Real Data on Expenditures			
Food				900
Clothing	50	75	25	150
Shelter				300
Transport				500
Total				1,850
Resource shares				

The dream data also reveal resource shares, defined as ratios of individual-level total expenditure to household-level total expenditure. The man's resource share is 43% (800/1,850), and the woman and child's are 32% and 24%, respectively. Note that resource shares do not correspond to consumption ratios for any particular good. For example, the man's resource share is higher than the woman's, but the woman's clothing expenditure is higher than the man's. Resource shares measure access to the total household budget.

Thus, the dream data reveal within-household inequality in resource shares and the fact that some household members are poor while others are not.

If data like those in table 1A were widely available, poverty measurement at the person level would be straightforward. Cherchye et al. (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 similar data to investigate individual-level poverty and food deprivation. Bargain, Lacroix, and Tiberti (2020) 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 most (or all) consumption categories are collected.

To consider the case where some goods are not private, we use the distinction between shareable and nonshareable goods introduced by Browning, Chiappori, and Lewbel (2013). *Nonshareable* goods are private goods: they have 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 that if two members eat one unit each, the household must buy 2 units. In contrast, *shareable* goods have the property that the sum of the quantities consumed by all household members is greater than the quantity purchased by the household. For example, if two people ride a motorcycle together, they each consume a motorcycle ride, but the household has to purchase gasoline for only 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. This is where the distinction between expenditure and consumption is critical. In this example, two people consume a motorcycle ride each, so that total consumption of rides is 2, but household expenditure is 1, and individual expenditures on rides are 1/2.

Much of the literature on collective-household models emphasizes the distinction between private goods, which are nonshareable, and public goods. In contrast, we emphasize the distinction between nonshareable goods and shareable goods. There are similarities between public goods and shareable goods. Public goods are fully shareable in the sense that

the household can attain a quantity consumed equal to q for each household member by spending pq on that good. (For nonshareable goods, it would have to spend Npq, where N is the number of members.) But for public goods, all members must consume the same amount (equal to q). In contrast, shareable goods can have any degree of shareability, and household members can consume different amounts of them. Another difference between shareable goods and public goods is the way we represent their price. Public goods have Lindahl ([1919] 1958) prices, which are different for each person. Shareable goods have shadow prices, which embody the scale economies associated with their consumption and are the same for all members of the household.

Because sharing of goods results in more consumption by individuals than the nominal value of what the household purchases, the (shadow) price of consumption of shared goods is lower than the market price. That is, shareable goods feel cheap within the household. In contrast, for nonshareable goods, the household must purchase the sum of the what each individual consumes. Therefore, nonshareable goods feel just as expensive to individuals living in households as they do to individuals living alone.

Consider the individual-level shelter expenditures in table 1A. Suppose that shelter is fully shareable, so that each individual can consume what the household purchases. In this case, each person spends $100 but consumes $300 worth of shelter at market prices. It is as though individuals consuming shelter within the household pay a shadow price equal to one-third the market price (see sec. A1 for a full description of this).

Regardless of how shareable different goods are, resource shares have the same interpretation: they are individual fractions of total household expenditure, and these fractions of total expenditure are spent at shadow prices, not at market prices. In table 1A, the household's expenditure on the man is $800, but if shelter is fully shareable and all other goods are nonshareable, his consumption is valued at $1,000 at market prices.

Shareability thus governs the total quantities that may be consumed by household members and therefore the size of the pie to be allocated to household members. In contrast, resource shares are exclusively about the allocation of that pie, regardless of its size. In this work, we focus on the estimation of resource shares (allowing for shareability), but we do not show how to estimate the shareability of goods. Existing methods for estimating individual poverty similarly compute shares of expenditure and may make ad hoc adjustments for shareability.

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 because of international research efforts such as the 100-plus data sets in the 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 a subset of commodities collected at the individual level.

In this paper, we show how to estimate the resource shares—43%, 32%, and 24% in table 1A—with incomplete data, as in table 1B. However, we do not fill in all the missing cells of table 1B, nor do we make assumptions on which goods are shared or how much they are shared. We estimate only the resource shares, and the estimated resource shares are compatible with any scale economies, for any good, including none. If there are shareable goods, resource shares are in terms of expenditure, not in terms of consumption.

We base our work on the BCL and DLP models, which model the allocation problem of the household and "back out" the individual resource shares from these incomplete data. The DLP model uses information on individual-level spending on nonshareable assignable goods (clothing in the tables) to infer the resource share. They do not infer the individual-level expenditure on any particular good (e.g., the individual cells of the "Transport" row in table 1A). Importantly, the fact that the man has less clothing expenditure than the woman does not imply that he has a smaller resource share than she. Instead, the link between individual assignable-goods expenditure and individual-level total expenditure is driven by the response of the former to the total household budget. If the man's clothing expenditure responds more to a change in the household budget than does the woman's, then he has a larger claim on household resources than she does.

B.   The BCL and DLP Models

Efficient collective-household models are those in which the individuals in a household are assumed to reach the (household) Pareto frontier (see Becker 1981 and Chiappori 1988, 1992). As in earlier results in general equilibrium theory, the assumption of Pareto efficiency is very strong:[1] it means that the household-level allocation problem is observationally equivalent to a decentralized, person-level allocation problem. Collective-household models therefore let us picture the household as a machine that makes budget constraints for its members. Each person's budget constraint is characterized by a shadow budget and a shadow

---

[1] Much effort has gone into testing the restriction of efficiency by Cherchye, De Rock, and Vermeulen (2007), Attanasio and Lechene (2014), Rangel and Thomas (2019), and many others. In this paper, we take efficiency as a maintained hypothesis.

price vector. They are "shadow" budgets and prices because they govern each person's consumption demands but are not observed. Shadow budgets add up to the household budget. Shadow prices may be different from market prices.

Browning, Chiappori, and Lewbel (2013) provide a general collective-household model (the BCL model) whose parameters (shadow budgets and shadow prices) may be estimated with data on the consumption behavior of single individuals and collective households at many budgets and price vectors. Dunbar, Lewbel, and Pendakur (2013; the DLP model) provide sufficient restrictions on the BCL model that shadow budgets may be identified with data on the consumption behavior of collective households observed at a single price vector. We now present the core of the BCL model and the identification restrictions imposed by the DLP model, pointing out where our restrictions are less restrictive than theirs.

Let t index the types of individuals, in our case, "m" for adult male, "f" for adult female, and "c" for child. Let the household consist of N^t individuals of each type t, and let $N = \Sigma_t N^t$. The types are in some sense defined by the data, as we see below. Let y denote the observed household budget.

The share of the household budget allocated to persons of type t in the household is called their resource share, denoted η^t and satisfying $\Sigma_t \eta^t = 1$. Resource shares may in general depend on household budgets, prices, and household and individual characteristics (including *distribution factors*, defined as variables that affect resource shares but not individual preferences). Most importantly, they can vary across the types of individuals in the household: for example, men's and women's resource shares may be unequal.

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

The BCL model was written for a childless adult couple, and the DLP model allowed for children of any number but not for multiple adults of a given gender. With the notation above, we allow for multiple members of any type (see also Calvi 2020). This extension is trivial mathematically but is vital to allowing the model to accommodate the complex households observed in many developing countries, including multigenerational, multifamily, polygamous, and single-parent households.

Shadow prices for goods are the within-household prices of consumption. Shadow prices are the same for all household members. If they were not the same, then there could be gains from trade across household

members, a violation of the assumption of efficiency. Let p denote the market price vector of goods, and let $\tilde{p}$ denote the shadow price vector of goods.

In the BCL model, all consumption is private (there are no public goods), but some goods are shareable. For nonshareable goods, the household purchases a quantity equal to the sum of individual consumption. For shareable goods, the household purchases a quantity less than the sum across individuals of consumption. If each of the N^t individuals of type t in household consumes a quantity vector q^t, then the household purchases quantity vector Q, given by

$$Q = A\sum_t N^t q^t, \tag{1}$$

where A is a square matrix that embodies the consumption technology relating quantities purchased to goods consumed by individuals. This implies that the shadow price of consumption within the household is Ap:

$$\tilde{p} = Ap.$$

A good that is not shareable has a shadow price equal to its market price. A good that is shareable has a shadow price less than its market price. The diagonal elements of A have a direct effect on the size of the shadow price relative to the market price. Its off-diagonal elements capture complementaries in the household consumption technology.

In table 1A, if shelter is fully shared, then when the household buys $Q = 300$, each individual consumes $q^t = 300$. With $N^t = 3$ household members, we have that, with some abuse of notation, $A = 1/3$. Therefore, the shadow price of shelter is one-third of the market price, so that each individual spends \$100 on shelter but consumes \$300 worth of shelter.

Browning, Chiappori, and Lewbel (2013) show that we can identify resource shares and shadow prices from consumption data like those in table 1B. In general, this is possible if we observe Engel curves at many observed price vectors and assume that single individuals have the same preferences as individuals who 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, in many countries, children and unmarried men and women rarely live alone.

Dunbar, Lewbel, and Pendakur (2013) provide sufficient restrictions on the BCL model that resource shares are identified from data on just Engel curve functions for assignable goods of collective households facing a single price vector. Thus, identification in the DLP model does not hinge on the observability of singles' demand functions or of price variation. To achieve this, they use the BCL model in combination with data

on assignable goods, restrictions on the consumption technology (A), restrictions on the resource shares (η^t), and restrictions on preferences.

An assignable good is a good for which we can observe the expenditure of each (type of) individual. Such goods are very useful for identification of household models (see, e.g., Chiappori and Ekeland 2009). The DLP model assumes that there is a single assignable good observed for each type of person. Generally speaking, the available data on assignable goods will define the typology of individuals. In our data, assignable spending on clothing is recorded for adult men, adult women, and children, so these are our types of people. If data were recorded by gender for children as well as for adults, there would be four types of individuals: adults or children and males or females.

The BCL and DLP models allow for caring preferences, where one person's utility level affects another person's utility (see also Cherchye, De Rock, and Vermeulen 2007). However, in these models, full identification is not possible in the presence of direct consumption externalities. Consequently, the DLP model does not allow for externalities in the consumption of the assignable good, where one person's assignable-good consumption affects another person's demand functions. Our main assignable good is clothing, so ruling out externalities may be unpalatable (we come back to this in the empirical discussion). Nonetheless, the absence of consumption externalities between the assignable good of a given member and the demand functions of any other household member is a maintained assumption in our work.

Sort the market price vector p so that its first T elements, denoted $p_1 = (p_1^1, \ldots, p_1^T)'$ are the market prices for the assignable good for types $1, \ldots, T$. Whereas the BCL model allows for an unrestricted A matrix (with any numbers in both diagonal and off-diagonal elements), the DLP model considers a restricted A matrix, where A is a diagonal matrix, with 1 in the elements corresponding to the assignable good of each person. This means that the assignable good of each person must be nonshareable (have no scale economies) and that its shadow price equals its market price. It also means that there are no complementarities in the household consumption technology.

