

Children's Resources in Collective Households: Identification, Estimation and an Application to Child Poverty in Malawi

By GEOFFREY R. DUNBAR, ARTHUR LEWBEL AND KRISHNA PENDAKUR*

The share of household resources devoted to children is hard to identify, because consumption is measured at the household level, and goods can be shared. Using semiparametric restrictions on individual preferences within a collective model, we identify how total household resources are divided up among household members, by observing how each family member's expenditures on a single private good like clothing varies with income and family size. Using data from Malawi we show how resources devoted to wives and children vary by family size and structure, and we find that standard poverty indices understate the incidence of child poverty.

JEL: D13, D11, D12, C31, I32

Keywords: Collective Model, Cost of Children, Bargaining Power, Identification, Sharing rule, Demand Systems, Engel Curves

Most measures of economic well-being rely, to some degree, on individual consumption. Yet the measurement of individual consumption in data is often confounded because consumption is typically measured at the household, not the individual, level. Dating back at least to Becker (1965, 1981), 'collective household' models are those in which the household is characterised as a collection of individuals, each of whom has a well-defined objective function, and who interact to generate household level decisions such as consumption expenditures. Given household-level data, useful measures of individual consumption expenditures are *resource shares*, defined as each member's share of total household consumption. If there is intra-household inequality, these resource shares will be unequal so standard per-capita calculations (assigning equal resource shares to all household members) are invalid measures of individual well-being.

Children differ from other household members in that they do not enter households by choice, they have little ability to leave, and generally bring little income or other resources to the household. Children may therefore be the most vulnerable of household members to intra-household inequality. It is thus imperative to measure children's resource shares in households in order to assess inequality and child poverty. This paper shows identification of children's resource shares in a collective household model, and

* Dunbar: Simon Fraser University, 8888 University Drive, Burnaby, BC, Canada, V5A 1S6, geoffrey_dunbar@sfu.ca. Lewbel: Department of Economics, Boston College, 140 Commonwealth Ave., Chestnut Hill, MA, 02467, USA. lewbel@bc.edu. Pendakur: Simon Fraser University, 8888 University Drive, Burnaby, BC, Canada, V5A 1S6, pendakur@sfu.ca. We would like to thank Martin Browning, Pierre-Andre Chiappori, Lei Han, Federico Perali, Frederic Vermeulen, and some anonymous referees for helpful comments, and the National Statistics Office of Malawi for providing data.

offers simple methods to estimate them.

While many papers exist on identification of collective household models, very few identify the main question we address, which is the share of total household resources devoted to children. See, e.g., Browning (1992) for a survey of the cost of children literature. Most collective household models, if they include children at all, treat children either as household attributes or as consumption goods for parents, rather than as separate economic agents with individual utility functions and associated resource shares. See, e.g., Blundell, Chiappori and Meghir (2005). The implication of these types of models is that children only have utility functions that affect household purchasing decisions once they reach adulthood. It seems more reasonable to assume as we do that households behave as if children have utility functions.

Dauphin et al (2008) and Cherchye, De Rock and Vermeulen (2008) test whether observed household demand functions are consistent with children having separate utility functions, and find evidence in favor of this hypothesis, though they do not estimate children's shares. Cherchye, De Rock and Vermeulen (2010) do consider estimation, but their method generally only yields bounds on resource shares. Our approach offers point identification of children's resource shares.

Based on the collective household model of Chiappori (1988, 1992), a series of papers show identification of *changes* in resource shares as functions of distribution factors, defined as variables which affect bargaining power, but which do not affect preferences over goods or scale economies. See, e.g., Bourguignon and Chiappori (1994), Browning, Bourguignon, Chiappori, and Lechene (1994), and Browning and Chiappori (1998). However, these papers (along with more recent variants such as Vermeulen 2002) do not identify the *level* of resource shares. In contrast, we identify the *levels* of resource shares and do not require distribution factors.

Almost all of these collective household models impose strong restrictions on how goods may be shared among household members. Specifically, they typically assume that all goods are either purely private or purely public within the household. An exception is Browning, Chiappori and Lewbel (2007) (hereafter BCL), who provide a model that nonparametrically identifies the levels of resource shares of all individual household members and which allows for very general forms of sharing of goods (see also Lewbel 2003). We extend BCL, and so allow for these more general types of scale economies.

BCL show identification only when the demand functions of individuals can be separately observed, which is not the case for children since they are always in households along with adults. In practice, BCL observe the demand functions of individuals by observing data from single men and single women living alone, and combine those demand functions with data on the demands of men and women living together as (childless) couples. Accordingly, they assume very limited differences between the utility functions of single and married men and between those of single and married women. Similar restrictions are required by Lise and Seitz (2004).

Lewbel and Pendakur (2008) propose some restrictions on BCL that permit the identification of the levels of adult's resource shares in a model based on comparing the Engel curves (demand equations holding prices constant) of single men and women and men

and women in childless married couples. Bargain and Donni (2009) extend the Lewbel and Pendakur (2008) model to identify the resource shares of children. Both of these papers impose the identification restriction of BCL, namely that single men and single women have the same preferences as their married counterparts.

The contribution of this paper to extend the model of BCL to include children, semi-parametrically identifying the levels of children's resource shares without requiring that single men and women have the same preferences as fathers and mothers.

We identify resource shares using Engel curves. Model and data requirements are also reduced by only needing the demand functions for one type of private good, like clothing. Basing identification and estimation on Engel curves also substantially reduces model complexity, in part since we do not require price data. In our empirical model, we use Engel curves linear in the log of total expenditures, wherein structural parameters relating to resource shares are computed from the slopes of those Engel curves.

Our identification uses private assignable goods. A good is defined to be *private* if it cannot be shared or consumed jointly by more than one person, and is defined to be *assignable* if it is consumed by one individual household member that is known to the researcher. Examples could include toys and diapers which are private goods assignable to children, or alcohol and tobacco which are private goods assignable to adults. In our empirical application we use clothing.

Our identification assumes that resource shares do not vary with total expenditure¹, and assumes one of two semiparametric restrictions on individual preferences (these are similar to, but weaker than, those proposed by Pendakur (1999) and Lewbel and Pendakur (2008)). Given these restrictions, resource shares are identified via comparison of Engel curves. With the first semiparametric restriction, we assume that preferences for a particular good are similar in certain limited ways across people (within household types), and use this similarity to identify resource shares within households with a given number of children. With the second, we assume that a person's preferences for a particular good are similar across household types, and compare the consumption choices of people across households with varying numbers of children.

In comparison with BCL (and with Lewbel and Pendakur 2008 and Bargain and Donni 2009), we do not need to use information on childless households (either couples or singles). In that respect, our identification strategies impose milder conditions on preference stability across household types, since e.g. we would assume that fathers of two children have similar preference to fathers of three children, rather than assume that either are similar to single men. Related identification ideas go back at least to Lazear and Michael

¹Samuelson (1956) shows that resource shares cannot in general be constant for a large class of household social welfare functions. While we assume resource shares independent of household expenditures y , we do not require them to be constant, but instead allow them to vary arbitrarily with prices p (and other household characteristics). For example, a social welfare function that sums the utilities of individuals will satisfy our assumptions (making resource shares independent of y but not of p), if indirect utility functions are linear in $\ln y$ or if they are linear in y^k for any k . This class of indirect utility function (called PIGL and PIGLOG) and its implications for social welfare maximisation is explored by Muelbauer (1974, 1976). In a supplemental online appendix we prove formally that the PIGLOG functional form we use in our empirical application can satisfy this (and other) assumptions we require for identification. Finally, we note that the assumption that resources shares not depend on y (at low levels of y) still permits resource shares to depend on other variables closely related to y , such as household income, wealth, or member's wages. We would like to thank Martin Browning for alerting us to this general issue.

(1988, chapter 4). We also impose milder functional restrictions on demands and preferences than Pendakur and Lewbel (2008) and Bargain and Donni (2009). In particular, we only place restrictions on the demand functions for one set of goods like clothing, instead of imposing restrictions on the demand functions for all goods.

In models where goods are purely private or purely public, what we call private, assignable goods are known as exclusive goods. See, e.g. Bourguignon, Browning, and Chiappori (2009). Chiappori and Ekeland (2008) and Cherchye, De Rock and Vermeulen (2008) among others show how assignable goods can aid in the identification of resource shares. Our strategy follows this line in assuming the presence, and observability, of a small number of private assignable goods, and uses these to identify childrens' resource shares. The end result is that we identify how total household expenditures on *all* goods are divided up among household members, just by observing how family expenditures on each member's single private good (like clothing) vary with total expenditures.

Some previous papers (e.g. Lundberg, Pollak and Wales (1997)) have used private assignable goods to address children's resources without invoking a full structural model of the household like ours. We provide a structural model for calculating the child's economic well-being, defined as the total amount of the household's resources consumed by the child, which is based on budget share equations for private assignable goods like clothing. Our structural model shows that the level of budget shares mixes both a price response, coming in part from the extent to which some goods are consumed jointly, and an income response, coming from the child's share of household expenditure. In contrast to non-structural approaches, our identification of children's resources accounts for these two types of responses.

We present empirical results for children's resource shares in Malawi using data from the Second Integrated Household Survey (IHS2), conducted by the National Statistics Office in conjunction with the International Food Policy Research Institute and the World Bank. We use the Malawi data for two reasons: Malawi is one of the poorest countries in the world, with per-capita (2005 PPP) GDP of US\$773 in 2008, and the IHS2 data are particularly rich in terms of household-level detail, which we exploit in our empirical work. Given the extreme poverty of most Malawian households, one may suspect that children are vulnerable to intra-household inequality.

We find that children command a reasonably large share of resources – roughly 20 percent for the first child – and that this share rises with the number of children – 5-10 percentage points per additional child. Moreover, fathers command a larger share of resources than mothers, and mothers seem to sacrifice more resources than fathers to their children. Indeed, we cannot reject the hypothesis that the father's shares do not respond to the number of children. These patterns are evident even if household size is taken as endogenous and the model is estimated using instrumental variable techniques. These findings are in the spirit of Duflo (2003), who finds evidence that male household heads tend not to allocate additional resources to children while female household heads do.

We find some evidence of gender asymmetry within the household, similar to Rose

(1999).² We find that mothers' resource shares rise, and childrens' resource shares fall, as the proportion of children that are girls rises. Indeed, if all children are girls, then the mother's resource share rises, and the children's share falls, by roughly five percentage points. We also find that higher mother's education is associated with higher resource shares for women and children.

Finally, we use our estimates of resources shares to construct estimates of the poverty incidence of men, women and children in Malawi. Using the World Bank \$2/day per-capita poverty measure, which assumes equal resource shares across people, yields an overall poverty rate of 91%. In contrast, we find that allowing for unequal resource shares across people shows sharp differences in the incidence of poverty. In particular, we find that the incidence of poverty is roughly 60% for men, 85% for women, and over 95% for children.

I. Collective Households and Resource Shares

In the version of the BCL model we consider, each household member is allocated a *resource share*, that is, a share of the total resources (total expenditures) the household has to spend on consumption goods. Within the household, each member faces this total resource income constraint and a vector of Lindahl (1919) type shadow prices for goods. Each household member's resource share may differ from those of other members, but all members face the same shadow price vector. The resource share of a person and shadow price vector of the household together define a shadow budget constraint faced by each individual within the household. Each household member then determines their own demand for each consumption good by maximizing their own utility function.

These shadow prices differ from market prices because of economies of scale to consumption. In particular, shadow prices will be lower than market prices for goods that are shared or consumed jointly by multiple household members. Goods that are not shared (i.e., private goods) will have shadow prices equal to market prices. Each member faces the same shadow prices because the degree to which a good can be shared is an attribute of the good, rather than an attribute of the consumer.

The shadow budget constraint faced by individuals within households can be used to conduct consumer surplus exercises relating to individual well-being. One example of this is the construction of 'indifference scales', a tool BCL develop for comparing the welfare of individuals in a household to that of individuals living alone, analogous to an equivalence scale.

Resource shares for each individual may also be of interest even without knowledge of shadow prices. The resource share times the household expenditure level gives the extent of the individuals' budget constraint for consuming resources within the household, and is therefore an indicator of that individual's material well-being. For example, Lise and

²By asymmetry here we only mean unequal treatment regarding allocation of resources. We do not claim that these allocations are necessarily unfair or imply inequality in welfare. For example, a large fraction of total expenditures in Malawi are devoted to food, so if women and girls are on average smaller and have lower caloric requirements, then they might be equally well off in a welfare sense to men and boys despite having smaller resource shares. We would like to thank Frederic Vermeulen for pointing this out.

Seitz (2004) use estimated resource shares to construct national consumption inequality measures that account for inequality both within and across households.

In addition, because within-household shadow prices are the same for all household members, resource shares describe the relative consumption levels of each member. Consequently, they can be used to evaluate the relative welfare level of each household member, and are sometimes used as measures of the bargaining power of household members. BCL show a one to one relationship between resource shares and collective household model "pareto weights" on individual utility, which are also used as measures of member bargaining power. Since we focus on the estimation of children's resource shares, we do not interpret our results in terms of bargaining power.

A. The Model

We begin by summarizing the BCL model, extended to include children. In general, we use superscripts to index goods and subscripts to index people and households. We consider three types t of individuals: m , f , and c , indicating male adult, female adult, and child. Our results readily extend to more types of individuals, such as younger and older children or boys and girls, but to simplify the presentation consider only households consisting of a mother, a father, and one or more children, so we can index households by the size measure $s = 1, 2, \dots$ where s is the number of children in the family. Also to simplify notation, for now we suppress arguments corresponding to attributes like age, location, etc., that may affect preferences. We also suppress arguments corresponding to distribution factors, that is, variables like relative education levels that may help to determine bargaining power and hence resource shares devoted to each household member. All of our identification results may be conditioned on these types of variables, and when it comes to the empirical section, we will introduce them explicitly.

Households consume K types of goods. Let $p = (p^1, \dots, p^K)'$ be the K -vectors of market prices and $z_s = (z_s^1, \dots, z_s^K)'$ be the K -vectors of quantities of each good k purchased by a household of size s . Let $x_t = (x_t^1, \dots, x_t^K)'$ be the K -vectors of quantities of each good k consumed by an individual of type t .

Let y denote total expenditure, which may be subscripted for households or individuals. Let $U_t(x_t)$ denote an ordinal measure of the utility that an individual of type t would attain if he or she consumed the bundle of goods x_t while living in the household. An individual's total utility may depend on the well being of other household members, on leisure and savings, and on being a member of a household, so $U_t(x_t)$ should be interpreted as a subutility function over goods this period, which may be just one component of member t 's total utility. For children, $U_c(x_c)$ could either represent a child's actual utility function over the bundle of goods x_c that the child consumes, or the utility function that parents believe the child has (or think he or she should have).

For their identification, BCL assume that for a person of type t , $U_t(x_t)$ also equals the utility function over goods of a single person of type t living alone. The Marshallian demand functions of a person t living alone are then obtained by choosing x_t to maximize $U_t(x_t)$ under the linear budget constraint $p'x_t = y$. We do not impose this assumption, so for us $U_t(x_t)$ only describes the preferences over goods of individual t as a member of

a family, which may be completely different from that person's preferences if he or she were living alone. In particular, it would not be sensible to define $U_c(x_c)$ as the utility function of a child living alone.

For simplicity, we assume that each child in a family is assigned the same utility function $U_c(x_c)$. The model will later be extended to include parameters that allow $U_c(x_c)$ to vary by, e.g., the age and sex of the child, but these like other observed household and individual characteristics are omitted for the time being. However, up to the inclusion of such observable characteristics, we assume that the individual household member utility functions $U_f(x_f)$, $U_c(x_c)$, and $U_m(x_m)$ are the same regardless of whether the household has one, two, or three children. So, e.g., in a household with given observed characteristics, mothers have the same preferences over privately consumed consumption goods regardless of how many children are in the household.

In our model and application below we assume each child has the same utility function, but it is straightforward in theory to extend the model to allow each child to have a different utility function. In this case, using arguments analogous to those presented below, a separate private assignable good for each child is needed to achieve identification of each child's resource share. The choice of whether to allow utility to vary across children is data-driven, specifically, it is possible to do so if one has data on private assignable goods for each child (rather than for all the children together).

We assume that the total utility of person t is weakly separable over the subutility functions for goods. So, e.g., a mother who gets utility from her husband's and child's well-being as well as her own would have a utility function of the separable form $U_f^*[U_f(x_f), U_c(x_c), U_m(x_m)]$ rather than being some more general function of x_f, x_m , and x_c .

Following BCL, assume that the household has economies of scale to consumption (that is, sharing and jointness or consumption) of a Gorman (1976) linear technologies type. The idea is that a bundle of purchased goods given by the K vector of purchased quantities z_s is converted by a matrix A_s into a weakly larger (in magnitude of each element) bundle of 'private good equivalents' x , which is then divided among the household members, so $x = x_f + x_m + x_c$. Specifically, there is assumed to exist a K by K matrix A_s such that $x_f + x_m + x_c = x = A_s^{-1}z_s$. This "consumption technology" allows for much more general models of sharing and jointness of consumption than the usual collective model that categorizes goods only as purely private or purely public.

For example, suppose that a married couple without children ride together in a car (sharing the consumption of gasoline) half the time the car is in use. Then the total consumption of gasoline (as measured by summing the private equivalent consumption of each household member) is $3/2$ times the purchased quantity of gasoline. Equivalently, if there had been no sharing of auto usage, so every member always drove alone, then the couple would have had to purchase 50% more gasoline to have each member travel the same distance as before. In this example, we would have $x^k = (3/2)z^k$ for k being gasoline, so the k 'th row of A would consist of $2/3$ in the k 'th column and zeros elsewhere. This $2/3$ can be interpreted as the degree of "publicness" of good k within the household. A purely private good k would have $x^k = 1$. Nonzero off diagonal elements

of A_s may arise when the extent to which one good is shared depends upon other goods, e.g., if leisure time is a consumption good, then the degree to which auto use is shared may depend on the time involved, and vice versa.

BCL assume the household is Pareto efficient in its allocation of goods, and does not suffer from money illusion. This implies the existence of a monotonically increasing function \tilde{U}_s such that a household of type s buys the bundle of goods z_s given by

$$(1) \quad \max_{x_f, x_m, x_c, z_s} \tilde{U}_s [U_f(x_f), U_m(x_m), U_c(x_c), p/y]$$

such that $z_s = A_s [x_f + x_m + x_c]$ and $y = z'_s p$

Solving the household's maximization problem, equation (??) yields the bundles x_t of "private good equivalents" that each household member of type t consumes within the households. Pricing these vectors at within household shadow prices $A'_s p$ (which differ from market prices because of the joint consumption of goods within the household) yields the fraction of the household's total resources that are devoted to each household member.

Let η_{ts} denote the *resource share*, defined as fraction of the household's total expenditure consumed by a person of type t in a household with s children. This resource share has a one-to-one correspondence with the "pareto-weight", defined as the marginal response of \tilde{U}_s to U_t .

In this paper, we lean heavily on existence of private assignable goods for identification of resource shares. A *private good* for our purposes is defined as a good with its corresponding diagonal element of A equal to 1 and all off-diagonal elements in that row or column are equal to 0. This means that private goods are goods that do not have any economies of scale in consumption. For example, food is private to the extent that any unit consumed by one person cannot also be eaten by another.³ A private good is *assignable* if it is consumed exclusively by one known household member. So, e.g., a sandwich would be assignable if we could observe who ate it. Note that if a good is private, assignability has no further consequence for preferences. For example, preferences (and resource shares) determine who in the household eats a sandwich, but given that the sandwich is privately consumed, it is assignable if and only if the data on who ate it is collected and provided for analysis. In our application we observe separate expenditures on men's, women's, and children's clothing, which we take to be private and assignable.

Our definition of a private assignable good is quite strict, but we do not need to rule out all externalities. In particular, we can allow for externalities of private assignable goods onto the *utilities* of other household members, but we cannot allow for any externalities that affect household resource allocations or the expenditure patterns of other household members. So, e.g., smoking could be used as an identifying private assignable good even if the smoke made other household members unhappy, but not if the smoke made other household members spend more than otherwise on household cleaning products.

³This ignores possible economies of scale in food from reduced waste associated with preparation of larger quantities.

Suppose there exists a private assignable good for a person of type t . This good is not jointly consumed, and so appears only in the utility function U_t , not in the utility functions of any other type of household member. Let $W_{ts}(y, p)$ be the share of total expenditures y that is spent by a household with s children on the type t private good. For example $W_{cs}(y, p)$ could be the fraction of y that a household with s children spends on toys or children's clothes. Also let $w_t(y, p)$ be the share of y that would be spent buying the type t private good by a (hypothetical) individual that maximized $U_t(x_t)$ subject to the budget constraint $p'x_t = y$. Unlike in BCL, these individual demand functions need not be observable.

While the demand functions for goods that are not private are more complicated (see the online supplemental appendix for derivations and details, especially equation (2) in Appendix A.1), the household demand functions for private assignable goods, derived from equation (??), have the simple forms

$$(2) \quad \begin{aligned} W_{cs}(y, p) &= s \eta_{cs}(y, p) w_c(\eta_{cs}(y, p) y, A'_s p) \\ W_{ms}(y, p) &= \eta_{ms}(y, p) w_m(\eta_{ms}(y, p) y, A'_s p) \\ W_{fs}(y, p) &= \eta_{fs}(y, p) w_f(\eta_{fs}(y, p) y, A'_s p) \end{aligned}$$

This solution to BCL for the case of private assignables states that the household's budget share for a person's private assignable good is equal to her resource share multiplied by the budget share she would choose herself if facing her personal shadow budget constraint. Household demand functions W_{ts} , the left side of equation (2), are in principle observable by measuring the consumption patterns of households with various y facing various p regimes. Our goal is identification of features of the right side of equation (2), in particular η_{cs} , and moreover we wish to obtain identification using only data from a single price regime.

Two problems prevent us from using the BCL identification strategy in our setting with children. First, unlike adults, we cannot observe the demand functions for children living alone. BCL exploited data on adults living alone by assuming that single and married individuals have the same underlying utility functions. We replace this questionable assumption with the milder assumption that parents (and individual children) have utility functions over goods that do not depend on whether the number of children in the household is one, two, or three. (Our formal assumptions are even weaker, as described below, and in an online supplemental appendix.)