We derive the same demand equations as Dunbar, Lewbel, and Pendakur (2013), with a weaker restriction on the matrix A.[2] Here, we require that A is a block-diagonal matrix satisfying

[2] Specifically, Dunbar, Lewbel, and Pendakur (2013) used an unnecessarily strong restriction on A_1 and A_2 for identification of resource shares, given their similar-across-people (SAP) preference restriction. For that, the restriction in eq. (2) is sufficient for identification. But for identification of resource shares, given their similar-across-types (SAT) restriction, allowing nondiagonal A_1 does not make sense in the context of the model. For identification, given both SAP and SAT restrictions, Dunbar, Lewbel, and Pendakur could have allowed for nondiagonal A_2.

$$A = \begin{bmatrix} A_1 & 0 \\ 0 & A_2 \end{bmatrix} \tag{2}$$

where A_1 is a diagonal matrix with elements A_1^t for persons $t = 1, \dots, T$, each giving the price scale that multiplies the market price of the assignable good to get the shadow price of that person's assignable good. The matrix A_2 is unrestricted.

Unlike Dunbar, Lewbel, and Pendakur (2013), we allow for the possibility that the assignable good has economies or diseconomies of scale ($A_1^t \neq 1$). Consider the example of food waste in food preparation. Suppose that one portion of food is wasted regardless of the number of portions prepared. For example, to prepare three portions, the household must buy four portions (since one is wasted), so that one-fourth of food purchases are wasted. Then, for a household with three people, $3\tilde{p} = 4p$, so that $\tilde{p} = (4/3)p$. For a household with N people, $1/(N + 1)$ of food purchases are wasted, and $\tilde{p} = [(N + 1)/N]p$. These are diseconomies of scale that decrease with household size.

Unlike Dunbar, Lewbel, and Pendakur (2013), we allow for the possibility that the nonassignable goods exhibit complementarities in consumption (nondiagonal A_2). For example, if food is the assignable good, there could be complementarities in the consumption technology between clothing and household heating.

The restriction that the off-diagonal blocks of A equal 0 rules out complementarities in consumption between the assignable goods and all other goods. This is a maintained assumption in Dunbar, Lewbel, and Pendakur (2013) and in this work.

The structure on the matrix A plays a sideshow role here. We do not try to estimate it in this paper. Tractable estimation of A is a job for future research (see, e.g., Calvi et al. 2020).[3] 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 A takes: the resource share is the fraction of the (observed) household budget enjoyed by a type of person and spent at (unobserved) shadow prices Ap.

For now, treat the numbers of household members N^t as constants. Below, we condition the entire model on other observed covariates, but we suppress that here. So, imagine data where N^t and other covariates are constant across households but where prices p and budgets y vary. Let $\eta^t(p, y)$ be the resource share of type t.

---

[3] Our methodology estimates resource shares at a given price vector, without knowledge of prices. Since we do not observe market prices, we cannot estimate shadow prices. Other methodologies use observed price variation to identify shadow prices and thus scale economies (e.g., Browning, Chiappori, and Lewbel 2013; Pendakur 2018). Still others use budget variation to identify welfare-relevant features of shadow prices (e.g., Lewbel and Pendakur 2008, 2021; Calvi et al. 2020).

Assume that each person demands their own assignable good and demands zero of any other person's assignable good. Let $q^t(\boldsymbol{p}, y)$ be the scalar-valued demand function of a person of type t for their assignable good. Individual demand within the household is evaluated at their shadow budget constraint and so equals $q^t(\boldsymbol{Ap}, \eta^t(\boldsymbol{p}, y)y/N^t)$. Substituting this demand and the restriction (2) into equation (1) gives household quantity demand functions for assignable goods:

$$Q^t(\boldsymbol{p}, y) = A_1^t \sum_t N^t q^t = A_1^t N^t q^t(\boldsymbol{Ap}, \eta^t(\boldsymbol{p}, y)y/N^t).$$

Much work on consumer demand focuses on 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).[4] Engel curve functions hold prices constant at some vector $\bar{\boldsymbol{p}}$ and evaluate the fraction of expenditure as a function of the total household budget. Denote the resource share at the fixed price vector $\bar{\boldsymbol{p}}$ as $\eta^t(y) = \eta^t(\bar{\boldsymbol{p}}, y)$. At the fixed price vector $\bar{\boldsymbol{p}}$, the household Engel curve function for the assignable good of type t, $W^t(y)$, is given by

$$W^t(y) = \bar{p}_1^t A_1^t N^t q^t(\boldsymbol{A\bar{p}}, (\eta^t(y)y/N^t))/y.$$

Let $w^t(y) = A_1^t \bar{p}_1^t q^t(\boldsymbol{A\bar{p}}, y)/y$ be the Engel curve function at the fixed shadow price vector $\boldsymbol{A\bar{p}}$ for a person of type t for their assignable good at budget y. Substituting in the shadow budget and then substituting into the above equation,[5] we get equation (3) of Dunbar, Lewbel, and Pendakur (2013):

$$W^t(y) = \eta^t(y)w^t(\eta^t(y)y/N^t). \tag{3}$$

The relationship (3) says that the household's Engel curves (at market prices, held fixed) for the assignable goods are equal to the resource share of the relevant type times the Engel curve of a person of that type facing the shadow price vector and their shadow budget. This model is not identified without further structure: there are $2T - 1$ unobserved functions ($\eta^t(y)$ and $w^t(y)$), but only T observed functions ($W^t(y)$).

Dunbar, Lewbel, and Pendakur (2013) provide sufficient restrictions on the model that resource shares are identified from data on just Engel curve functions for assignable goods of collective households facing a single price vector. Sufficient restrictions are (1) the matrix $\boldsymbol{A}$ is block diagonal, as in equation (2); (2) resource shares do not depend on the

---

[4] 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."

[5] First, note that plugging the shadow budget into $w^t(y) = A_1^t \bar{p}_1^t q^t(\boldsymbol{A\bar{p}}, y)/y$ gives $w^t(\eta^t(p, y)y/N^t) = A_1^t \bar{p}_1^t q^t(\boldsymbol{A\bar{p}}, (\eta^t(p, y)y/N^t))/(\eta^t(p, y)y/N^t)$, and therefore $q^t(\boldsymbol{A\bar{p}}, (\eta^t(p, y)y/N^t)) = (\eta^t(p, y)y/N^t)w^t(\eta^t(p, y)y/N^t)/A_1^t \bar{p}_1^t$. Substituting this into the equation for $W^t(y)$ gives $W^t(y) = \bar{p}_1^t A_1^t N^t(\eta^t(p, y)y/N^t)w^t((\eta^t(p, y)y/N^t))/A_1^t \bar{p}_1^t y$. Canceling terms gives eq. (3).

household budget,[6] so that $\eta^t(y) = \eta^t$; (3) individual Engel curve functions are linear in $\ln y$, so that $w^t(y) = \alpha^t + \beta^t \ln y$ (a case of such Engel curves is the "Almost Ideal" demand system of Deaton and Muellbauer 1980); and (4) preferences are similar—but not identical—across people, such that $\beta^t = \beta$.[7] Substituting these assumptions into equation (3) gives

$$W^t(y) = \eta^t[\alpha^t + \beta(\ln y + \ln \eta^t - \ln N^t)]. \tag{4}$$

The econometric model defined by equation (4) is nonlinear as a result of the fact that η^t multiplies α^t and β and requires positive resource shares, because of the $\ln \eta^t$ term. Its estimation requires nonlinear optimization subject to bounding restrictions on parameters. Although such estimators are feasible, our linear reframing below makes them unnecessary.

C. OLS Estimation of Resource Shares

We now present a theory-consistent linear reframing of the DLP model. Let the subscript $h = 1, \dots, H$ index households. Consider first the case with no demographic covariates (the entire model can be written

---

[6] The assumption that resource shares do not depend on the household budget is strong: it implies that if a household's total expenditure increases, the intrahousehold relative consumption distribution does not change. Surprisingly, there is some empirical support for this restriction (see Menon, Perali, and Pendakur 2012; Cherchye et al. 2015). Note, however, that Dunbar, Lewbel, and Pendakur (2013) require only that resource shares are invariant to expenditure over some range of household expenditure. If this invariance held only for the poorest households, we could still identify resource shares for the very poor. Further, the independence of resource shares from total household expenditure is conditional on other observed covariates, which may include, e.g., income and/or wealth.

[7] Dunbar, Lewbel, and Pendakur (2013) 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 2007) or if cost functions satisfy "independence of base" (Lewbel 1989) or "equivalence-scale exactness" (Blackorby and Donaldson 1993). Dunbar, Lewbel, and Pendakur show that if $\tilde{p} = Ap$, where A is diagonal with a 1 for the assignable good, resource shares do not depend on household budgets, and SAP holds, then resource shares are identified from the Engel curves of collective households at a single price vector. They do not require loglinear Engel curves for identification. When applied to the loglinear Engel curves, SAP implies $\beta^t = \beta$. We (and they) use loglinear Engel curves to make estimation easier, not to achieve identification. The linear reframing we develop below also works with quadratic Engel curves corresponding to the Quadratic Almost Ideal (QAI) model of Banks, Blundell, and Lewbel (1997). In that case, resource shares are identified by ratios of OLS coefficients on $(\ln y)^2$, rather than by ratios of OLS coefficients on $\ln y$. We show this in sec. A3. Dunbar, Lewbel, and Pendakur 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 (1) SAP is consistent with our more general consumption technology (more general A matrix), but SAT is not; (2) given SAP, resource shares are exactly identified with just one household type, e.g., just nuclear households with one child, whereas identification, given SAT, requires at least three different compositions and may be overidentified; and (3) the solution for resource shares, given SAT, is more complex and has multiple solutions when overidentified.

conditionally on covariates, which we do below) and where $N_h^t = N^t$ is constant across households. Rewrite equation (4) with a subscript $h = 1, \ldots, H$ indexing households and an additive error term ε_h^t, as the following linear model:

$$W_h^t = a^t + b^t \ln y_h + \varepsilon_h^t, \tag{5}$$

where

$$a_h^t = \eta^t \alpha^t + \eta^t \beta \ \ln \eta^t - \eta^t \beta \ \ln N^t$$

and

$$b^t = \eta^t \beta. \tag{6}$$

This model may be estimated by linear regression[8] of the observed household-level assignable-good expenditure share, W_t (in Stata notation), on a constant and the log of the household budget, lny. For example, for data on households with $t = m, f, c$, one could implement the linear seemingly unrelated regression system in Stata via

```
sureg (W_m ln y)  (W_f ln y)  (W_c ln y).
```

Rearranging equation (6), we have

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

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

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

One could implement this estimator for η^m in households with this fixed value of N^t in Stata via[9]

```
generate eta_m = [W_m] ln y/([W_m] ln y +[W_f] ln y + [W_c] ln y).
```

The intuition for identification of resource shares is as follows. Note that the estimated resource share does not depend on the estimate of the level term $\hat{a}^t$. It is budget responses of Engel curves, not levels of Engel curves, that identify resource shares. The observable budget semielasticity of household-level Engel curves for assignable goods, $b^t = \partial W^t(y)/\partial \ln y$, is equal to $\eta^t \beta$. Since η^t sum to 1, the sum of this semielasticity across types

[8] We describe only OLS estimators in this paper. But if some regressors are endogenous and instruments are available, then 2SLS estimators are analogous.

[9] This command would generate the resource share for men in all households. It would be identical for all households, because it depends only on estimated parameter values and not on the value of any regressor.

is β. Thus, we have that $\eta^t = b^t / \Sigma_t b^t$. So it is the relative magnitudes of budget semielasticities (b^t) that determine resource shares. If, for example, the household Engel curve for the men's assignable good has twice the slope (twice the value of b^t) of that for the women's assignable good, then the men have twice the resource share of the women.