A second problem with BCL is that identification of the household consumption technology A_s requires observable price variation and the measurement of price responses in household demand functions. The measurement of price responses in demand is typically difficult for at least two reasons: first, the rationality restrictions of Slutsky symmetry and homogeneity typically require that price effects enter demand functions in complicated nonlinear ways; and second, there is often not much observed relative price variation in real data, so estimated price responses can be very imprecise. Indeed, many data sources on household consumption of commodities have no information at all on the prices of those commodities.

We get around these two problems in two steps. First, we restrict the resource share functions η_{fs} to be independent of household expenditures y , at least at low expenditure levels (though they may depend arbitrarily on prices p). This restriction has real bite, but one can at least write down sensible parametric household objective functions over reasonable parametric utility functions whose resulting resource shares satisfy this restriction (see footnote 1; in addition, we present a class of such models in an online supplemental appendix). Moreover, while resource shares cannot depend on total expenditures y , they can depend on closely related variables such as income, wages, or wealth. Similar to Lewbel and Pendakur (2008) and Bargain and Donni (2009), this restriction allows us to recast the BCL model into an Engel-curve framework where price variation is not exploited for identification.

Second, we invoke some semiparametric restrictions on the shapes of individual Engel curves. These restrictions allow us to identify individual resource shares by comparing household demands for private assignables across people within households, or by comparing these demands across households for a given type of person. Unlike Bargain and Donni (2009), who also identify children's resource shares from Engel curves, we only place restrictions on the shapes of Engel curves for the assignable goods rather than on all goods, and we only need to assume similarity of preferences of individuals in households with varying numbers of children, rather than equality of preference of all adults regardless of whether they are single, couples without children, or couples with children.

II. Identification of Children's Resource Shares Using Engel Curves

In this section, we offer a brief nontechnical description of how we achieve identification of each person's resource share in the collective household, using only data on Engel curves for private assignable goods in households with children. Technical discussion and formal identification proofs are in an online supplemental appendix.

An Engel curve is defined as the functional relationship between a budget share and total expenditure, holding prices constant. In a slight abuse of notation, we may write the BCL solutions for private assignables given by equation (2) in Engel curve form as

$$(3) \quad \begin{aligned} W_{cs}(y) &= s\eta_{cs}w_{cs}(\eta_{cs}y) \\ W_{ms}(y) &= \eta_{ms}w_{ms}(\eta_{ms}y) \\ W_{fs}(y) &= \eta_{fs}w_{fs}(\eta_{fs}y). \end{aligned}$$

Here, the Engel curve function w_{ts} gives the demand function for person t when facing the price vector $A'_s p$ for one particular value of p . For example, the children's Engel curve function, $w_{cs}(\eta_{cs}y) = w_c(\eta_{cs}(p)y, A'_s p)$, gives their (latent) clothing demand for that one value of p . The resource share η_{ts} does not depend on y by assumption, and its dependence on p is suppressed in the Engel curve $w_{cs}(\eta_{cs}y)$ because prices are held constant.

The main difficulty for identification is that for every observable budget share function subscripted by ts on the left side of (3), there are two unobservable functions subscripted by ts on the right side. BCL achieve identification by assuming that w_{ts} on the right-

hand side is observable via the behaviour of single people, leaving just one subscribed unobserved function to worry about: the resource shares η_{ts} . There are no single children, so we cannot use this method.

One extreme alternative would be to assume that people have identical preferences so that w_{ts} does not vary across t . In this case, for any household size s , we would use the 3 observable functions W_{ts} (for $t = m, f, c$) to identify 2 resource shares η_{ts} (the third may be computed because they add up to 1) and 1 budget share function w_s . A different extreme alternative would be to assume that people have preferences which do not vary across household type, so that w_{ts} does not vary across s . In this case, if we had enough household sizes, we would similarly have enough observable household budget share functions W_{ts} to identify the unobserved resource shares η_{ts} and unobserved individual budget share functions w_t . Unfortunately, both of these extreme restrictions are unreasonable. The first assumes that preferences are completely identical across people. The second is roughly equivalent to forcing w_t to be unresponsive to prices.

Our identification is based on the insight that one does not need the entire function w_{ts} to be independent of t or of s . It is enough for a separable part of w_{ts} to be independent of t or of s . Consider budget share functions w_t that are linear in functions of expenditure:

$$w_t(y, p) = h_{t0}(p) + h_{t1}(p)g_1(y) + h_{t2}(p)g_2(y) + \dots + h_{tL}(p)g_L(y),$$

where $h_{tl}(p)$ are price-varying functions which multiply the functions of expenditure $g_l(y)$. Then, observed private assignable budget share equations would be given by

$$W_{ts}(y) = \eta_{ts}h_{ts0} + \eta_{ts}h_{ts1}g_1(\eta_{ts}y) + \eta_{ts}h_{ts2}g_2(\eta_{ts}y) + \dots + \eta_{ts}h_{tsL}g_L(\eta_{ts}y),$$

where $h_{tsl} = h_{tl}(A_s p)$ for m, f and with $W_{cs}(y)$ defined analogously. We could achieve identification if any $h_{tl}(p)$ was independent of t so the coefficient h_{tsl} would drop its dependence on t . In this case, preferences would not be identical across people (indexed by t), but would be *similar* across people, due to the fact that one separable part of the budget share function is the same for all people. Identification would be analogous to the case where people had completely identical budget share functions.

Alternatively, we would achieve identification if any $h_{tl}(p)$ was independent of p so that the corresponding coefficient h_{tsl} would drop its dependence on s . In this case, preferences would not be identical across household types, but for any given person they would be *similar* across household types. Identification would be analogous to the case where preferences don't vary across household types.

Although the formulation above is useful for seeing how identification works, it is well-known that not all such formulations can be rationalised with a utility function (that is, not all are integrable). In the next sections, we describe restrictions which give individual budget share functions that can be rationalised with individual utility functions, and which permit identification of individual resource shares.

A. Identification if Preferences are Similar Across People

Here, we consider identification when people have similar preferences. We restrict how preferences for the private assignable goods vary across people, so we consider the same good for all people. For example, the private assignable good could be clothing, so that the demand function $w_t(y, p)$ gives person t 's (unobserved) budget-share function for clothing when facing the constraint defined by y, p . In particular, we impose the restriction that Engel curves for the private assignable good have the same shape across people, at least at low expenditure levels⁴:

$$(4) \quad w_t(y, p) = d_t(p) + g\left(\frac{y}{G_t(p)}, p\right) \text{ for } y \leq y^*(p),$$

where $y^*(p)$ is a real expenditure threshold. The budget share functions for all people have the same shape, given by the function g , and differ only by the person-specific additive term $d_t(p)$ and the person-specific expenditure deflator $G_t(p)$. If $d_t(p)$ and the person-specific expenditure deflator $G_t(p)$ were the same for all people t , then preferences would be identical across people. These functions may differ across people, so we say that preferences are *similar across people* (SAP) if equation (4) holds.

SAP is similar to the shape-invariance restriction of Pendakur (1999) and Lewbel (2010), except that we apply it only to the Engel curves for the private assignable goods and we apply it only at low expenditure levels. Pendakur (1999) shows that if people have costs that differ only by (price-dependent) multiplicative equivalence scales, then budget share functions must satisfy a condition like SAP for all goods and at all expenditure levels. When SAP is applied to all goods and at all expenditure levels, the result is a much stronger condition, known in the consumer demand literature as "shape-invariance". Many empirical consumer demand analyses impose this shape-invariance restriction on budget share functions. See, e.g., Blundell, Duncan, and Pendakur (1998), Blundell, Chen, and Kristensen (2007), and Lewbel (2010). Some have tested the restriction of shape-invariance, and found that it does not do great violence to the data (see, e.g., Pendakur 1999 and Blundell, Chen, Kristensen 2007). In our model, we only assume SAP for a single good and only at real expenditure levels below a threshold $y^*(p)$.

⁴Our assumptions do not rule out applying these conditions (and the corresponding condition for SAT) at all expenditure levels. This corresponds to an infinite threshold $y^*(p)$. One could also specify and estimate a model that relaxes these conditions above the threshold, and then estimate the cutoff threshold along with the other parameters of the model. The cutoff would generally be identified assuming that the model was correctly specified and included other parameters that are nonzero at expenditure levels where the conditions do not hold.

Substituting the SAP restriction (4) into (3) we get, for $y \leq y^*$,

$$\begin{aligned} W_{cs}(y) &= s\eta_{cs}\delta_{cs} + s\eta_{cs}\gamma_s \left(\frac{\eta_{cs}y}{\Gamma_{cs}} \right), \\ W_{ms}(y) &= \eta_{ms}\delta_{ms} + \eta_{ms}\gamma_s \left(\frac{\eta_{ms}y}{\Gamma_{ms}} \right), \\ W_{fs}(y) &= \eta_{fs}\delta_{fs} + \eta_{fs}\gamma_s \left(\frac{\eta_{fs}y}{\Gamma_{fs}} \right), \end{aligned}$$

where $\delta_{ts} = d_t(A'_s p)$, $\gamma_s(y) = g(y, A'_s p)$ and $\Gamma_{ts} = G_t(A'_s p)$. The key here is that g does not vary across people. All these functions are evaluated at the same shadow price vector $A'_s p$, and as a result the function γ_s does not vary across people either (it does not have a t subscript). Theorem 1 in the online supplemental appendix shows the class of individual utility functions that satisfy SAP, and shows that if the function g has sufficient nonlinearity, then the resource shares η_{ts} are identified from the Engel curve functions $W_{ts}(y)$ for any household size s .

A simple example (which we will use in our empirical work) shows how this identification works. Suppose that each person has preferences over goods given by a PIGLOG (see the online supplemental appendix and Muellbauer 1979) indirect utility function, which has the form $V_t(p, y) = b_t(p) [\ln y - \ln a_t(p)]$. An example is the popular Almost Ideal demand system (Deaton and Muelbauer 1980). With PIGLOG preferences, a sufficient restriction for SAP is $b_t(p) = b(p)$.

By Roy's identity, corresponding budget share functions for each person's private assignable are then given by

$$w_t(y, p) = d_t(p) + \beta(p) \ln y,$$

where d_t is a function of $a_t(p)$ and $b(p)$, and $\beta(p)$ is minus the price elasticity of $b(p)$ with respect to the price of the private assignable good.

Plugging these budget share functions into (3) yields

$$(5) \quad \begin{aligned} W_{cs}(y) &= s\eta_{cs}(\delta_{cs} + \beta_s \ln \eta_{cs}) + s\eta_{cs}\beta_s \ln y, \\ W_{ms}(y) &= \eta_{ms}(\delta_{ms} + \beta_s \ln \eta_{ms}) + \eta_{ms}\beta_s \ln y, \\ W_{fs}(y) &= \eta_{fs}(\delta_{fs} + \beta_s \ln \eta_{fs}) + \eta_{fs}\beta_s \ln y, \end{aligned}$$

for any household size s , where $\delta_{ts} = d_t(A'_s p)$ and $\beta_s = \beta(A'_s p)$. These three household Engel curves are linear in $\ln y$, with slopes that can be identified by linear regressions of the household budget shares W_{ts} on a constant and on $\ln y$. The slopes of these three Engel curves are proportional to the unknown resource shares η_{ts} , and the constant of proportionality is identified by the fact that resource shares must sum to one. Equivalently, we have four equations (three Engel curves and resource shares summing to one) in four unknowns (three resource shares and the preference parameter β_s). Consequently, resource shares are exactly identified from a single household's Engel curves for the

private assignable good for each of its three members.