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.

D. Adding Covariates

The model above does not include any covariates, such as demographic preference shifters, and it holds the numbers of household members N^t constant. Including them does not affect identification but does require some additional notation. Let z be all variables that affect preferences and/or resource shares, including the numbers of household members of each type $N = \{N^t\}$. Let $\tilde{z}$ be the subvector z of that excludes N, so that $z = [N \quad \tilde{z}]$.

Assume that resource shares, η^t, and preference parameters, α^t and β, all depend on z. Substituting this into equation (4) and expanding out the terms, we have

$$W^t(y, z) = \eta^t(z)\alpha^t(z) + \eta^t(z)\beta(z) \ln y + \eta^t(z)\beta(z) \ln \eta^t(z)$$
$$- \eta^t(z)\beta(z) \ln N^t. \tag{7}$$

This nonlinear structural model (eq. [7]) with bounded parameter spaces has been estimated by several researchers on data from several countries (e.g., Dunbar, Lewbel, and Pendakur 2013 in Malawi; Bargain, Donni, and Kwenda 2014 in Côte D'Ivoire; Calvi 2020 in India; De Vreyer and Lambert 2016 in Senegal; Calvi et al. 2020 in Bangladesh). In these papers, bounding $\eta^t(z)$ is often addressed by wrapping it in a logit function, and estimation is typically by nonlinear least squares.[10]

---

[10] As with eq. (4), this model contains a linear term in the log of the resource share, $\ln \eta^t(z)$. If η^t is parameterized as a linear index (and especially if z contains an unbounded variable), then search algorithms trying to find the minimum/maximum of the sum of squares, likelihood function, or generalized method of moments 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 consequence that it may induce numerical problems for nonlinear solvers. For example, they may try to evaluate the function in a region of the parameter space where $\eta^t(z)$ is negative, yielding a missing value for $\ln \eta^t(z)$. Alternatively, η^t may be parameterized as a bounded function of z, but then the behavior of the function near the boundaries may present problems for nonlinear solvers. An additional problem relevant to eq. (7) comes from the fact that the term $\eta^t(z)\beta(z) \ln y$ has quadratic interactions in z multiplying $\ln y$. These make it difficult to precisely identify the dependence of resource shares $\eta^t(z)$ on z, because z affects both η^t and β.

Our linear reframing strategy also works when there are covariates. Rewrite equation (7) with a subscript h on all observed variables and an additive error term ε_h^t as the following linear model:

$$W_h^t = a_h^t + b_h^t \ln y_h + \varepsilon_h^t, \tag{8}$$

where

$$a_h^t = \eta^t(z_h)\alpha^t(z_h) + \eta^t(z_h)\beta(z_h) \ln \eta^t(z_h) - \eta^t(z_h)\beta(z_h) \ln N_h^t, \tag{9}$$

and

$$b_h^t = \eta^t(z_h)\beta(z_h). \tag{10}$$

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

E. Approximation

Unfortunately, neither of these approaches is practical with a high-dimensional conditioning vector z_h. For example, with a constant and nine conditioning variables in z_h, third-order interactions require 444 regressors.[11] So we recommend approximating the model. We approximate the a_h^t term with the linear index

$$a_h^t = a_0^t + a_{\ln N^t}^t \ln N_h^t + \mathbf{a}_z^{t'} z_h. \tag{11}$$

Similarly, we approximate the slope term with

$$b_h^t = b_0^t + \mathbf{b}_z^{t'} z_h. \tag{12}$$

From inspection of equation (10), it is easy to see that this approximation for b_h^t is exact if η^t is linear in z_h and β is independent of z_h (i.e., if β is a constant).

---

[11] That is, having $10^3 = 1{,}000$ triples, deleting permutations, means that there are 222 unique combinations, times 2 (intercept and slope).

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 $z_h \cdot \ln y$. The estimated coefficients on $\ln y$ and $z_h \cdot \ln y$ are estimates of b_0^t and b_z^t, respectively. These may be used to construct an estimate $\hat{b}_h^t$ of b_h^t:

$$\hat{b}_h^t = \hat{b}_0^t + \hat{\boldsymbol{b}}_z^{t'} \boldsymbol{z}_h.$$

Regardless of the specification of b_h^t, and regardless of whether it is taken to be an approximation or exact (because of 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 $\Sigma_t \hat{b}_h^t$ as an estimate of $\beta(z_h)$, which implies that an estimate of the resource share of type t in a household with characteristics z_h is given by

$$\hat{\eta}_h^t = \hat{\eta}^t(\boldsymbol{z}_h) = \frac{\hat{b}_h^t}{\sum_{t=1}^T \hat{b}_h^t} = \frac{\hat{b}_0^t + \hat{\boldsymbol{b}}_z^{t'} \boldsymbol{z}_h}{\sum_t \left(\hat{b}_0^t + \hat{\boldsymbol{b}}_z^{t'} \boldsymbol{z}_h\right)}. \tag{13}$$

Engel curves may be estimated by equation-by-equation OLS or with linear seemingly unrelated regression (SUR).[12] Resource shares may then be computed via equation (13).

Suppose that we have a data set on childless couples, so that $T = \mathrm{m}, \mathrm{f}$, $N^\mathrm{m} = 1$, $N^\mathrm{f} = 1$, $\ln N^\mathrm{m} = 0$, and $\ln N^\mathrm{f} = 0$. Let the data be budget shares `W_t`, log budgets `lny`, 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 equation (13) 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

$$\sum_t \boldsymbol{b}_z^t = 0, \tag{14}$$

---

[12] OLS and SUR coincide if the regressor lists are identical across equations. But since the regressor lists differ across equations in our application ($\ln N_h^t$ shows up in the regressor list only for W_h^t), and since SUR is asymptotically efficient, we recommend using SUR.

implying that $\Sigma_t b_h^t = \Sigma_t(b_0^t + b_{N^m}^t N_h^m + b_{N^w}^t N_h^f + b_{N^c}^t N_h^c)$. Then, estimated resource shares are equal to

$$\hat{\eta}^t(\mathbf{z}_h) = \frac{\hat{b}_0^t + \hat{\mathbf{b}}_z^{t\prime} \mathbf{z}_h}{\sum_t(\hat{b}_0^t + \hat{b}_{N^m}^t N_h^m + \hat{b}_{N^w}^t N_h^f + \hat{b}_{N^c}^t N_h^c)}.$$

Here, we expect all $b_{N^t}^t$ to 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 share of each type equals their per capita share. In particular, if $b_0^t = 0$ for all t, $\mathbf{b}_z^t = 0$ for all t, $b_{N^{t\prime}}^t = 0$ for all $t' \neq t$, and $b_{N^t}^t = \kappa$ for all t, then we get per capita resource shares $\eta^t(\mathbf{z}_h) = N^t\kappa/\Sigma_t N^t\kappa = N^t/\Sigma_t N^t$.[13] 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 four 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 ($\mathbf{z}$, $\ln y$, and $\mathbf{z} \cdot \ln y$); 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 one man, one woman, and one 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 one woman and one child, and analogously for the other two compositions. All test statistics—for example, the Wald test of the per capita model—are simply sums of χ^2 test statistics across the samples for each composition.

F.   A Linear Test of Model Identification

As noted above, if $\beta(\mathbf{z}_h) = 0$, then resource shares are not identified. In this case, the estimated value of the denominator may be close to zero, and the resulting estimated resource shares would be unreliable. If it were the case in the limit, then inference is polluted by weak-identification problems (see Han and McCloskey 2019). Consequently, it is valuable to have a test of identification to tell us whether these methods will work at all. Previous papers (Dunbar, Lewbel, and Pendakur 2013, 2021; Han

---

[13] One could additionally restrict $\Sigma_t \mathbf{b}_N^t = 0$, implying that $\hat{\eta}^t(\mathbf{z}_h) = b_h^t/(\Sigma_t b_0^t)$. This further simplifies the denominator, but at the cost of not nesting the per capita model.

and McCloskey 2019) have tested this identifying restriction, but their tests all involve estimating nonlinear models. Our linear reframing of the DLP model straightforwardly delivers an OLS-based test of whether this identifying restriction for resource shares is supported by the data.

Let the overall assignable-goods Engel curve of the household be given by $W_h = \Sigma_t W_h^t$, and let $a_h = \Sigma_t a_h^t$, $b_h = \Sigma_t b_h^t$, and $\varepsilon_h = \Sigma_t \varepsilon_h^t$. If W_h^t is the fraction of the household budget 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. Note that a_h depends on $\ln N_h$, the vector of logs of numbers of members. Then, our approximate model above implies

$$W_h = a_h + b_h \ln y_h + \varepsilon_h, \tag{15}$$

and OLS regression of W_h on 1, $\ln N_h$, z_h, $\ln y_h$, and $z_h \ln y_h$ yields an estimate $\hat{b}_h$ of $\beta(z_h)$. We propose that an easy and useful test of identification for this model is to test whether overall assignable-goods 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 $E[\hat{b}_h] = \hat{b}_0 + \hat{b}_z' \bar{z}_h$, where $\bar{z}_h$ is the sample average of z_h, as a test statistic. This is a test of the economic hypothesis that the overall assignable-goods Engel curve, evaluated at the mean value 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.[14]

Second, for every observation in the data, we test whether or not $\hat{b}_h = \hat{b}_0 + \hat{b}' z_h$ 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).[15]

---

[14] Because our estimator for resource shares is a ratio of estimated coefficients, it is mathematically similar to an exactly identified 2SLS estimator with possible weak identification. With that view, our test of identification is similar to a weak-identification test in 2SLS. Although in this paper we use a 5% two-sided normal critical value of 1.96, if one thinks in terms of weak identification, a critical value of 3.2 (the square root of Staiger and Stock's 1997 recommended threshold of 10 for the F-statistic) may be more relevant. We thank Isaiah Andrews for noticing this. In our empirical work, we reach the same conclusions even if the more stringent critical value of 3.2 is used.

[15] In our empirical work below, no country has a fraction of households between 64% and 83%, so any cutoff value between these values would have yielded the same set of countries passing the test.

III.   Data

In most countries in the world, national statistical offices regularly collect household expenditure survey data. These data are used as inputs in national accounts, for the calculation of the gross domestic product, to measure inflation, to analyze household spending patterns and behavior, 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 LSMS. These data are standardized to some extent and are the best tool available for cross-country comparisons of poverty in low- and middle-income countries.

The LSMS surveys exist for about 40 countries, and often several waves exist. There are, in total, nearly a hundred country-waves potentially available for the analysis of household consumption behavior. We analyze the most recent waves from 11 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).[16] 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 Lewbel and Pendakur 2021 for details on how person-level food data are aggregated from food diary data).[17] Children are defined as members aged 15 or less, except in Iraq, where we define them as members aged 11 or less. Data on children's clothing expenditures in Iraq have some classification errors not present in other country data sets, and so we take our estimates for Iraq with a little bit of caution.[18] We discuss this further in section A4.

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 Bangladeshi data) for the most recent survey year (shown in the column "Survey Year"), 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 the column "Total H." There is a wide range of sample sizes after this initial cleaning, from 1,503

---

[16] A variety of reasons make the data from the other countries unusable. In some cases, no data on assignable goods are collected; in others, information on elements of nondurable expenditure is missing.

[17] 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 are available on request.

[18] In the Iraqi data, the age of children is not specified in questions asking about expenditures on children's clothing. Thus, there is likely classification error in the Iraqi data on children's clothing expenditures. We investigate this further in sec. A4, where we show that our main conclusions regarding resource shares and gendered poverty go through even in a sample of adults-only households in Iraq.

TABLE 2
DESCRIPTIVE STATISTICS

Country	Survey Year	Total H	Single H	Compositions	Our H	Nuclear H	Budget Mean	Budget Standard Deviation
Albania	2008	3,599	239	mf, mfc	3,279	612	11,084	6,477
Bangladesh	2015	6,434	219	mf, fc, mfc	6,120	2,122	6,416	6,268
Bulgaria	2003	3,018	801	mf, mfc	2,099	412	13,117	7,954
Ethiopia	2016	4,717	503	mf, fc, mfc	3,845	1,481	3,092	3,645
Ghana	2006	8,687	1,922	mf, mc, fc, mfc	6,313	2,195	5,096	4,835
Iraq	2007	17,513	309	mf, fc, mfc	13,935	3,763	25,780	14,501
Malawi	2011	12,271	1,030	mf, fc, mfc	10,873	5,488	3,189	3,758
Nigeria	2013	4,600	349	mf, fc, mfc	3,556	1,013	6,656	20,322
Tajikistan	2009	1,503	54	mf, mfc	1,275	192	10,483	6,250
Tanzania	2015	3,352	320	mf, fc, mfc	2,677	1,133	7,219	5,164
Timor Leste	2008	4,477	229	mf, fc, mfc	3,788	1,577	4,954	4,116
Uganda	2014	3,117	257	mf, fc, mfc	2,468	1,014	2,462	2,262
Bangladesh (food)	2015	6,434	219	mf, fc, mfc	4,990	1,916	6,445	6,287

households in Tajikistan to 17,513 households in Iraq. Below, we pay attention to whether sample size matters to the feasibility of the method.

In the column "Single H," we report the number of households that are composed of a single adult man or woman. Since these households have only 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.[19]

For the estimation of the resource shares, we use all household compositions apart from households with only a single type of individual (i.e., we exclude households composed of men only, whatever their number, and similarly for women), and we allow more than one individual of each type, up to four men, four women, and six children. The possible compositions are mf, mfc, fc, and mc. These indicate that individuals of the type m for men, f 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 the column "Compositions," and the 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.

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

The column "Nuclear H" shows the number of nuclear households (those composed of one man, one woman, and some children) 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 six of our 12 countries.

We then provide the mean and standard deviation in our sample (excluding singles) of the overall budget in (PPP [purchasing power parity]) 2011 US dollars. In some countries in our data, the average household budget is close to the World Bank extreme-poverty line of $2,774 per annum for a four-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 sample on assignable food in Bangladesh. It is different from the clothing sample because fewer observations have valid assignable food data than have valid clothing data. In the analysis below, we compare estimated resource shares from assignable food data with those from assignable clothing data.

IV. Results

We estimate equations (8), (11), and (12) under the restrictions of equation (14) via SUR in Stata. Our observed vector of demographic variables z_h consists of the numbers of men, women, and children (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.[20] Resource shares are then computed via equation (13).

A.   Test of Identification

The statistical significance of the slope of the Engel curve for the sum of household assignable goods provides a test of whether or not resource shares are identified. 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 z-test for its difference from zero. In the rightmost column, we give the fraction of observations whose estimated slope (conditional

---

[20] 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.

TABLE 3
Test of Identification

Country	Sample	Budget Share			z-Test of Slope	% of Sample Significant
		Mean	Standard Deviation	slope at z̄		
Clothing						
Albania	3,279	.041	.042	.014	4.7	84
Bangladesh	6,120	.039	.021	−.016	−21.4	100
Bulgaria	2,099	.036	.040	.014	5.2	90
Ethiopia	3,845	.072	.064	−.011	−3.5	63
Ghana	6,313	.048	.040	−.002	−1.0	63
Iraq	13,935	.070	.046	.020	14.3	99
Malawi	10,873	.025	.036	.009	10.0	98
Nigeria	3,556	.017	.023	−.002	−2.0	51
Tajikistan	1,275	.058	.050	.008	1.9	12
Tanzania	2,677	.044	.058	−.002	−.9	14
Timor Leste	3,788	.022	.020	−.002	−1.8	24
Uganda	2,468	.055	.052	−.004	−1.1	5
Bangladesh						
Clothing	4,990	.038	.020	−.015	−18.3	99
Food	4,990	.571	.149	−.118	−15.8	100

on their observed covariates) is statistically significantly different from zero. The bottom panel of table 3 gives statistics for clothing and food in Bangladesh, using the sample of observations with valid food data.

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

Using a standard critical value of 1.96 for a two-sided z-test, we see that 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 might not be identified.

We also report the percentage of the sample for which the slope is statistically significantly different from zero. For our method to work, this must be high enough, so 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 five countries that pass the test, hence for which the model is identified and resource shares can be estimated. We note that if a threshold of 75% is considered too lax, then a threshold of 84% would result in a change in the set of countries considered to pass the test and

would remove Albania. In the other direction, if 75% is considered too tight, then a value of 63% would bring in Ethiopia.

For Bangladesh, we also have assignable data on food consumption. We have fewer observations (4,990) on food than on clothing (6,120) because there is some nonresponse in the daily food diary data and because some household members are absent on the diary day in some households. Food budget shares are much larger than clothing budget shares: whereas clothing accounts for only 4% of total household expenditure, 57% of household expenditure in our Bangladeshi sample is for food.

A long history of demand analysis, dating back to Engel (1895), 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 −16, and 100% of the sample having 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 downward, while clothing Engel curves sometimes slope upward, sometimes slope downward, and are sometimes flat. Both of these differences suggest that food is a preferable assignable good for our methods.

The model does not specify which assignable good to use to estimate resource shares—using any assignable good provides an estimate of the same underlying resource shares. Consequently, in our analysis below, we pay special attention to the difference—or lack thereof—between estimates of Bangladeshi resource shares based on clothing and food Engel curves.

B.   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.[21] We report both the resource shares estimated at the mean of observed covariates, $\bar{z}$, and the mean of the resource shares evaluated at all z_h. For the former we give the standard error and for the latter the standard deviation. The bottom panel of table 4 gives statistics for clothing and food, using the sample of observations with valid food data. The bottom line gives the difference between estimates using food and clothing.[22]

In Albania, the estimated men's and women's per-person resource shares at $\tilde{z}$ (the average z_h in Albania) are 28% and 25%, respectively, with small standard errors of 3 percentage points. Because resource

TABLE 4
PREDICTED RESOURCE SHARES, SELECTED COUNTRIES

COUNTRY	SAMPLE SIZE	EVALUATED AT $\bar{z}$ (Estimated) Men (SE)	Women (SE)	Children (SE)	EVALUATED AT ALL z_h (Mean) Men (SD)	Women (SD)	Children (SD)	η OUTSIDE [0, 1]	WALD PER CAPITA TEST, df (p-Value)
					Clothing				
Albania	3,279	.282 (.032)	.247 (.033)	.134 (.030)	.280 (.369)	.256 (.340)	.126 (.166)	.062	45, 35 (.129)
Bangladesh	6,120	.312 (.011)	.286 (.014)	.120 (.010)	.311 (.114)	.284 (.118)	.122 (.059)	.000	387, 41 (.000)
Bulgaria	2,099	.304 (.038)	.372 (.041)	.188 (.061)	.292 (.14)	.387 (.218)	.173 (.214)	.079	49, 35 (.058)
Iraq	13,935	.249 (.009)	.210 (.010)	.045 (.007)	.249 (.121)	.211 (.110)	.045 (.035)	.030	473, 47 (.000)
Malawi	10,873	.312 (.028)	.274 (.030)	.124 (.011)	.31 (.179)	.267 (.154)	.127 (.089)	.015	267, 45 (.000)
					Bangladesh (with Valid Food Data)				
Clothing	4,990	.325 (.013)	.290 (.017)	.132 (.012)	.322 (.109)	.290 (.118)	.135 (.064)	.002	288, 41 (.000)
Food	4,990	.309 (.014)	.256 (.015)	.174 (.011)	.313 (.110)	.250 (.114)	.176 (.082)	.042	232, 41 (.000)
Difference: Clothing − Food		.016 (.019)	.034 (.022)	−.043 (.016)	.009 (.001)	.004	−.039	−.018	

shares are nonlinear functions of estimated OLS regression coefficients, the estimate of resource shares at average z_h does not equal the average of estimated resource shares over all z_h. However, they are similar: the sample averages of the resource shares are 28% and 26%, respectively, for men and women. Variation in estimated resource shares is driven by variation in observed covariates z_h. The standard deviations of these estimated resource shares are 37 and 34 percentage points, indicating quite a lot of heterogeneity in resource shares driven by the sample variation in observed covariates.

The rightmost column of table 4 gives the fraction of resource shares that fall outside of the [0, 1] interval. In Albania, this fraction is 6%. This fraction provides another rough assessment of the model. Resource shares are fractions between 0 and 1, so if many point estimates are outside this range, it may signal that some part of the model is wrong. In practice, estimates fall outside (0, 1) when the slope for one type is different in sign from the slope of another type. For example, if clothing was a luxury for men (Engel curves sloped upward) but a necessity for women (Engel curves sloped downward), one of the types would have a negative resource share.

In these five countries, only a small fraction of observations have a predicted resource share outside (0, 1). Bulgaria has the highest density of such observations, with 8% of the point estimates of resource shares outside the valid range. Although some countries have point estimates outside (0, 1), in no country is any individual point estimate statistically significantly outside (0, 1).

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 5% of resources each. In all five countries, the per-person resource shares of children are smaller than those of adults. Resource shares may be driven by differential needs across people as well as by inequity or power imbalance. That the resource shares of children are smaller than those of adults is consistent with the practice of using a lower poverty line for children than for adults, on the basis of presumed lower needs. When we come to measuring child poverty below, we use a lower poverty threshold for children than for adults.

We note that these are per-child resource shares and that Iraqi households with children have an average of 4.1 children, whereas Bulgarian households with children have an average of 2.1 children. Further, Iraqi children are younger, and thus possibly less needful, than children in other countries (by definition). Even so, the estimated per-child resource share of 5% in Iraq seems counterintuitively small and may be related to misclassification of children's clothing expenditures in this particular data set. We elaborate on this in section A4. The key message from that

section is that our main results regarding the gender gap and gendered poverty are seen in an adults-only subsample, where such misclassification is minimized.

A standard resource share in current use by the World Bank and other agencies is the per capita share of household members, that is, $\eta_h^t / N_h^t = 1/\Sigma_s N_h^t$. This would assign each person their per capita share of household expenditure. The estimates in table 4 exhibit quite a lot of cross-country variation. If they correspond to true variation in resource shares, this suggests that a universal tool such as the per capita resource share would leave out a lot of cross-country heterogeneity.