With more complex Engel curves for private assignable goods, identification is achieved by taking higher-order derivatives of the household Engel curves with respect to y or $\ln y$, but the spirit of the identification is the same. By assuming that individuals have budget share functions for their private goods that have the same shape across people for a given price vector, we are able to compare the shape of household Engel curves across people when they face the common within-household shadow price vector. Formal identification theorems are provided in an online supplemental appendix.⁵

B. Identification if Preferences are Similar Across Types

Our second, alternative shape restriction for identifying resource shares invokes comparability across household types (or, equivalently, across shadow-price vectors) for a given person, rather than across people for a given household type. In particular, here we assume that cross-price effects load onto an expenditure deflator for the shadow-price vectors associated with households with one, two, or three children.

Let $p = [p_m, p_f, p_c, \bar{p}, \tilde{p}]$ where \bar{p} is the subvector of p corresponding to purely private goods other than the assigned private goods, and \tilde{p} is the subvector of p corresponding to all the other goods. Note that \bar{p} includes goods like food that are private but may not be assignable. Let L be the total number of private goods. The matrix A_s is block-diagonal, with an upper left block \bar{A}_s equal to the identity matrix and a lower-right block \tilde{A}_s which is unspecified. For the private goods, the corresponding elements of $A_s p$ are p_m, p_f, p_c and \bar{p} , since by definition the shadow prices of private goods equal their market prices. The shadow price of non-private goods is $\tilde{A}_s \tilde{p}$. Thus, for private goods, the difference in a person's budget shares across household sizes is driven by two factors: changes in their resource share, and their cross-price demand responses.

Now we invoke the restriction that preferences are "similar across types" (SAT) as follows:

$$(6) \quad w_t(y, p) = g_t \left(\frac{y}{G_t(\tilde{p})}, p_t, \bar{p} \right) \text{ for } y \leq y^*(p).$$

Again, $y^*(p)$ is a real expenditure threshold, so the restriction is applied only at low expenditure levels. Here, the scale-economies associated with non-private goods load onto the person-specific expenditure deflator $G_t(\tilde{p})$. If $G_t(\tilde{p}) = 1$, then preferences would be identical across household types. But, we allow preferences to vary through the expenditure deflator $G_t(\tilde{p})$, so we say that preferences are only similar across types.

If SAT were applied to all price effects, rather than just the cross-price effects of non-private goods, so that $w_t(y, p) = g_t \left(\frac{y}{G_t(p)} \right)$, and if it were applied to all goods at all expenditure levels, then preferences would be homothetic, which is clearly undesirable.

⁵The online supplemental appendix also provides more details regarding the construction of PIGLOG preference models and household models that are consistent with all of our assumptions, including, e.g., that resource shares be independent of y .

Here, we apply it only to the cross-price effects of non-private goods on the private assignable good, and we apply it only at low expenditure levels.

Lewbel and Pendakur (2008) apply a restriction like SAT to all price effects for all goods at all expenditure levels. They avoid the implication of homotheticity by requiring that the restriction hold for just one set of price changes rather than for all possible price vectors. In contrast, we assume that the restriction holds for all price changes, but only for the Engel curves of the assignable good.

Substituting the SAT restriction (6) into (3), we get

$$(7) \quad \begin{aligned} W_{cs}(y) &= s\eta_{cs}\gamma_c \left(\frac{s\eta_{cs}y}{\Gamma_{cs}} \right) \\ W_{ms}(y) &= \eta_{ms}\gamma_m \left(\frac{\eta_{ms}y}{\Gamma_{ms}} \right), \\ W_{fs}(y) &= \eta_{fs}\gamma_f \left(\frac{\eta_{fs}y}{\Gamma_{fs}} \right), \end{aligned}$$

where $\gamma_t(y) = g_t(y, p_t, \bar{p})$ and $\Gamma_{ts} = G_t(\tilde{A}'_s \tilde{p})$. The key here is that the functions g_t , and therefore $\gamma_t(y)$, do not depend on household size s . We show in Theorem 2 in the supplemental appendix that if private assignable good budget shares don't asymptote to zero when expenditures get too low (that is, if $\lim_{u \rightarrow 0} \gamma_t(u) \neq 0$) and there is sufficient variation in resource shares across individuals and household sizes, then the resource shares η_{ts} are identified from the Engel curve functions $W_{fs}(y)$ for any three household sizes.

To illustrate, suppose again that each person has PIGLOG preferences over goods, so the indirect utility is given by $V_t(p, y) = b_t(p) [\ln y - \ln a_t(p)]$. This utility function satisfies SAT if $b_t(p) = \bar{b}_t(\bar{p}/p_t)$ and $a_t(p) = \bar{a}_t(\tilde{p})$, so \bar{b}_t is some function of private good prices and \bar{a}_t is some function of the prices of other goods.⁶ By Roy's identity, the corresponding budget share functions for each person's private assignable good are given by

$$w_t(y, p) = d_t(p) + \beta_t(\bar{p}/p_t) \ln y,$$

where $d_t(p)$ is a function of $\bar{a}_t(\tilde{p})$ and $\bar{b}_t(\bar{p}/p_t)$, and $\beta_t(\bar{p}/p_t)$ is minus the own-price elasticity of $\bar{b}_t(\bar{p}/p_t)$. Plugging these budget share functions into (3) yields

$$(8) \quad \begin{aligned} W_{cs}(y) &= s\eta_{cs} (\delta_{cs} + \beta_c \ln \eta_{cs}) + s\eta_{cs}\beta_c \ln y, \\ W_{ms}(y) &= \eta_{ms} (\delta_{ms} + \beta_m \ln \eta_{ms}) + \eta_{ms}\beta_m \ln y, \\ W_{fs}(y) &= \eta_{fs} (\delta_{fs} + \beta_f \ln \eta_{fs}) + \eta_{fs}\beta_f \ln y, \end{aligned}$$

⁶Assumption B3 of Theorem 2 in the supplemental online appendix provides a general class of utility functions that yield equation (6). For PIGLOG preferences, Assumption B3 holds if $b_t(p) = \bar{b}_t(\bar{p}/p_t)$ and $a_t(p) = \bar{a}_t(\tilde{p})$. However, Assumption B3 is sufficient but not necessary for equation (6), and in the case of PIGLOG, this equation will hold under the weaker restriction that $b_t(p) = \bar{b}_t(\bar{p}/p_t)\bar{b}_t(\tilde{p})$ and $a_t(p)$ is unrestricted, so the only required restriction for PIGLOG is that $b_t(p)$ be multiplicatively separable into a function of private goods $\bar{b}_t(\bar{p}/p_t)$ and a function of public goods $\bar{b}_t(\tilde{p})$. Either way, the Engel curve system to be estimated takes the form (8).

where $\delta_{ts} = d_t(A'_s p)$ and $\beta_t = \beta(\bar{p}/p_t)$. These Engel curves are linear in $\ln y$, with slopes that vary across household size s for any person t . The coefficient of $\ln y$ for person t in a household with s children (which can be identified by linearly regressing W_{ts} on a constant and on $\ln y$) is $\eta_{ts}\beta_t$. The ratio of $\ln y$ coefficients for person t 's assignable good in two different households equals the ratio of that person's resource shares in the two households. Given three household sizes we have a total of twelve equations (three Engel curves for each of three households, plus three sets of resource shares summing to one) in twelve unknowns (three sets of three resource shares, plus three β_t parameters), so the order condition for identification is satisfied. The corresponding rank condition for identification is provided in an online supplemental appendix. A nice feature of the SAT restriction is that with more than 3 household sizes, the model is overidentified. Thus, the information from additional household sizes can be used to test the model, or to improve the precision of the estimates.

One drawback of using the SAT restriction is that the identification hinges on the summation restriction on the resource shares, and hence may not be very strong in practice. To see this, observe that SAT with PIGLOG preferences identify resource shares by having derivatives of observable budget shares that satisfy

$$\begin{aligned}\partial W_{cs}(y) / \partial \ln y &= s\eta_{cs}\beta_c \\ \partial W_{ms}(y) / \partial \ln y &= \eta_{ms}\beta_m \\ \partial W_{fs}(y) / \partial \ln y &= \eta_{fs}\beta_f.\end{aligned}$$

for multiple values of s . Since the β_t coefficients are also unknown, the only thing that identifies the levels of η_{ts} from the observed budget share functions is the restriction that the resource shares η_{ts} sum to 1. If we instead had the restriction that the product of η_{ts} was 1, then identification would fail, because then we could for example replace each η_{ts} and β_t with $\tilde{\eta}_{ts} = \eta_{ts}\lambda_t$ and $\tilde{\beta}_t = \beta_t/\lambda_t$ for any positive constants λ_t such that $\lambda_m\lambda_f\lambda_c = 1$, without changing any of the observed budget share derivatives. Thus SAT identification is as fragile as the difference between a restriction on the sum versus a restriction on the logged sum. This suggests that although identification is possible given the SAT restriction alone, it may take a lot of data to get precise estimates just from SAT.

The point of this example is that the model provides the restriction that shares sum to one, and if the model had instead provided the restriction that shares multiply to one (or equivalently, that the sum of logged shares were zero) then identification based on SAT would fail. We are not claiming that shares multiplying to one are likely or unlikely, we are only pointing out that SAT identification is as fragile as the difference between a restriction on the sum versus a restriction on the logged sum.

C. Combining restrictions

Our two restrictions, (4) and (6), can be used separately for identification, or combined to strengthen the identification. Either restriction is partly testable (with price variation) because one can test whether or not household demands fit into the structures given by equation (4) or equation (6). Semiparametric testing may follow the lead of Pen-

dakur (1999) or Blundell, Chen and Kristensen (2007). In this paper, we briefly explore parametric testing via overidentification with more than one private assignable good per person and overidentification from having more than three household sizes.⁷

With PIGLOG preferences, SAP holds if $V_i(p, y) = b(p) [\ln y - \ln a_i(p)]$ and SAT holds if $V_i(p, y) = \bar{b}_i(\bar{p}/p_i) \tilde{b}_i(\tilde{p}) [\ln y - \ln a_i(p)]$, so the combination of both holds if $V_i(p, y) = \bar{b}(\bar{p}/p_i) \tilde{b}(\tilde{p}) [\ln y - \ln a_i(p)]$ for some functions \bar{b} and \tilde{b} and if the private assignable goods all have the same price, so $p_c = p_f = p_m$. Equal prices would hold if each member is buying the same type of private assignable good, like similar clothing. By Roys identity, corresponding budget share functions for each person's private assignable will then be given by

$$w_t(y, p) = d_t(p) + \beta \ln y,$$

for some functions $d_t(p)$, and household demands for the private assignables are then

$$(9) \quad \begin{aligned} W_{cs}(y) &= s\eta_{cs}(\delta_{cs} + \beta \ln \eta_{cs}) + s\eta_{cs}\beta \ln y, \\ W_{ms}(y) &= \eta_{ms}(\delta_{ms} + \beta \ln \eta_{ms}) + \eta_{ms}\beta \ln y, \\ W_{fs}(y) &= \eta_{fs}(\delta_{fs} + \beta \ln \eta_{fs}) + \eta_{fs}\beta \ln y \end{aligned}$$

for all household sizes s and for all persons c, m, f . Essentially, here we take the household demands (5), which may have different slopes for each household size, and impose the SAT restriction that the shapes are the same across different household sizes.