Given our model, the per capita share obtains if $\boldsymbol{b}_z^t = 0$ for all t, $b_{n_{t'}}^t = 0$ for all $t' \neq t$, and $b_{n_{t'}}^t = \kappa$ 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 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, the latter two countries have the smallest samples by a factor of about 2, and consequently, the estimated standard errors of estimated resource shares are larger in these two countries than in the others. This suggests to us that rather large sample sizes are needed to estimate these models and, for example, to test the per capita model.[23]

That the per capita model is rejected suggests that there is substantial within-country variation in resource shares. This can be seen clearly in the standard deviations of estimated resource shares, which are quite large relative to mean resource shares for each type of person in every country. For example, in Bangladesh, the mean women's resource share (based on clothing) is 0.284, with a standard deviation of 0.118. This means that the per capita model not only overestimate women's resource shares but also drastically underestimates the heterogeneity in the per-woman resource share (because in the per capita model, they would all be $1/N_h$). We show below that the failure of the per capita model implies the existence of both gender gaps in resource shares and gendered poverty.

We now turn to the comparison of the resource shares estimated using clothing as the assignable good with the resource shares estimated using food as the assignable good, shown in the bottom panel of table 4. These estimates use the subsample of observations with valid food data. In the

---

[23] 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 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).

leftmost columns, we consider estimated resource shares to be the sample mean values of the covariates z. The estimated per-man resource share in rural Bangladeshi households, using clothing as the assignable good in the sample with valid food data, is 32.5%. Using food as the assignable good in the same sample gives an estimated resource share of 30.9%. The point estimates are close to each other, and the estimated difference of 1.6 percentage points has a standard error of 1.9 percentage points, so we cannot reject the hypothesis that the estimated men's resource share does not depend on which good is used to identify it.

Similarly, the point estimates of the per-woman resource share based on assignable clothing are close to those based on assignable food. The difference between these two estimates is 3.4 percentage points, with an estimated standard error of 2.2 percentage points, so, as with men, we cannot reject the hypothesis that the two assignable goods yield the same resource shares for women.

However, the differences are positive for both men and women: using food as the assignable good delivers lower point estimates for both men and women. As a consequence, estimated children's resource shares are larger when we use food as the assignable good, and the 4.3 percentage point gap between the food-based estimate and the clothing-based estimate is statistically significant. There are two possibilities here: either the model is wrong or either clothing or food is not really assignable for children. In section V, we offer some speculative reasons as to why food is more likely to satisfy the restrictions of the model than clothing, for example, because of hand-me-downs (see also online app. A.5.4 in Dunbar, Lewbel, and Pendakur 2013).

Turning to the right-hand columns in the bottom rows, we consider the difference in the distribution of predicted values of resource shares between estimates based on food and those based on clothing. The mean values of resource shares are nearly identical for men and a bit different for women and children. Turning to the estimated standard deviation of resource shares, we see that for children, the estimated resource shares have greater dispersion in estimates based on food. But for men and women, the estimated standard deviations of estimates based on clothing and those based on food hardly differ at all.

The estimates based on food data versus clothing data are correlated but separately identified. That is, while the regressors in the reduced-form regressions based on food and those based on clothing are the same, the outcome variables are different, and no restrictions are imposed across the food versus clothing equations. Although the estimated resource shares are statistically significantly different (for children), the estimated magnitudes are quite close. For example, using either estimate tells us that children's resource shares are smaller than either men's or women's resource shares.

We take the similarity of resource shares across the two estimators as substantial evidence in favor of the model. But estimates such as these must always be digested with some caution: they are contingent on the modeling assumptions described above, which may be wrong. On the other hand, in the absence of direct information on resource shares (as in the dream data), these estimates might be the best we have.

C. 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 do not have an adult man and those that do not have an adult woman. 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 means and standard deviations of estimated resource shares evaluated at all values of the covariates.

TABLE 5
ESTIMATED RESOURCE SHARES AND GENDER GAPS, SELECTED COUNTRIES:
HOUSEHOLDS WITH BOTH MEN AND WOMEN PRESENT

| | | EVALUATED AT ALL z_h | | EVALUATED AT $\bar{z}$ | | GENDER GAP AT $\bar{z}$ | |
| | | Mean | | Estimate | | | |
COUNTRY	SAMPLE H	Men (SD)	Women (SD)	Men (SE)	Women (SE)	Estimate (SE)	Significance (p)
				Clothing			
Albania	3,279	.333 (.298)	.287 (.259)	.282 (.032)	.247 (.033)	.035 (.059)	
Bangladesh	5,427	.348 (.111)	.302 (.090)	.312 (.011)	.267 (.011)	.045 (.020)	<.05
Bulgaria	2,099	.312 (.150)	.447 (.212)	.304 (.038)	.372 (.041)	−.068 (.070)	
Iraq	13,805	.310 (.138)	.258 (.125)	.249 (.009)	.209 (.008)	.041 (.015)	<.01
Malawi	9,490	.362 (.167)	.279 (.135)	.312 (.028)	.253 (.029)	.059 (.054)	
				Bangladesh (with Valid Food Data)			
Clothing	4,391	.354 (.107)	.303 (.089)	.325 (.013)	.270 (.013)	.055 (.024)	<.05
Food	4,391	.334 (.115)	.248 (.105)	.309 (.014)	.236 (.013)	.073 (.024)	<.01
Difference: Clothing − Food		.020 −.008	.055 −.016	.016 (.019)	.034 (.018)	−.018 (.034)	

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 statistical significance. The bottom panel of table 5 gives statistics for clothing and food, using the sample of observations with valid food data. The bottom line gives the difference between estimates using food and those using clothing.

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. Consequently, the difference between them—the gender gap—is statistically indistinguishable from zero.

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 see a statistically significant gender gap only in Bangladesh and Iraq, and both of these show larger resource shares for men. The Bangladeshi clothing data (full sample) suggest a gender gap of 4.5 percentage points, and the Iraqi data suggest a gender gap of 3.5 percentage points.

In the Bangladeshi data, the estimated gender gaps from assignable clothing and food data are 5.5 and 7.3 percentage points, respectively. The similarity between the estimates of the gender gap coming from clothing data and food data is striking (and they are not statistically significantly different from each other) and provides more evidence that the methodology is valid.

Estimated resource shares are functions of the covariates z (ages and education levels of household members). The left-hand columns give the means and standard deviations of estimated resource shares of men and women over the observed values of covariates. The standard deviations of estimated resource shares are large relative to their respective means. For example, our estimates for Bangladesh using food as the assignable good show standard deviations of roughly 11 percentage points for men's and women's per-person resource shares.

Our methodology makes it easy to see how resource shares depend on covariates. Figure 1 shows a scatter plot of 4,391 estimates of Bangladeshi women's resource shares based on assignable food data in households with at least one man and one woman present, plotted against the household budget measured in 2011 US dollars (on a log scale). Each estimate is a function of demographic covariates, given by equation (13). Here, we separate households into those with one woman (black dots) and those with two or more women (gray dots). From the figure, we can see several patterns. First, households with just one woman have larger per-woman

F<small>IG</small>. 1.—Estimated women's resource shares in Bangladesh (food based).

resource shares than households with two or more women (roughly twice the size), suggesting that women crowd each other out when more than one are present. Second, there is considerable heterogeneity in resource shares conditional on total household expenditure, suggesting that, even if the household budget is right at the poverty line, there may be many poor and nonpoor women. Third, there is a mild negative correlation with total household expenditure (Pearson's $r \sim -0.3$). Recall that our identification strategy assumes that, conditional on covariates, resource shares do not vary with household budgets. The unconditional correlation we observe is driven by age and education, which are positively correlated with household budgets and negatively correlated with resource shares.

D. Individual Poverty

The per capita approach, a standard approach to poverty measurement used by the World Bank and other international organizations, assumes equal resource shares. In this approach, per capita household budgets are compared to a person-level poverty threshold. In table 6, we use the extreme-poverty threshold of $1.90 per day and two other "societal poverty" thresholds of $3.20 and $5.50 per day.

The per capita approach does not account for scale economies in consumption. Because our measure of resource shares simply divides the pie differently than the per capita approach, our poverty rates are directly

TABLE 6
ESTIMATED POVERTY RATES

| | $1.90 PER PERSON PER DAY | | | | | | | $1.90, $3.20 OR $5.50 PER PERSON PER DAY | | | | | |
| | Per Capita Estimate | | Estimated Using Resource Shares | | | | | Per Capita Estimate | Estimated Using Resource Shares | | | | |
COUNTRY (INCOME CLASS)	World Bank (year)	Our Data (SE)	Men (SE)	Women (SE)	Children (SE)	All (SE)		Our Data (SE)	Men (SE)	Women (SE)	Children (SE)	All (SE)
Albania	.003	.003	.056	.035	.081	.054		.319	.282	.392	.471	.369
(Upper middle)	(2008)	(.001)	(.042)	(.041)	(.076)	(.031)		(.008)	(.051)	(.054)	(.082)	(.021)
Bangladesh	.143	.109	.020	.059	.284	.131		.464	.252	.338	.631	.421
(Lower middle)	(2016)	(.004)	(.008)	(.011)	(.052)	(.018)		(.006)	(.026)	(.022)	(.034)	(.008)
Bulgaria	.056	.003	.028	.026	.276	.063		.080	.185	.157	.416	.206
(Upper middle)	(2006)	(.001)	(.041)	(.031)	(.119)	(.031)		(.005)	(.058)	(.035)	(.11)	(.031)
Iraq	.015	.000	.000	.001	.116	.042		.074	.017	.052	.680	.265
(Upper middle)	(2006)	(.000)	(.001)	(.002)	(.066)	(.024)		(.002)	(.005)	(.012)	(.058)	(.020)
Malawi	.711	.629	.469	.586	.727	.627		.629	.469	.586	.727	.627
(Lower)	(2010)	(.004)	(.025)	(.035)	(.030)	(.009)		(.004)	(.025)	(.035)	(.03)	(.009)
Bangladesh (with Valid Food Data; Lower-Middle Income)												
Clothing		.103	.016	.065	.224	.111		.454	.277	.342	.581	.400
		(.004)	(.008)	(.014)	(.053)	(.019)		(.007)	(.028)	(.026)	(.041)	(.011)
Food		.103	.044	.134	.107	.098		.454	.266	.470	.338	.362
		(.004)	(.016)	(.025)	(.023)	(.011)		(.007)	(.031)	(.024)	(.035)	(.009)

comparable to the per capita poverty rates. Our approach can accommodate any scale economies, in particular an absence of scale economies, which is the assumption underlying the per capita approach. Thus, the poverty rates here could be taken as an upper bound on actual poverty rates. We use this approach here to highlight the contrast with the standard per capita approach. (In sec. A3, we provide estimates of poverty rates that assume scale economies following OECD [Organization for Economic Co-operation and Development] standard methods.) Our poverty rates differ from, and are larger than, the per capita rates because our approach allows for inequality within households, so that some members may be poor even if the household budget exceeds the per capita threshold.

In figure 2, we take the subsample of 764 households from figure 1 with exactly one man, one woman, and two children. We add a poverty line of $4,672 ($3.20 per person per day), shown with a thick vertical line. If resource shares were equal, women in households with less than this level of expenditure would be poor. We also add, with a black curve, the critical value of the woman's resource share necessary to keep the woman out of poverty, at any given household expenditure level.