It is important to stress that by invoking either or both of our identifying restrictions, we identify the levels of the resource shares themselves, not just how they vary with distribution factors, and we identify children's resource shares, not just those of adults. These features are not provided in the existing literature on resource share/pareto-weight identification (as discussed in the introduction). Both are crucially important for our policy analysis, which is to measure the relative welfare of children in households of varying composition.

Another feature of our identification results is that the associated estimators can be easy to implement. We do not require any data on prices, we do not require a breakdown of household total expenditures into many different goods (only some private, assignable goods are needed), and we do not require a division of household characteristics into 'distribution factors' versus preference shifters. When using the PIGLOG specification for individual utility functions (which includes the Almost Ideal model as a special case), the equations to be estimated are linear in the variables. With identification using SAP, the reduced form parameters may be obtained via OLS estimation of these equations for any particular household size, with the structural parameters being given by nonlinear functions of the reduced form parameters. With identification using SAT, this "OLS identification" holds for estimation with any 3 household types, and with more than three

⁷ Either restriction is compatible with large classes of indirect utility functions as described in the supplemental online appendix, though obviously the intersection of these restrictions is smaller. Using both restrictions together should provide more efficient estimates, assuming both restrictions hold.

types the model is still linear, but there are nonlinear restrictions on the parameters that, for efficiency, should be imposed upon estimation. Consequently, estimation is far less onerous, both computationally and in terms of data requirements, than other empirical collective household models such as BCL, and is more in the spirit of the econometric shortcuts offered by Lewbel and Pendakur (2008).

III. Engel Curve Estimation

A. Malawian Expenditure Data

We use Malawian household expenditure and demographic data. Malawi is one of the poorest countries on earth, with an average per capita income level of less than one US dollar per day. It is a former British protectorate in southern Africa which achieved independence in 1964. The population of Malawi is roughly 16 million as of 2009 with a population density of approximately 120 persons per sq. km. It is one of the most densely populated countries in Africa. Half of Malawians live in the Southern region, 40% in the Central region and 10% in the Northern region, with more than 90% of the population living in rural areas. The economy of Malawi is largely based on agriculture and fishing with its chief exports being tobacco and sugarcane. Its recent political history has been remarkable for the absence of military coups and occasional multi-party elections, most recently in 2009. Despite its relative political stability, Malawi has numerous socio-economic tensions including extreme poverty (over 90% living under two US dollars per person per day), a high incidence of HIV/AIDs, high infant mortality and one of the lowest life expectancies in the world (51 years). In 2005, Malawi received almost \$600 million in foreign aid, equivalent to roughly 50 percent of government spending. Malawi is a good case study for our empirical exercise of measuring intra-household inequality, because with so much of the population having low household expenditure, inequality within households could substantially change the assessment of individual level poverty.

The data come from the second Malawi Integrated Household Survey (IHS2), conducted in 2004-2005, made available to us by the National Statistics Office of Malawi (NSOW 2011). The Survey was designed by the National Statistics Office of the Government of Malawi with assistance from the International Food Policy Research Institute and the World Bank in order to better understand poverty at the household level in Malawi. The survey includes roughly 11,000 households, drawn randomly from a stratified sample of roughly 500 strata.⁸

Enumerators were sent to individual households to collect the data. Enumerators were monitored by Field Supervisors in order to ensure that the random samples were followed and also to ensure data quality. Cash bonuses, equivalent to roughly 30 per cent of average household income in Malawi, were used as an incentive system in the IHS2 for all levels of workers. Roughly 5 per cent of the original random sample was resampled because dwellings were unoccupied. Only 0.4 per cent of initial respondents refused to

⁸For computational reasons, we do not use the complex sampling information associated with stratification in our estimation. This means that our estimates are unbiased and consistent, but not efficient. However, the robust nonlinear SUR and GMM estimated standard errors that we report remain consistent.

answer the survey in the IHS2, so endogenous selection of reporters is not likely to be a problem in these data.

In the Survey, households are asked questions from a number of modules relating to health, education, employment, fertility and consumption. Households are asked to recall their food consumption (one week recall) and their non-food expenditure broken into four recall categories (one week, one month, three months and one year). Consumption amounts also include the value of home produced goods and services imputed at the value of those services consumed in the market.

The consumption data include (in the three month recall questionnaire) household expenditures on clothing and shoes for the household head, spouse(s), boys and girls. These are our assignable goods which we construct for each household from the detailed module data. For almost all the empirical work, we use a single private assignable good for each person equal to the sum of clothing and footwear expenditures for that person. As distribution and demographic factors, we use information from the remaining modules to construct measures of education, age, marital status, etc.

B. Estimation

In this section, we estimate Engel curve systems in an environment without price variation using the identification results provided in Theorems 1 and 2. Our sample consists of 2794 households comprised of married couples with one to four children all under 15 years of age. These households (drawn from the database of approximately 11,000 households) satisfy the following additional sample restrictions: (1) polygamous marriages are excluded; (2) observations with any missing data on the age or education of members are excluded; (3) households with children aged 15 or over are excluded; (4) households with any member over 65 are excluded; and (5) urban households are excluded. Our private assignable good is the sum of clothing and footwear expenditures. Table 1 gives summary statistics of our assignable goods and some general expenditure and wealth data for our sample of nonurban families with 2 parents and 1-4 children.

| | couples with | | | | all |
|----------------------------------|--------------|------------|------------|------------|------|
| | 1 child | 2 children | 3 children | 4 children | |
| Number of Observations | 845 | 825 | 667 | 457 | 2794 |
| clothing plus | | | | | |
| men | 1.46 | 1.34 | 1.21 | 1.00 | 1.29 |
| footwear | | | | | |
| women | 2.10 | 1.92 | 1.61 | 1.52 | 1.84 |
| (in per cent) | | | | | |
| children | 1.06 | 1.50 | 1.69 | 1.89 | 1.48 |
| food (in per cent) | 58.8 | 58.8 | 59.3 | 59.2 | 59.2 |
| transportation (in per cent) | 2.3 | 2.3 | 2.2 | 2.3 | 2.3 |
| housing (in per cent) | 21.7 | 21.7 | 21.8 | 21.3 | 21.9 |
| log-total-expenditure (demeaned) | -0.09 | -0.03 | 0.05 | 0.15 | 0 |
| log-livestock-value (demeaned) | -0.39 | -0.16 | 0.24 | 0.67 | 0 |
| log-durables-value (demeaned) | -0.23 | 0.04 | -0.04 | 0.41 | 0 |

Because the Malawian data are very rich, we also include some demographic variables, which may affect preferences and/or resource shares. If they were to affect resource shares and not preferences, they would be called "distribution factors" in much of the collective household model literature. Our theorems show identification for models without these variables, so one can apply the theorems conditioning on each value these additional variables can take on, and thereby prove identification when these variables are included. As in Browning and Chiappori (1998) the presence of distribution factors may help identification of resource shares, but, unlike Browning and Chiappori (1998) (and most other empirical collective household models), we do not require distribution factors for identification. This also means that we do not have to take a stand on whether any particular demographic variable affects only resource shares and hence is a distribution factor, versus affecting either resource shares, preferences or both. All demographic variables can be allowed to affect both the allocation of resources across individuals, and the preferences of all individuals in the households.

We include 14 demographic variables in our models: region of residence (non-urban North and non-urban Central with non-urban South as the left-out category); the average age of children less 5; the minimum age of children less 5; the proportion of children who are girls; the age of the man less 28 and the age of the woman less 22 (the average ages of men and women in the sample); the education levels of the household head and spouse (ranging from -2 to 4 , where 0 is the modal education level); the log of the distance of the village to a road and to a daily market; a dummy indicating that the 3 month recall period for consumption occurred over the dry season; and dummy variables indicating that the household is christian or muslim (with animist/other as the left-out category). We allow all demographic factors to affect both the preferences and the resource shares of every household member.

We estimate models corresponding to individuals with PIGLOG indirect utility functions and their resulting log-linear Engel curves. Household budget share equations are given by

$$(10) \quad \begin{aligned} W_{cs}(y) &= s\eta_{cs}(\delta_{cs} + \beta_{cs} \ln \eta_{cs}) + s\eta_{cs}\beta_{cs} \ln y, \\ W_{ms}(y) &= \eta_{ms}(\delta_{ms} + \beta_{ms} \ln \eta_{ms}) + \eta_{ms}\beta_{ms} \ln y, \\ W_{fs}(y) &= \eta_{fs}(\delta_{fs} + \beta_{fs} \ln \eta_{fs}) + \eta_{fs}\beta_{fs} \ln y. \end{aligned}$$

Implementation requires imposition of one or both of our identification restrictions. We impose $\beta_{ts} = \beta_s$ for all t to satisfy SAP as in Equation (5) or we impose $\beta_{ts} = \beta_t$ for all s to satisfy SAT as in Equation (8), or both. Both conditions are satisfied when $\beta_{ts} = \beta$ for all t, s .

Let a be a vector of 4 dummy variables for the 4 household types (indexed by s), and let z indicate the 14 demographic variables. For each person t , the resource shares η_{ts} and the intercept preference parameters δ_{ts} are specified as linear in a and z , so they have 18 coefficients each. There are no constant terms in the resource share functions or the intercept preference parameters—the levels are captured by the four household size dummies for households with 1-4 children.

The slope preference parameters β_{ts} are specified according to the identifying restriction: given SAP, β_{ts} is linear in a and z for a total of 18 coefficients; and given SAT, β_{ts} is linear in only a constant and z (15 coefficients) for each person t (3 people) for a total of 45 coefficients. Given both SAP and SAT, β_{ts} is linear in a constant and z for a total of 15 coefficients.

We implement the model by adding an error term to each equation of (10). These errors may covary across equations, so in the case with exogenous regressors, we estimate the model via nonlinear Seemingly Unrelated Regression (SUR), and with endogenous regressors, we use Hansen's (1982) Generalised Method of Moments (GMM). Both SUR and GMM estimators are iterated until the estimated parameters and error/orthogonality condition covariance matrices settle. Iterated SUR is equivalent to maximum likelihood with multivariate normal errors. We use the sum of clothing and footwear expenditures for each person as the private assignable good.

C. Results

We present estimates for η_{ts} in Table 2. Asymptotic standard errors are robust to heteroskedasticity of unknown form, and are given in *italics*. All estimated values of the coefficients on the constant term in beta β_{ts} are statistically significantly different from zero, which is important because nonzero latent slopes are required for identification of the resource shares. The leftmost block of Table 2 gives estimates using the SAP restriction, the middle block gives estimates using the SAT restriction, and the rightmost block imposes both SAT and SAP restrictions. We report only coefficients relating to the levels of resource shares in different household sizes, and coefficients relating to a few key demographic factors which potentially relate to policy levers: age and gender composition of children, and the education level of the parents (full estimation results are available on request from the authors). Parameters related to children's resource shares are computed off of the estimated values for adult resource share parameters, based on the restriction that resource shares sum to one.