With very low women's resource shares, even households of substantial means may have women members in poverty. We note that variation in resource shares is especially important to individual poverty measurement


Fig. 2.—Women's resource shares and poverty, Bangladesh.

when there are many households near the poverty line. The figure makes clear that there are two types of misclassification that arise when we use the per capita method in the presence of unequal resource shares: some women are classified as poor even though they have a personal budget $\eta^f y$ that exceeds the poverty threshold (top-left region), and some women are classified as nonpoor even though their personal budget is below the poverty threshold (bottom-right region).

In table 6, we present poverty rates. In the left-hand block, we use the extreme-poverty threshold of \$1.90 per person per day (evaluated in 2011 real dollars, PPP exchange rates). In the right-hand block, we use the "societal poverty thresholds" advocated by World Bank (2018). They suggest poverty lines 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. Since the societal poverty line equals the extreme-poverty threshold for Malawi, the figures in those rows are the same in the left- and right-hand blocks.

In the leftmost column, we give the published estimated extreme-poverty rate from the World Bank Development Indicators database (World Bank 2021, series SI.POV.DDAY) for the available year closest to our survey data. In the next column ("Our Data"), we compute for each person in the household their per capita expenditure, y_h/N_h, compare this to the poverty line, and report the poverty rate. The World Bank estimates count households as poor if their income is below the threshold. In contrast, our estimates compare expenditures (including imputed expenditure for home-produced food) to the poverty threshold. Thus, our estimates differ from the World Bank estimates, even when they are based on the same survey data (as in Albania, Iraq, and Malawi). However, they are reasonably close to each other.

In the next three columns, we use our estimated resource share estimates rather than the per capita share. In 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 compute, for each man, woman, and child in the data set, $y_h \eta_h^t/N_h^t$, compare this to the poverty line, and report the poverty rate. Like Dunbar, Lewbel, and Pendakur (2013), we use a poverty line 40% lower for children. In the next column ("All"), we report the overall poverty rate, at the person level and using our resource shares for the entire sample. The right-hand block does the same exercise but compares individual expenditure $y_h \eta_h^t/N_h^t$ to the societal poverty thresholds of \$3.20 per person per day for Bangladesh and \$5.50 per person per day in Albania, Bulgaria, and Iraq. We provide asymptotic standard errors, computed via the bootstrap.[24]

---

[24] 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

Looking first at the extreme-poverty measures, 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. From estimates based on clothing, in Bangladesh, women are 4 percentage points more likely than men to experience extreme poverty; in Malawi, they are 12 percentage points more likely. In Iraq, which is classified as a middle-income country, essentially no men or women experience extreme poverty.

Turning to the right-hand panel of the table, we use higher poverty thresholds for all countries except Malawi, corresponding to the World Bank's (2018) prescriptions for "societal poverty lines." In Bangladesh, if we use a poverty threshold of $3.20 per person per day, we find that 25.2% of men and 33.8% of women are poor, a gender gap of 8.6 percentage points. In Iraq, if we use a poverty threshold of $5.50 per person per day, we find that 1.7% of men and 5.2% of women are poor, a gender gap of 3.5 percentage points.

When we turn to child poverty, one key feature pops out. Estimated child poverty rates are higher, sometimes much higher, than estimated poverty rates for either adult men or adult women. This is due to the fact that estimated per-child resource shares are below $1/N_h$ in every country.

Part of this pattern is driven by the assumption, used in table 6, that there are no scale economies in household consumption. In section A3, we present poverty estimates constructed using the same estimated resource shares and allowing for scale economies in household consumption via a standard OECD scale economy adjustment. There, we show that estimated children's poverty rates are still higher than those of adults, but not by nearly as much.

The estimated levels of child poverty are especially high in Iraq, where the very low estimated per-child resource share of 5% corresponds to the very high estimated extreme-poverty rate of 12%. When we use the societal poverty threshold of $5.50 per person per day, we see 68% of children falling below the threshold. As noted above, there are classification error issues in the Iraqi data, which we think render the child poverty estimates for Iraq invalid, as detailed in section A4. However, we also show in section A4 that the gender gap in poverty in Iraq is evident in a subsample where such classification error is minimized. In that adults-only subsample, we find a gender gap in poverty of 3.9 percentage points, which is similar to what we see in table 6.

_____

nonlinear functions of estimated OLS regression coefficients. For an alternative, see Woutersen and Ham (2013).

Turning to the bottom panel of table 6, we see the difference between estimated poverty rates across Bangladeshi resource shares based on assignable clothing data and those based on assignable food data. Focusing on the right-hand panel (poverty line of $3.20 per person per day), the estimated poverty rate for men is roughly 27% regardless of which assignable good is used to identify resource shares. Further, we see that estimated women's and children's poverty rates are higher than estimated men's poverty rates regardless of which assignable good is used. But the estimates based on clothing imply higher children's poverty than do the estimates based on food.

These differences in estimated poverty rates are driven by the differences in estimated resource shares described in tables 4 and 5. In that discussion, we argued that estimated resource shares are pretty close to each other across the two assignable goods. But small differences in resource shares can get magnified in poverty estimates, especially if (1) there are many households with budgets near the poverty line and/or (2) there is a lot of heterogeneity in resource shares for households near the poverty line.

V.   Discussion

We provide evidence of substantial within-household expenditure inequality. This suggests that the current standard practice for poverty measurement in developing countries—asking whether or not per capita household income or expenditure falls below a threshold—can be misleading. This current practice ignores within-household inequality and so mischaracterizes poverty rates. For example, if a household has income slightly above the poverty line, then by the per capita method we would call it nonpoor, 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 five countries for which we estimate resource shares, we see statistically significant gender gaps in resource shares that favor men over women in two countries, and we see no statistically significant evidence of gender gaps that favor women. Further, these gender 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 (i.e.,

the 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 one or two categories of goods and services that can be measured well and 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, Lacroix, and Tiberti 2020).

This recommendation applies similarly to field experiments where an outcome variable is individual poverty or consumption. If information on total household expenditure (or consumption) is already being gathered, this may require only adding a few questions to a questionnaire. With total household expenditure and assignable-goods expenditure in hand, field experimentalists and survey designers can add resource shares and therefore within-household inequality to their list of interesting outcome variables. Understanding the determinants of intrahousehold inequality and its consequences can help us understand better a wide variety of phenomena concerning individuals and help design better policies.

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% of the household budget, and food, which is roughly 56% 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 2020). However, the use of clothing is due to its availability in public-use data sets, 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 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 know only expenditure. There may thus be more unobserved heterogeneity in clothing than in food. Third, food budget shares are known to be downward sloping (e.g., Engel 1857) and therefore satisfy the identifying restriction of our model. Fourth, clothing is much more durable than food. Consequently, observed clothing expenditure may not equal its flow value, because of 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 auxiliary feature. All together, this suggests that the collection of person-level food consumption is desirable, even if costly, given the benefits it brings to our understanding of individual outcomes.

Appendix

A1.   Dream Data with Scale Economies

Table A1A presents the dream data in terms of expenditure, which we discussed in section II.A. Table A1B presents the corresponding dream data in terms of quantities consumed by each person. Normalize market prices to 1 for all goods, so that we can think of consumed quantities as measured in dollars.

TABLE A1
DREAM DATA

	Man	Woman	Child	Total	Total Expenditure
	A. Expenditures				
Food	400	300	200	900	
Clothing	50	75	25	150	
Shelter	100	100	100	300	
Other	250	125	125	500	
Total	800	600	450	1,850	
Resource shares (%)	43	32	24	100	
	B. Quantities				
Food	400	300	200	900	900
Clothing	50	75	25	150	150
Shelter	300	300	300	900	300
Other	500	250	250	1,000	500
Total	1,250	925	775	2,950	1,850
Consumption shares (%)	43	31	26	100	

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 had to purchase only $300 of housing to accomplish this. This means that the household purchased only one-third of the total housing consumption of the three 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 A from these data: the element of A corresponding to shelter is $1/3$, because the household needs to buy only one-third of the total consumed quantities of all the members.

Goods do not have to be either fully shared or nonshareable in the BCL model or here; 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 A1B are the individual-level numbers of kilometers 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 that the man is the only member who knows how to drive a motorcycle. If the man rode 250 km with the woman and 250 km with the child, then their consumed quantities would be as in table A1B, with 1,000 person-km driven. But the motorcycle would have traveled only 500 km, so the household would have purchased only 500 km. Here, the element of A corresponding to transport (other) would be 1/2. The value of the sum of quantities consumed at market prices ($2,950) is greater than total expenditure ($1,850) because of shared goods.

In table A1A, individual-level expenditures are obtained by multiplying quantities by shadow prices. Since market prices p are normalized to 1 within household prices given by Ap, this means that we multiply quantities by the diagonal matrix A. Since nonshareable goods have an element of A equal to 1, for the nonshareable goods of food and clothing, the corresponding rows of tables A1A and A1B are identical. The elements of 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 A1A, 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 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 p and facing a price vector Ap. Browning, Chiappori, and Lewbel (2013) show how to identify resource shares and the matrix 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).

Dunbar, Lewbel, and Pendakur (2013) do not attempt to identify the matrix 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 nonshareable. The DLP model does not make any assumptions about how shareable the nonassignable goods are. In terms of the matrix A, the DLP model assumes that the single element of A corresponding to the nonshareable assignable good equals 1 and make no assumptions about the other elements of A.

Although the DLP model 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 $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 $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.

A2. 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.[25] 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 $y_h/(N_h)^{1/2}$, 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 $(N_h)^{1/2}$ and then dividing equally among household members, assigning $(N_h)^{1/2}y_h/N_h = y_h/(N_h)^{1/2}$ to each member. In table A2, we pursue this approach using our resource shares: we first inflate the household budget by $(N_h)^{1/2}$ to account for scale economies in consumption and then multiply by the resource shares to assign a consumption level $(N_h)^{1/2}y_h\eta_h^t$ to each member. In table A2, we give extreme-poverty rates using this scale-economy adjustment. These estimates, which use a $1.90 per person per day poverty threshold, are comparable to those in table 6, in section IV.D.

TABLE A2
ESTIMATED POVERTY RATES WITH SCALE ECONOMIES

| | $1.90 PER PERSON PER DAY | | | | | | $1.90, $3.20, OR $5.50 PER PERSON DAY | | | | | |
| | Per Capita Estimate | | Estimated Using Resource Shares | | | | Per Capita Estimate | Estimated Using Resource Shares | | | |
COUNTRY (INCOME CLASS)	World Bank (Year)	Our Data (SE)	Men (SE)	Women (SE)	Children (SE)	All (SE)	Our Data (SE)	Men (SE)	Women (SE)	Children (SE)	All (SE)
Albania (Upper middle)	.003 (2008)	.000 (.000)	.037 (.042)	.020 (.041)	.058 (.076)	.035 (.031)	.016 (.002)	.069 (.046)	.066 (.046)	.111 (.082)	.078 (.033)
Bangladesh (Lower middle)	.143 (2016)	.003 (.001)	.000 (.000)	.003 (.003)	.007 (.010)	.004 (.004)	.048 (.003)	.007 (.002)	.034 (.007)	.112 (.040)	.055 (.015)
Bulgaria (Upper middle)	.056 (2006)	.000 (.003)	.015 (.036)	.004 (.028)	.244 (.119)	.043 (.030)	.022 (.003)	.058 (.048)	.052 (.034)	.316 (.119)	.092 (.033)
Iraq (Upper middle)	.015 (2006)	.000 (.000)	.000 (.001)	.001 (.001)	.008 (.021)	.004 (.009)	.000 (.000)	.000 (.000)	.002 (.002)	.130 (.070)	.056 (.030)
Malawi (Lower)	.711 (2010)	.238 (.004)	.192 (.029)	.228 (.040)	.298 (.047)	.253 (.022)	.238 (.004)	.192 (.029)	.228 (.040)	.298 (.047)	.253 (.022)
Bangladesh (with Valid Food Data; Lower-Middle Income)											
Clothing		.001 (.000)	.000 (.001)	.004 (.004)	.007 (.013)	.004 (.005)	.040 (.003)	.009 (.003)	.036 (.009)	.074 (.039)	.042 (.014)
Food		.001 (.000)	.006 (.007)	.014 (.012)	.040 (.013)	.021 (.008)	.040 (.003)	.023 (.011)	.060 (.021)	.074 (.018)	.055 (.011)

The poverty rates shown in table A2 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 $(N_h)^{1/2}$ times the poverty threshold of $5.50 per person per day, resulting in an estimated poverty rate of zero. 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 is unchanged by allowing for varying poverty thresholds or by 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.

A3. Quadratic Engel Curves

Consider first the case with no demographic variation and a fixed household composition, as in section II.B. Let indirect utility for type t be given by the QAI model of Banks, Blundell, and Lewbel (1997):

$$V^t(\boldsymbol{p}, y) = \left[\left(\frac{\ln y - \ln \tilde{\alpha}^t(\boldsymbol{p})}{\tilde{\beta}^t(\boldsymbol{p})} \right)^{-1} - \tilde{\gamma}^t(\boldsymbol{p}) \right]^{-1}$$

for some homogeneous of degree 1 function $\tilde{\alpha}^t(\boldsymbol{p})$ and homogeneous of degree 0 functions $\tilde{\beta}^t(\boldsymbol{p})$ and $\tilde{\gamma}^t(\boldsymbol{p})$. Then, the vector of budget-share functions are given by Roy's identity as

$$\boldsymbol{w}^t(\boldsymbol{p}, y) = \nabla_{\ln \boldsymbol{p}} \ln \tilde{\alpha}^t(\boldsymbol{p}) + \nabla_{\ln \boldsymbol{p}} \ln \tilde{\beta}^t(\boldsymbol{p})(\ln y - \ln \tilde{\alpha}^t(\boldsymbol{p}))$$
$$+ \frac{\nabla_{\ln \boldsymbol{p}} \tilde{\gamma}^t(\boldsymbol{p})}{\tilde{\beta}^t(\boldsymbol{p})} (\ln y - \ln \tilde{\alpha}^t(\boldsymbol{p}))^2.$$

Let $w^t(y)$ be the assignable-good element of $\boldsymbol{w}^t(\boldsymbol{p}_0, y)$ at fixed price vector $\boldsymbol{p}_0$. These scalar-valued Engel curve functions for the assignable good of person t are given by

$$w^t(y) = \alpha^t + \beta^t \ln y + \gamma^t (\ln y)^2, \tag{A1}$$

where the scalar-valued coefficients α^t, β^t, and γ^t are functions of $\tilde{\alpha}^t(\boldsymbol{p}_0)$, $\tilde{\beta}^t(\boldsymbol{p}_0)$, and $\tilde{\gamma}^t(\boldsymbol{p}_0)$, respectively:

$$\alpha^t = \nabla_{\ln \boldsymbol{p}} \ln \tilde{\alpha}^t(\boldsymbol{p}_0) - \nabla_{\ln \boldsymbol{p}} \ln \tilde{\beta}^t(\boldsymbol{p}_0) \ln \tilde{\alpha}^t(\boldsymbol{p}_0) + \ln \tilde{\alpha}^t(\boldsymbol{p}_0)^2 \frac{\nabla_{\ln \boldsymbol{p}} \tilde{\gamma}^t(\boldsymbol{p}_0)}{\tilde{\beta}^t(\boldsymbol{p}_0)},$$

$$\beta^t = \nabla_{\ln \boldsymbol{p}} \ln \tilde{\beta}^t(\boldsymbol{p}_0) - 2 \ln \tilde{\alpha}^t(\boldsymbol{p}_0) \frac{\nabla_{\ln \boldsymbol{p}} \tilde{\gamma}^t(\boldsymbol{p}_0)}{\tilde{\beta}^t(\boldsymbol{p}_0)}, \quad \text{and}$$

$$\gamma^t = \frac{\nabla_{\ln \boldsymbol{p}} \tilde{\gamma}^t(\boldsymbol{p}_0)}{\tilde{\beta}^t(\boldsymbol{p}_0)}.$$

This quadratic Engel curve system satisfies the SAP of Dunbar, Lewbel, and Pendakur (2013) if either $\gamma^t = \gamma$ and $\gamma \neq 0$ or $\beta^t = \beta$ and $\gamma^t = 0$. This condition is satisfied if $\tilde{\beta}^t(\boldsymbol{p}) = \tilde{\beta}(\boldsymbol{p})$ and $\tilde{\gamma}^t(\boldsymbol{p}) = \tilde{\gamma}(\boldsymbol{p})$.

Suppose that (1) the matrix $\boldsymbol{A}$ is block diagonal, as in equation (2); (2) resource shares do not depend on the household budget, so that $\eta^t(y) = \eta^t$; (3) individual Engel curve functions are QAI, as in equation (A1); and (4) preferences satisfy the SAP such that $\gamma^t = \gamma$. Substituting into equation (3) gives

$$W^t(y) = \eta^t \left[\alpha^t + \beta^t (\ln y + \ln \eta^t - \ln N^t) + \gamma(\ln y + \ln \eta^t - \ln N^t)^2 \right].$$

Rewrite with a subscript h indexing households and an additive error term ε_h^t:

$$W_h^t = a^t + b^t \ln y_h + c^t (\ln y_h)^2 + \varepsilon_h^t,$$

where

$$a^t = \eta^t \left[\alpha^t + \beta^t (\ln \eta^t - \ln N^t) + \gamma(\ln \eta^t - \ln N^t)^2 \right],$$
$$b^t = \eta^t [\beta^t + 2\gamma(\ln \eta^t - \ln N^t)], \quad \text{and}$$
$$c^t = \eta^t \gamma.$$

Analogously to the case with linear Engel curves, we can estimate resource shares as

$$\hat{\eta}^t = \frac{\hat{c}^t}{\sum_t \hat{c}^t}.$$

Estimated quadratic terms (c^t) are likely to have larger estimated standard errors than estimated linear terms (b^t; for an empirical example of this, see Banks, Blundell, and Lewbel 1997). Thus, resource shares estimated from quadratic Engel curves may have large standard errors. Thus, we recommend the linear model presented in the body of the paper, if the linear model is thought to be acceptable.

Now consider the case with variation in observed demographic characteristics and household composition. Add dependence on $\boldsymbol{z}$ to the structural parameters η^t, α^t, β^t, and γ:

$$W^t(y, z) = \eta^t(z) \left[\alpha^t(z) + \beta^t(z)(\ln y + \ln \eta^t(z) - \ln N^t) + \gamma(z)(\ln y + \ln \eta^t(z) - \ln N^t)^2 \right].$$

Rewrite with a subscript h indexing households and an additive error term ε_h^t:

$$W_h^t = a_h^t + b_h^t \ln y_h + c_h^t (\ln y_h)^2 + \varepsilon_h^t,$$

where

$$a_h^t = \eta^t(\boldsymbol{z}_h) \left[\alpha^t(\boldsymbol{z}_h) + \beta^t(\boldsymbol{z}_h)(\ln \eta^t(\boldsymbol{z}_h) - \ln N_h^t) + \gamma(\boldsymbol{z}_h)(\ln \eta^t(\boldsymbol{z}_h) - \ln N_h^t)^2 \right],$$
$$b_h^t = \eta^t(\boldsymbol{z}_h) [\beta^t(\boldsymbol{z}_h) + 2\gamma(\boldsymbol{z}_h)(\ln \eta^t(\boldsymbol{z}_h) - \ln N_h^t)], \quad \text{and}$$
$$c_h^t = \eta^t(\boldsymbol{z}_h) \gamma(\boldsymbol{z}_h).$$

Analogously to the case with linear Engel curves, we can approximate a_h^t, b_h^t, and c_h^t with, respectively, equations (11) and (12) and

$$c_h^t = c_0^t + \boldsymbol{c}_z^{t\prime} \boldsymbol{z}_h.$$

Finally, an estimate of the resource share is given by

$$\hat{\eta}_h^t = \frac{\hat{c}_h^t}{\sum_t \hat{c}_h^t}.$$

A4. Misclassification of Children in Iraq

In the expenditure module of the Iraqi Household Socio-Economic Survey 2006–2007 (IHSES), respondents are asked to recall spending on men's clothing, women's clothing, and children's clothing, in 12 categories, over the previous 3 months. However, unlike in other LSMS household surveys, they are not given an age range for children versus adults. (In most other LSMS surveys, the age range is 15 or less for children and 16 or more for adults.) In other modules of IHSES, questions regarding children specify age ranges of under 1 (ration cards), under 5 (breastfeeding, vaccination), under 6 (language, schooling, labor force activity), and under 12 (marital status, births). This suggests that, without guidance in the expenditure module, households could reasonably classify their members as children if they are under 6 or if they are under 12.

We can use data to understand how households classify people as children or not by examining the pattern of purchasing versus not purchasing children's clothing across households with different age compositions. Table A3 gives the fraction of households with zero expenditures on children's, men's, and women's clothing by the minimum age of household members for households that contain at least one adult male over 15 years old and at least one adult female over 15 years old.

TABLE A3
PROBABILITY OF REPORTING ZERO CLOTHING EXPENDITURE, BY TYPE

Minimum Age of Members	Observations	Children	Men	Women
0	3,515	.06	.03	.03
1	2,724	.05	.04	.04
2	1,899	.05	.05	.04
3	1,259	.05	.04	.03
4	936	.05	.07	.03
5	766	.09	.05	.03
6	559	.16	.05	.04
7	487	.22	.04	.03
8	450	.35	.07	.03
9	367	.35	.06	.04
10	328	.47	.06	.04
11	296	.59	.04	.04
12	277	.73	.07	.03
13	260	.78	.04	.03
14	266	.89	.05	.02
15	265	.92	.05	.04

All these households have at least one adult man and at least one adult woman. Here, we see that for adults, only about 5% of households report spending exactly zero on clothing for men or women. Similarly, for households with children aged under 6, about 5% of households report spending exactly zero on

children's clothing. This is consistent with a world where all households have roughly the same rate of purchase frequency for all types of people (95% over 3-month recall) and where all respondents classify members under 6 as children. As we move to households where the youngest member is older than 5, we see the frequency of nonpurchase rising smoothly to roughly 90% for households with some 15-year-old members but no members under 15 years old.

If all households have roughly the same rate of purchase frequency for all types of people, then nonpurchase exceeding 50% signals that more than half of households have classified their youngest members as adults. This switch occurs at age 11 in these data. Since classification into adulthood is a binary choice, we think that it is therefore reasonable to classify members under 12 as children. Roughly speaking, more than half of respondents would agree with this classification.