Define a reference household as one in which $z = 0$, which is the case for animist/other households living in a village with both a daily market or a road, whose consumption recall was during the wet season, in which the man is aged 28 and woman is aged 22, and both have the modal level of education, and the children are all boys aged 5 (so that the average and minimum are both 5). For a reference household, the resource share is given by the number-of-children term in η_{ts} . In the Table, we report the level of the resource share in households of various sizes for the man η_{ms} , woman η_{fs} , all children $s\eta_{cs}$, and each child η_{cs} . For the demographic factors, we report only the effect on the resource shares of all children.

Consider first the rightmost block which presents the estimates given both the SAP and SAT (SAP&SAT) restrictions. Looking at the coefficients giving the level of resource shares in reference households of different sizes, we see that, roughly speaking, as the number of children increases, the total share of household resources devoted to children goes up, but the average share devoted to each child declines. A reference household with one child directs 22.7 per cent of its expenditures to children's consumption. With

two children, this share rises to 31.7 per cent, and four children, to 43.4 per cent. Even with three or four children, the resource share per child remains about 11 per cent or more.

Table 2: Estimates from Malawian Clothing (inc Footwear) Budget Shares

| | | SAP | | SAT | | SAP&SAT | |
|-----------------------------|------------|----------|---------|----------|---------|----------|---------|
| | | Estimate | Std Err | Estimate | Std Err | Estimate | Std Err |
| one child | man | 0.443 | 0.048 | 0.378 | 0.076 | 0.400 | 0.045 |
| | woman | 0.308 | 0.041 | 0.368 | 0.062 | 0.373 | 0.042 |
| | children | 0.249 | 0.037 | 0.254 | 0.072 | 0.227 | 0.036 |
| | each child | 0.249 | 0.037 | 0.254 | 0.072 | 0.227 | 0.036 |
| two children | man | 0.423 | 0.051 | 0.436 | 0.090 | 0.462 | 0.051 |
| | woman | 0.222 | 0.042 | 0.212 | 0.056 | 0.221 | 0.043 |
| | children | 0.355 | 0.045 | 0.352 | 0.100 | 0.317 | 0.045 |
| | each child | 0.177 | 0.022 | 0.176 | 0.050 | 0.158 | 0.023 |
| three children | man | 0.427 | 0.057 | 0.437 | 0.099 | 0.466 | 0.053 |
| | woman | 0.185 | 0.046 | 0.166 | 0.054 | 0.176 | 0.044 |
| | children | 0.388 | 0.050 | 0.397 | 0.114 | 0.358 | 0.050 |
| | each child | 0.129 | 0.017 | 0.132 | 0.038 | 0.119 | 0.017 |
| four children | man | 0.318 | 0.070 | 0.352 | 0.112 | 0.384 | 0.063 |
| | woman | 0.214 | 0.054 | 0.168 | 0.062 | 0.182 | 0.052 |
| | children | 0.468 | 0.061 | 0.479 | 0.133 | 0.434 | 0.059 |
| | each child | 0.117 | 0.015 | 0.120 | 0.033 | 0.109 | 0.015 |
| min. age of children | man | -0.005 | 0.010 | 0.007 | 0.010 | 0.008 | 0.009 |
| | woman | -0.005 | 0.008 | -0.014 | 0.008 | -0.014 | 0.008 |
| | children | 0.010 | 0.006 | 0.007 | 0.007 | 0.007 | 0.006 |
| avg. age of children | man | 0.006 | 0.010 | -0.007 | 0.010 | -0.008 | 0.009 |
| | woman | 0.006 | 0.008 | 0.017 | 0.008 | 0.017 | 0.008 |
| | children | -0.012 | 0.006 | -0.010 | 0.008 | -0.009 | 0.006 |
| proportion girl children | man | 0.006 | 0.029 | 0.001 | 0.031 | -0.003 | 0.028 |
| | woman | 0.053 | 0.024 | 0.058 | 0.027 | 0.056 | 0.026 |
| | children | -0.059 | 0.020 | -0.059 | 0.025 | -0.053 | 0.019 |
| man education | man | 0.021 | 0.009 | 0.008 | 0.010 | 0.008 | 0.010 |
| | woman | -0.009 | 0.008 | 0.003 | 0.009 | 0.002 | 0.009 |
| | children | -0.012 | 0.006 | -0.011 | 0.007 | -0.010 | 0.006 |
| woman education | man | -0.022 | 0.012 | -0.050 | 0.012 | -0.049 | 0.011 |
| | woman | 0.007 | 0.010 | 0.030 | 0.012 | 0.032 | 0.011 |
| | children | 0.015 | 0.008 | 0.020 | 0.010 | 0.017 | 0.008 |

Although the total resources of parents roughly decline with the number of children, this is not spread evenly across men and women. Men absorb between 40 per cent and 47 per cent of household resources if there are 3 or less children. Given the standard errors, this is a relatively small amount of variation. In contrast, women see their resource shares drop by about 20 percentage points as the number of children goes from 1 to 3. These

patterns are reasonably consistent under the SAP and SAT assumptions individually. One difference is that there is a large (but statistically insignificant) drop in men's resource share in households with 4 children under just the SAP assumption. A second difference is that the estimated levels of resource shares are much less precisely estimated under SAT than under the SAP or SAT&SAP cases, with standard errors that are almost twice as large. This is consistent with our earlier discussion regarding the comparative weakness of SAT identification.

Turning to the covariates, three observations stand out. First, the coefficients relating to the proportion of children who are girls are important. In particular, if all children in the household are girls, then their combined resource share is about 6 percentage points lower than if the children are all boys. These resources are almost fully diverted to the woman (the man's resource share is almost unaffected).⁹ Thus, unlike Deaton (1989, 1997) but similar to Rose (1999), we find statistically significant evidence of gender asymmetry in consumption within the household. One difference between our finding and that of Rose (1999), is that we find that gender asymmetry is the *status quo* and does not arise only in response to household income shocks.

Second, the higher the mother's level of education the more resources are diverted from fathers – with these extra resources being allocated 2/3 to mothers and 1/3 to children. These effects have reasonably large magnitudes. If a woman moves from the median to the top decile of education (from 0 to 2), the man's resource share declines by 10 percentage points. In contrast, we see little difference in resource shares from differences in men's education. The magnitude of education effects depends on which identifying assumption is used: the effects of women's education are much smaller (though still statistically significant) given the SAT assumption alone.

Third, a higher variance in the age distribution of children tends to increase the mother's share of resources. If the minimum age of children in the household rises by one year, women lose a 1.4 percentage point share of resources. Conversely, when the average age of children in the household rises by one year, women gain 1.7 percentage points. The estimates also suggest that these resources are diverted to men and children in roughly equal measure although this division is not statistically significant. These estimates imply that women tend to receive higher shares of household resources when both young and old children are present. As with the education effects, the estimated magnitudes are smaller given the SAT assumption alone.

One interesting hypothesis is whether the resource share functions depend linearly on the number of children. In this case, the resource share functions are linear in a constant, s , and the 14 demographic variables, which involves 2 exclusion restrictions in each of 2 resource share functions. We test this hypothesis using the SAP and SAT model (rightmost column of Table 2) and find the sample value of the likelihood ratio test statistic for this hypothesis is equal to 0.6 which is less than its χ^2_4 5 per cent critical value of 9.5. In contrast, the sample value of the Wald test statistic for this restriction is

⁹We note that part of this finding is somewhat specification dependent. For various specifications of the list of included demographic variables, we find that children's resources always respond negatively to the proportion of girls. However, whether these resources are diverted to the man or the woman is specification-dependent.

13.2 (and also a χ^2_4). Thus, it may or may not be reasonable to model resource shares as linear in the number of children. We conclude that imposing the restriction that resource shares are linear in the number of children does not do undue violence to the data, and will be a useful restriction later when we consider dealing with possible endogeneity in the number of children (see Section 4.5). It is difficult to find sufficient instruments to instrument for each of the dummy variables corresponding household size, but it is feasible to find instruments for a scalar-valued number of children.

D. Testing Model Assumptions

To check for possible violations of our modeling assumptions, we performed a large number of statistical specification tests, both directly on our data and in comparison with auxiliary data sets including single men and women, childless couples, and single mother households. To save space we only summarize the main results here. Details of test statistics and associated significance levels and tables of estimated parameters are provided in an online supplemental appendix.

IS HOUSEHOLD DECISION-MAKING EFFICIENT?

Like BCL, we assume that households divide resources Pareto-efficiently among members having individual utility functions, with a technology for sharing consumption goods. We show in the online supplemental appendix that, given our other assumptions and our chosen functional form, efficiency implies one of two restrictions on how the behavior of singles should compare to that of couples. With SAP, the slopes of household budget share Engel curves for men's clothes and women's clothes should have the same sign. With SAT, the slopes of household demands for men's or women's clothing will be proportional to individual demands, with factors of proportionality summing to one. Empirically we are unable to reject either restriction. We note that BCL is an example of a Pareto-efficient collective household model, and such models have repeatedly been found to be empirically satisfactory, particularly when compared to unitary household models. See, e.g., Browning and Chiappori (1998), among many others.

ARE RESOURCE SHARES INVARIANT TO EXPENDITURE?

Our identification strategy requires that resource shares are invariant to expenditure below a threshold-level of expenditure, $y^*(p)$. The invariance of resource shares to expenditure is often invoked for identification (Lewbel and Pendakur (2008) and Bargain and Donni (2009)) or is imposed in empirical application (Lise and Seitz (2004) and BCL). We show in the online supplemental appendix that there are reasonable structural models of household decision-making that imply invariance. Ultimately, however, whether resource shares are invariant to expenditure is an empirical question and thus we test whether resource shares are invariant to expenditure in our dataset.

We re-estimate the rightmost column of Table 2 (SAP and SAT) including an additional dummy variable in η_{ts} , δ_{ts} and β . This new dummy variable is equal to one if

the household is above the median of the total-expenditure distribution. This "higher-expenditure" dummy could be a legitimate demographic variable entering δ_{ts} and/or β , if resource shares. But, if η_{ts} depends on it then our identifying restriction is violated. We conduct a Wald test of the hypothesis that the dummy may be excluded from the resource shares. We do not reject the hypothesis and thus we conclude that the assumption that resource shares are invariant to expenditure is acceptable given our data.

ARE SAP AND/OR SAT VALID RESTRICTIONS?

SAT by itself yields overidentifying restrictions when there are more household sizes than types of household members. This is the case in our data. Empirically, we do not reject the overidentifying restrictions which implies that SAT is a reasonable assumption given our data.

BCL given either SAP or SAT implies overidentifying restrictions if there is more than one assignable good. Essentially, the restriction is that the resource shares must be the same no matter which assignables are used to identify them. Using clothing and footwear separately, we test these restrictions. The results of this test were mixed and hence indeterminate: the same hypothesis was rejected by the Wald test and not rejected by the likelihood ratio test.

SAP and SAT are restrictions on the preferences of individuals, so we additionally test if these restrictions are satisfied by single men and single women living alone. We only require SAP or SAT to hold for couples with children, but we can have more confidence in these restrictions if they are found also to hold for single men and single women. An added advantage of testing with single men and women is that the complications associated with the presence of shared and public goods within a household do not arise with singles. We can test SAP by comparing single men and single women to each other in one wave of data, and we can test SAT by separately comparing single men and single women across two time periods. Neither SAP nor SAT is rejected using our data on singles.