We then have two choices. First, we can define the number of children in Iraqi households as the number of members less than 12 and proceed as usual, recognizing that there is measurement error present because of the substantial numbers of households that would report clothing expenditures on persons aged 6–11 as adult clothing expenditures. This is what we do in the main text. Second, we can exclude all households with members under 12 years old, call all persons remaining as adults, and consider gender disparity only, as below.

In the main text, we allow for households with as many as 14 members (four men, four women, and six children). In the exercise below, we include households with up to seven men and seven women (14 members total) where all members are aged 12 or more. In table A4, we give numbers analogous to those in tables 2–4 and 6. There is no need to reproduce table 5, because for households with men and women only, it is identical to table 4.

TABLE A4

RESULTS FOR IRAQI ADULTS AGED 12 AND UP

TABLE 2: DESCRIPTIVE STATISTICS

	Compositions	Our H	Budget	SD
Total H				
17,513	mw	3,385	26,360	15,114

TABLE 3: TEST OF IDENTIFICATION

	Budget Share				
	Mean	SD	Slope at $\bar{z}$	t-Test of Slope	% of Sample Significant
Our H					
3,385	.065	.046	.022	14.6	99.6

TABLE 4: PREDICTED RESOURCE SHARES

	Predicted Resource Shares						
	Evaluated at $\bar{z}$			Evaluated at all z_h			
	Men (SE)	Women (SE)	Gap (SE)	Men (SD)	Women (SD)	% outside [0,1]	Wald test stat (p-value)
Sample Size							
3,385	.224 (.009)	.192 (.009)	.033 (.019)	.199 (.126)	.179 (.113)	0	33.3 (.0009)

TABLE 6: POVERTY RATES

	$1.90 per Person per Day				$5.50 per Person per Day			
	Per Capita	Use Resource Shares			Per Capita	Use Resource Shares		
	Our Data (SE)	Men (SE)	Women (SE)	All (SE)	Our Data (SE)	Men (SE)	Women (SE)	All (SE)
	0 (.000)	0 (.005)	0 (.002)	0 (.003)	.027 (.003)	.034 (.012)	.073 (.02)	.053 (.01)

The key messages from our work are evident in these results for a two-type (men and women) model in Iraq estimated on the sample of households that have no members under 12. First, there is a (mildly) statistically significant gender gap in resource shares: men's per-person resource shares are 22.4%; women's per-person resource shares are 19.2%. Note that these resource shares are far below 0.5 because there are five members in typical households, with an average of 2.4 men and 2.4 women in each household. The estimated gender gap in per-person resource shares is 3.3 percentage points, and it is mildly statistically significant, with a t-statistic of 1.76. Relatedly, the per capita model (in which there is no gender gap) is rejected. Finally, the gender gap in resource shares implies a gender gap in poverty. When using a poverty threshold of $5.50 per person per day, we see a poverty rate of 3.4% among men and poverty rate of 7.3% among women.

References

Alkire, Sabina, Ruth Meinzen-Dick, Amber Peterman, Agnes Quisumbing, Greg Seymour, and Ana Vaz. 2013. "The Women's Empowerment in Agriculture Index." *World Development* 52:71–91.

Anderson, Siwan, and Debraj Ray. 2010. "Missing Women: Age and Disease." *Rev. Econ. Studies* 77 (4): 1262–300.

Attanasio, Orazio, and Valérie Lechene. 2014. "Efficient Responses to Targeted Cash Transfers." *J.P.E.* 122 (1): 178–222.

Banks, James, Richard Blundell, and Arthur Lewbel. 1997. "Quadratic Engel Curves and Consumer Demand." *Rev. Econ. and Statis.* 79 (4): 527–39.

Bargain, Olivier, Olivier Donni, and Prudence Kwenda. 2014. "Intrahousehold Distribution and Poverty: Evidence from Côte d'Ivoire." *J. Development Econ.* 107:262–76.

Bargain, Olivier, Guy Lacroix, and Luca Tiberti. 2020. "Validating the Collective Model of Household Consumption Using Direct Evidence on Sharing." Working Paper no. 2018–06, Partnership Econ. Policy, Nairobi.

Becker, Gary S. 1981. *A Treatise on the Family.* Cambridge, MA: Harvard Univ. Press.

Blackorby, Charles, and David Donaldson. 1993. "Household Equivalence Scales and Welfare Comparisons: A Comment." *J. Public Econ.* 50 (1): 143–46.

Blundell, Richard, Xiaohong Chen, and Dennis Kristensen. 2007. "Seminonparametric IV Estimation of Shape-Invariant Engel Curves." *Econometrica* 75 (6): 1613–69.

Brown, Caitlin, Martin Ravallion, and Dominique van de Walle. 2019. "Most of Africa's Nutritionally Deprived Women Are Not Found in Poor Households." *Rev. Econ. and Statis.* 101 (4): 631–44.

Browning, Martin, Pierre-André Chiappori, and Arthur Lewbel. 2013. "Estimating Consumption Economies of Scale, Adult Equivalence Scales, and Household Bargaining Power." *Rev. Econ. Studies* 80 (4): 1267–303.

Calvi, Rossella. 2020. "Why Are Older Women Missing in India? The Age Profile of Bargaining Power and Poverty." *J.P.E.* 128 (7): 2453–501.

Calvi, Rossella, Jacob Penglase, Denni Tommasi, and Alexander Wolf. 2020. "The More the Poorer? Resource Sharing and Scale Economies in Large Families." IZA Discussion Paper no. 13948, Inst. Labor Econ., Bonn.

Campbell, Frances, Gabriella Conti, James J. Heckman, et al. 2014. "Early Childhood Investments Substantially Boost Adult Health." *Science* 343 (6178): 1478–85.

Cherchye, Laurens, Thomas Demuynck, Bram De Rock, and Frederic Vermeulen. 2017. "Household Consumption When the Marriage Is Stable." *A.E.R.* 107 (6): 1507–34.

Cherchye, Laurens, Bram De Rock, Arthur Lewbel, and Frederic Vermeulen. 2015. "Sharing Rule Identification for General Collective Consumption Models." *Econometrica* 83 (5): 2001–41.

Cherchye, Laurens, Bram De Rock, and Frederic Vermeulen. 2007. "The Collective Model of Household Consumption: A Nonparametric Characterization." *Econometrica* 75 (2): 553–74.

Chiappori, Pierre-André. 1988. "Nash-Bargained Households Decisions: A Comment." *Internat. Econ. Rev.* 29 (4): 791–96.

———. 1992. "Collective Labor Supply and Welfare." *J.P.E.* 100 (3): 437–67.

Chiappori, Pierre-André, and Ivar Ekeland. 2009. "The Microeconomics of Efficient Group Behavior: Identification." *Econometrica* 77 (3): 763–99.

Deaton, Angus, and John Muellbauer. 1980. *Economics and Consumer Behavior.* Cambridge: Cambridge Univ. Press.

De Vreyer, Philippe, and Sylvie Lambert. 2016. "Intra-household Inequalities and Poverty in Senegal." Manuscript, Paris School Econ.

Donni, Olivier, and José Alberto Molina. 2018. "Household Collective Models: Three Decades of Theoretical Contributions and Empirical Evidence." IZA Discussion Paper no. 11915, Inst. Labor Econ., Bonn.

Dunbar, Geoffrey R., Arthur Lewbel, and Krishna Pendakur. 2013. "Children's Resources in Collective Households: Identification, Estimation, and an Application to Child Poverty in Malawi." *A.E.R.* 103 (1): 438–71.

———. 2021. "Identification of Random Resource Shares in Collective Households without Preference Similarity Restrictions." *J. Bus. and Econ. Statis.* 39 (2): 404–21.

Engel, Ernst. 1857. "Die Productions- und Consumtionsverhältnisse des Königreichs Sachsen." *Zeitschrift des Statistischen Bureaus des Königlich-Sächsischen, Ministerium des Innern* 8 (9): 1–54. Reprinted in 1895 as *Anlage* (appendix) I in "Die Lebenskosten belgischer Arbeiter-Familien früher und jetzt." *Bulletin de l'Institut International de Statistique* 9 (1). Rome: Imprimerie Nationale.

———. 1895. "Die Lebenskosten belgischer Arbeiter-Familien früher und jetzt." *Bulletin de l'Institut International de Statistique* 9 (1): 1–74. Rome: Imprimerie Nationale.

Ewerling, Fernanda, John W. Lynch, Cesar G. Victora, Anouka van Eerdewijk, Marcelo Tyszler, and Aluisio J. Barros. 2017. "The SWPER Index for Women's Empowerment in Africa: Development and Validation of an Index Based on Survey Data." *Lancet Global Health* 5 (9): e916–e923.

Han, Sukjin, and Adam McCloskey. 2019. "Estimation and Inference with a (Nearly) Singular Jacobian." *Quantitative Econ.* 10 (3): 1019–68.

Jayachandran, Seema, and Rohini Pande. 2017. "Why Are Indian Children So Short? The Role of Birth Order and Son Preference." *A.E.R.* 107 (9): 2600–2629.

Lewbel, Arthur. 1989. "Household Equivalence Scales and Welfare Comparisons." *J. Public Econ.* 39 (3): 377–91.

Lewbel, Arthur, and Krishna Pendakur. 2008. "Estimation of Collective Household Models with Engel Curves." *J. Econometrics* 147 (2): 350–58.

———. 2021. "Inefficient Collective Households: Cooperation and Consumption." *Econ. J.*, forthcoming.

Lindahl, Erik. (1919) 1958. "Just Taxation—A Positive Solution." In *Classics in the Theory of Public Finance*, edited by Richard A. Musgrave and Alan T. Peacock, 168–76. London: Macmillan.

Menon, Martina, Krishna Pendakur, and Federico Perali. 2012. "On the Expenditure-Dependence of Children's Resource Shares." *Econ. Letters* 117 (3): 739–42.

Muellbauer, John. 1974. "Household Production Theory, Quality, and the 'Hedonic Technique'." *A.E.R.* 64 (6): 977–94.

———. 1975. "Aggregation, Income Distribution and Consumer Demand." *Rev. Econ. Studies* 42 (4): 525–43.

Pendakur, Krishna. 1999. "Semiparametric Estimates and Tests of Base-Independent Equivalence Scales." *J. Econometrics* 88 (1): 1–40.

———. 2018. "Welfare Analysis When People Are Different." *Canadian J. Econ.* 51 (2): 321–60.

Pulerwitz, Julie, Steven L. Gortmaker, and William DeJong. 2000. "Measuring Sexual Relationship Power in HIV/STD Research." *Sex Roles* 42 (7–8): 637–60.

Rangel, Marcos, and Duncan Thomas. 2019. "Decision-Making in Complex Households." Working Paper no. 26511 (November), NBER, Cambridge, MA.

Staiger, Douglas, and James H. Stock. 1997. "Instrumental Variables Regression with Weak Instruments." *Econometrica* 65 (3): 557–86.

World Bank. 2018. *Piecing Together the Poverty Puzzle.* Poverty and Shared Prosperity Series. Washington, DC: World Bank. https://www.worldbank.org/en/publication/poverty-and-shared-prosperity-2018.

———. 2021. "World Development Indicators" [database]. https://databank.worldbank.org/home.

Woutersen, Tiemen, and John C. Ham. 2013. "Calculating Confidence Intervals for Continuous and Discontinuous Functions of Parameters." CEMMAP Working Paper no. CWP23/13, Centre Microdata Methods and Practice, London.