We also test whether or not the combination of SAP and SAT is much worse than either restriction alone. The estimates given both SAP and SAT are the most precise of the estimates presented because more identifying restrictions are imposed than with either SAP or SAT alone. These estimates allow for the sharpest testing of hypotheses about the behaviour of resource shares across household size and so we would ideally prefer to use both restrictions in estimation if the data suggest that both restrictions hold.

We can estimate the model (for a single private assignable good) under SAP and conduct a Wald test on the hypothesis that the coefficients on the household size dummies in β are identical for the 4 household types. Alternatively, we can estimate under SAT and conduct a Wald test of the hypothesis that the β_t are the same for all persons t . We do not reject either hypothesis and thus we conclude that the combination of SAP and SAT is a reasonable restriction for our data.

Taken together, these test suggest that both SAP and SAT are acceptable assumptions for our data, so to minimize variance we impose both for most of our analyses. These results should not be surprising since, as discussed earlier both restrictions (but particu-

larly SAP) are closely related to shape invariance, which is a well documented empirical regularity in the Engel curve literature (see Blundell, Chen, and Kristensen (2007)).

IS CLOTHING A PRIVATE ASSIGNABLE GOOD?

For identification we assume that clothing is an observable private assignable good. In our context, observability and assignability is very likely to hold (apart from the usual possibility of measurement errors associated with recall surveys), since we exclude households with older children and our data separate men's, women's, and children's clothing. A more serious concern for our model is that privateness may be violated either by the direct sharing of clothing, or by the externality of some household members deriving utility specifically from the clothing worn by others.^{10, 11}

Clothing is often used as a private assignable good in household demand estimation. For example, Kooreman (2000) says that clothing "is generally considered to be an 'assignable' good with a low degree of publicness; cf., Browning et al. (1994)." Nevertheless, tests of the 'privateness' assumption for clothing appear equivocal. BCL and Logan (2008) suggest that clothing is private while Donni (2009) finds some statistical evidence against the privateness assumption (although he also reports that the relevant departures in elasticities are only marginally statistically significant). Ultimately, this question must be decided empirically with our data.

First consider a violation of the privateness assumption due to sharing of clothing. We argue that footwear is less shareable than clothing because wearing a shoe is less a question of style and more a question of fit. As noted above, we separate footwear from clothing in the data and use the overidentifying information to test whether the estimated resource shares are identical for both goods. The estimated household resource shares recovered from our model using just clothing Engel curves should equal those based on footwear Engel curves, and should equal those based on the sum of clothing and footwear. Likelihood ratio tests fail to reject these equalities. Footwear represents a very small fraction of total expenditures, and the footwear Engel curves are poorly estimated, so this test does not have much power to reject in our context. Nonetheless, this test does not suggest that clothing is shared in our data.

Now consider a violation of the privateness assumption due to externalities in clothing consumption. To test for externalities, we compare estimates of preferences, i.e., Engel curve coefficients, based on single-mother households versus two-parent households. This addresses the externality issue because dependence of husband's utility on wife's or children's clothing (or vice versa) would in our model appear as a difference in individual's preference Engel curve coefficients across these two household types. We find that the Engel curve coefficients β_i for women who are single mothers is not statistically significantly different from that of women who are in two-parent households. Further, the patterns we observe in two-parent household resource shares regarding children also ap-

¹⁰Note that our model does permit deriving utility from the total well-being of each other household member.

¹¹One other issue regarding clothing is durability. Our analysis equates expenditures with consumption, which will not be a problem as long as the rates at which clothes are purchased and the rates at which they are consumed and wear out are comparable.

pear in single-mother households. Thus, we find no statistical evidence that externalities are important in clothing demands in these data.

In addition to these direct tests, the extensive tests of the BCL, SAP, and SAT assumptions summarized above are, in our context, really joint tests of both preference restrictions and private assignability of clothing, and so many of these tests should also have failed if privateness of clothing was not a reasonable assumption in our data. We conclude that non-privateness of clothing, if present, is not large enough to be statistically significant in our data.

E. Dealing with Endogeneity

Our models can be readily extended to deal with endogeneity via instrumental variables. One source of endogeneity in our setting is that total expenditures can suffer from measurement error, either because of infrequency of purchases creating a wedge between total expenditures and actual consumption, or because of recall errors, since total consumption is measured by asking households to recall their past expenditures.

Both sources of endogeneity in total expenditures y can be dealt with using wealth measures as instruments. The utility function \tilde{U}_s in equation (??) applies to a single time period t , so denote that \tilde{U}_{st} . Assume that the household determines its true total consumption expenditures in each period by maximizing the expected value of an additively time separable utility function $\sum_t b^t \tilde{U}_{st}$ with rate of time preference b , subject to a budget constraint determined by wealth. Then true total expenditures will be a function of wealth. Wealth will be then a valid instrument (uncorrelated with within period consumption allocation errors) if the consumption allocation decisions within a period are separable from savings decisions across periods. This is why wealth measures are commonly used as instruments for total consumption in demand system estimation.

In our data, wealth measures are also suitable as instruments for recall errors. Unlike expenditures, wealth in our data is measured by enumerating physically observed assets of the households such as farm animals and capital goods like tools, etc., and therefore is not subject to recall error. While wealth may also be mismeasured, due to omission of some items or incorrect valuation of others, it will remain valid as an instrument if these omission or valuation mismeasures are independent of consumption recall errors and if true wealth is correlated with true total expenditures.

Another potential source of endogeneity is a possible correlation between the number of children in the household and the residuals in the clothing equations. In particular, if unobserved preference heterogeneity is connected to both fertility decisions and expenditure (on clothing) decisions, then the number of children in the household will be endogenous. We use measures of access to medical care and medical information as instruments for household size. These should be valid instruments because such access is known to affect fertility decisions, and it is hard to see why unobserved heterogeneity in clothing preferences would be correlated with medical access.

As we show below, statistical tests do not reject the hypothesis that these instruments are valid, and our key empirical findings are not much changed when we account for potential endogeneity in total expenditures and household size.

Begin with equation (10), imposing SAP and SAT so $\beta_{ts} = \beta$, meaning that the latent slope parameters do not vary with either the individual t or the number of children s . The subscript s can then be absorbed into regressors and coefficients (specifically, s appears in the household size dummy variables a which are inside the resource shares η_{ts} , and in the latent intercepts δ_{ts}) so that the error terms from estimation for each person's assignable good, e_t , do not need an s subscript.

Let $q_t = q_{t1}, \dots, q_{tJ_t}$ be an J_t -vector of instruments uncorrelated with the error terms, e_t . These instruments can be any functions of any variables that are conditionally exogenous with respect to e_t . Then, $E(e_t q_{tj}) = 0$ for all t, j implies for our model:

$$E[(W_{cs} - s\eta_{cs}(\delta_{cs} + \ln \eta_{cs}) - s\eta_{cs}\beta \ln y) q_{cj}] = 0,$$

for $j = 1, \dots, J_c$, and

$$E[(W_{ts}(y) - \eta_{ts}(\delta_{ts} + \ln \eta_{ts}) - \eta_{ts}\beta \ln y) q_{tj}] = 0,$$

for $t = m, f$ and $j = 1, \dots, J_m$ and $j = 1, \dots, J_f$. With these moment conditions, the parameters may be estimated by Hansen's (1982) generalised method of moments (GMM).

Optimal instruments for these moment conditions (based on the first order conditions for minimizing a quadratic criterion function) would correspond to the derivatives of the error terms e_t with respect to the model parameters η_{ts} , δ_{ts} and β . To improve efficiency, we construct instruments that are close to optimal by suitable transformations of our observed instruments. In particular, we evaluate these derivatives at SUR pre-estimates, and plug in "hat" versions of endogenous variables rather than their true values, where "hat" versions are (first-stage) OLS predictions of the endogenous variables on the basis of all observed exogenous variables. This is essentially equivalent to the first stage of two stage least squares, when the first stage equations are nonlinear. Note that our models are overidentified, since β is found in all 3 equations, and η_{ms} and η_{fs} are each found in 2 equations (due to the summation restriction on η_{ts}).

Our exogenous variables include: the log of expenditure (except in models where we treat it as endogenous), all 14 demographic variables, the log of the value of livestock holdings, the log of the value of durable goods holdings, the log of the sum of livestock and durable holdings, the presence in the village of a HIV-prevention oriented NGO office, the distance to a doctor's office and a dummy variable indicating that the woman has a chronic illness. As discussed earlier, the first two of these are wealth measures, while the rest are medical and health related instruments to correlate with fertility decisions.

Our endogenous regressors are either the number of children in the household or both the number of children in the household and the log of total expenditure. These instruments are not very strong in predicting the number of children in the household in that, conditional on the demographic variables and the log of expenditure, the F statistic on the excluded instruments in the first stage is only 2.5. However, these instruments are very strong in predicting the log of expenditure: the F statistic on the excluded instruments is 67.

Table 3: GMM Estimates

| | | SUR | | GMM | | GMM | |
|--------------------------|----------|----------|---------|--------------------|--------------------|-------------------------|-------------------------|
| | | Estimate | Std Err | endog: extra child | endog: extra child | endog: extra child, lny | endog: extra child, lny |
| | | Estimate | Std Err | Estimate | Std Err | Estimate | Std Err |
| one child | man | 0.456 | 0.045 | 0.407 | 0.056 | 0.341 | 0.074 |
| | woman | 0.358 | 0.044 | 0.427 | 0.054 | 0.408 | 0.071 |
| | child | 0.186 | 0.030 | 0.166 | 0.044 | 0.251 | 0.073 |
| extra child | man | -0.012 | 0.018 | 0.083 | 0.085 | -0.008 | 0.095 |
| | woman | -0.055 | 0.015 | -0.148 | 0.073 | -0.075 | 0.098 |
| | children | 0.068 | 0.014 | 0.065 | 0.040 | 0.083 | 0.042 |
| min. age of children | man | 0.003 | 0.009 | 0.056 | 0.040 | 0.004 | 0.043 |
| | woman | -0.007 | 0.008 | -0.056 | 0.034 | 0.000 | 0.044 |
| | children | 0.004 | 0.006 | 0.000 | 0.019 | -0.004 | 0.019 |
| avg. age of children | man | -0.004 | 0.009 | -0.058 | 0.040 | -0.010 | 0.043 |
| | woman | 0.009 | 0.008 | 0.058 | 0.035 | 0.007 | 0.044 |
| | children | -0.005 | 0.006 | -0.001 | 0.019 | 0.003 | 0.019 |
| proportion girl children | man | -0.015 | 0.030 | 0.030 | 0.033 | -0.026 | 0.038 |
| | woman | 0.063 | 0.029 | 0.026 | 0.027 | 0.090 | 0.040 |
| | children | -0.048 | 0.016 | -0.056 | 0.024 | -0.065 | 0.033 |
| man education | man | 0.008 | 0.010 | 0.020 | 0.010 | 0.013 | 0.012 |
| | woman | -0.001 | 0.010 | -0.016 | 0.010 | -0.006 | 0.012 |
| | children | -0.008 | 0.005 | -0.004 | 0.005 | -0.007 | 0.008 |
| woman education | man | -0.047 | 0.011 | -0.044 | 0.012 | -0.058 | 0.014 |
| | woman | 0.033 | 0.011 | 0.028 | 0.012 | 0.042 | 0.015 |
| | children | 0.014 | 0.006 | 0.016 | 0.007 | 0.016 | 0.009 |

Table 3 gives estimates of resource share parameters in models where children enter the resource shares and latent intercepts (η_{ts} and δ_{ts}) linearly, so that these functions have 16 parameters each (a constant, the number of children s , and the 14 demographic shifters). The leftmost column presents SUR estimates analogous to those presented in Table 2, and the middle and rightmost columns give GMM estimates corresponding to instrumenting either the number of children, or both the number of children and the log of expenditure. Hansen J-tests of overidentifying restrictions do not suggest that the instruments are endogenous and hence do not reject the hypothesis of instrument validity (with p-values of 51% and 60% for the middle and rightmost columns, respectively).

In the leftmost column, the SUR estimates show that most of the results in Table 2 are evident when we replace the household size dummies with the scalar-valued number of children variable. In particular, we see that men's resource shares do not respond to the number of children, but women's decline substantially and statistically significantly with the number of children, so that each additional child increases the children's resource share by almost 7 percentage points.

The middle and rightmost columns show GMM estimates which account for the possible endogeneity of household size and both household size and the log of expenditure, respectively. In general, the patterns relating to household size are still visible, but are

estimated very much less precisely (particularly so when both size and expenditure are treated as endogenous). Again, children's resources are marginally statistically significantly increasing in the number of children, with about the same magnitude as in the SUR regressions. Men's shares are not statistically significantly related to the number of children and women's shares are statistically significantly declining if only household size is considered as endogenous.

The fact that the GMM regression estimates are very similar to the OLS estimates suggests that there may not be significant endogeneity in the number of children in the household. Hausman tests support this claim: sample value of the Hausman test statistic for the hypothesis that the parameter estimates are the same in the SUR and the GMM regressions is equal to 80 in the middle regression and equal to 116 in the rightmost regressions. In either case, the test statistic is distributed as a χ^2_{95} with a 5 per cent critical value of 119 under the null hypothesis. The lack of significant endogeneity may be due in part to timing issues - resource allocations may be decided long after fertility decisions.

Our key take-away from this exercise is that the importance of the education and child gender covariates remains unchanged. As one might expect, the reduced precision of GMM relative to SUR does not take as large a toll on the estimates of parameters associated with the exogenous regressors. The gender bias in children's resources is evident across all specifications: if the children are all girls, they absorb about 5 – 7 percentage points less of household resources. Whether it is men or women who gain at the expense of girls is less clear, although in all specifications the change in the men's share is not statistically significant. We also see a substantial effect of women's education, diverting resources towards women and children. If the women's education increases from the median to the 90th percentile (2 units), men's share of household resources falls by roughly 10 percentage points. These resources are shared roughly 2/3 and 1/3 by women and children, respectively.

F. Resource Shares, Poverty Rates and Child Poverty

The empirical results described so far relate to the levels of resource shares for persons in reference households, and to the marginal effects of various demographic factors. However, this does not tell us how resource shares would change in aggregate across household sizes because the demographic factors themselves covary with household size. To evaluate, for example, whether men or women make the larger sacrifice of consumption for their children, it is illustrative to consider the average resource shares in households of different sizes, averaging over all the values of demographic factors observed in the population.

Table 4: Estimated Resource Shares and Poverty Rates

| | | Mean | Std Dev | Min | Max | Pov Rate <i>Unequal</i> | Pov Rate <i>Equal</i> |
|----------------|------------|-------|---------|-------|-------|----------------------------|--------------------------|
| one child | man | 0.463 | 0.087 | 0.245 | 0.762 | 0.686 | 0.850 |
| | woman | 0.402 | 0.071 | 0.168 | 0.587 | 0.766 | |
| | children | 0.135 | 0.047 | 0.008 | 0.260 | 0.954 | |
| | each child | 0.135 | 0.047 | 0.008 | 0.260 | | |
| two children | man | 0.516 | 0.078 | 0.282 | 0.786 | 0.547 | 0.916 |
| | woman | 0.273 | 0.063 | 0.075 | 0.475 | 0.885 | |
| | children | 0.211 | 0.044 | 0.059 | 0.326 | 0.970 | |
| | each child | 0.105 | 0.022 | 0.029 | 0.163 | | |
| three children | man | 0.521 | 0.081 | 0.219 | 0.795 | 0.522 | 0.948 |
| | woman | 0.244 | 0.065 | 0.002 | 0.512 | 0.889 | |
| | children | 0.236 | 0.042 | 0.112 | 0.374 | 0.996 | |
| | each child | 0.079 | 0.014 | 0.037 | 0.125 | | |
| four children | man | 0.441 | 0.080 | 0.170 | 0.701 | 0.538 | 0.972 |
| | woman | 0.267 | 0.066 | 0.043 | 0.532 | 0.838 | |
| | children | 0.293 | 0.037 | 0.178 | 0.402 | 0.989 | |
| | each child | 0.073 | 0.009 | 0.044 | 0.101 | | |
| All Households | man | 0.489 | 0.088 | 0.170 | 0.795 | 0.582 | 0.913 |
| | woman | 0.304 | 0.093 | 0.002 | 0.587 | 0.842 | |
| | children | 0.207 | 0.070 | 0.008 | 0.402 | 0.974 | |
| | each child | 0.103 | 0.038 | 0.008 | 0.260 | | |
| All Persons | all | 0.235 | 0.177 | 0.008 | 0.795 | 0.855 | 0.924 |

The leftmost columns of Table 4 presents summary statistics on the estimated values of resource shares for people in households of different sizes, using the SAP&SAT results from Table 2. It is comforting to see that the minima and maxima of estimated resource shares do not fall outside the zero to one range for any person in any household in the sample. The standard deviations are quite small in most cases. Interestingly, the standard deviations of resource shares are larger for men than for women in all household sizes. Thus, the demographic variables are not very important in terms of their effects on resource shares, though they do induce more variation for men than for women. Much more important than these factors are the household sizes themselves. This suggests that our ability to identify the *level* of resource shares, rather than just their response with respect to demographic or distribution factors, is particularly important.

The rightmost columns of Table 4 show the estimated poverty rates (at the household level) for households of different sizes. To define poverty, we use the World Bank's US\$2/person/day poverty threshold. This measure (Equal) assumes that each household member gets an equal share of household income.

The next column to the left (Unequal) uses our resource shares to construct person-level expenditures (equal to household expenditure times the resource share) and compares this to the US\$2/day threshold. To account for the possibility that children may have lower needs than adults, we use the OECD estimate of the relative needs of children

(60 % that of adults), and so for children, compare their expenditure to US\$1.20/day.

The bottom block and bottom row give the estimated poverty rate for all households together and for all persons. Here, we see a poverty rate of 91.3% for households in our sample under the "equal-division" rule. For comparison, the World Bank reported poverty rate for all households in Malawi in 2004 was 90.5%. We now consider how accounting for intrahousehold inequality changes these per capita estimates.

There are at least three features to note in our poverty estimates. First, Table 4 shows that there are a lot more households with poor women than with poor men. For example, looking at the rows for All Households, we see that 58.2 per cent of households have a poor man, but 84.2 per cent of households have a poor woman. Second, the poverty rates of men seem to drop with household size, but the poverty rate for women and children is roughly rising with household size. Third, more households have poor children than have poor adults. In households with 3 or 4 children, nearly all children are poor. Indeed our estimates are that incidence of child poverty in Malawi is over fifty percent higher than the rate of poverty for adult men, and roughly 1/6 higher than that of adult women.

We considered a similar exercise for the three regions of Malawi and find somewhat contrasting pictures of poverty. In the North, we observe lower incidences of poverty for men (38%), but higher incidences for women and children (93% and 99%, respectively). In contrast, in the South, we observe higher incidences of poverty for men (76%) and relatively similar incidences for women and children in comparison to the North (90% and 97%). In the Central region women tend to fare best in a relative sense (a poverty incidence of 75%) while men are slightly worse off than in the North (44%) and children are essentially identical to the South (97%).

The major conclusion here is that the intra-household allocation may be very important to measuring child poverty. We do not wish to emphasize the absolute levels of poverty too much since they depend on a measure of the relative needs of each household member. It may be the case that the relative needs of each member differ from the levels we have used above. For example, they may depend on daily calorie requirements; see S. Paul (1989). Thus, as a practical matter, the researcher should check for robustness with respect to other reasonable needs measures¹². However, the fact that we can now measure children's resource shares within households is a very useful step in measuring poverty, and in measuring the full effect of policy interventions aimed at poverty alleviation.

IV. Conclusions

Child poverty is at the root of much inequality. Children are also among the least able in society to care for themselves. Despite the apparent importance of understanding the intra-household dimension of child inequality, very little research has focused on children's shares of household resources. Most collective household models either ignore children, or treat them as public or private goods for adults.

¹²While we do not formally present robustness analysis along this dimension, our major conclusions regarding child poverty hold up when the needs of children are assumed to be either 50% or 70% as much as adults. In particular, in both alternative specifications, we observe that child poverty is higher than men's poverty and that child poverty rises with the number of children in the household.

We propose a collective household model in which children are people with their own utility functions. Children's resource shares within the household are identified given household level Engel curve data on private assignable goods. In particular, by looking at how the budget shares for men's, women's and children's clothing vary across households with differing income levels and numbers of children, our structural model allows us to back out an estimate of the fraction of *total* household expenditure that is consumed by each family member on *all* goods they consume.

Using household consumption data for Malawi, we find that children command a reasonably large share of household resources (though not enough to avoid having higher rates of poverty than their parents) and that the share of resources devoted to children rises with the number of children, though the average share per child tends to decrease. Mothers appear to contribute more resources than fathers to children, and we find some evidence of gender-bias in children's resource shares. We also find that there is substantial intra-household inequality, one consequence of which is that standard per-capita poverty indices, which by construction ignore intra-household inequality, present a misleading picture of poverty, particularly for children.

V. Appendix

This paper has a supplementary online appendix that contains five main sections:

Appendix A.1 states formally our main identification theorems, providing the general conditions for identification under either the Similar Across People (SAP) or Similar Across Types (SAT) conditions.

Appendix A.2 gives proofs of the Theorems in Appendix A.1.

Appendix A.3 provides an example of a class of indirect utility functions that satisfy the assumptions of both of our main identification Theorems, and yields Engel curves having the functional form we assume for our empirical work.

Appendix A.4 provides an example functional form within the general class of models given in appendix A.3. This functional form could be used if one wished to combine our results with other structural analyses, e.g., if one wished to introduce prices into the model.

Appendix A.5 provides the results of extensive statistical tests of the adequacy of our model's assumptions. These are divided into sets of tests focusing on the general BCL methodology, our SAP and SAT restrictions on preferences, invariance of resource shares with respect to total expenditures at low total expenditure levels, and privateness of clothing expenditures.

VI. References

Bargain, O. and O. Donni, (2009), "The Measurement of Child Costs: A Rothbarth-Type Method Consistent with Scale Economies," Institute for the Study of Labor (IZA) discussion paper No. 4654.

Becker, G. (1965), "A Theory of the Allocation of Time," *Economic Journal*, 75: 493-517.

Becker, G. (1981), "Altruism in the Family and Selfishness in the Market Place," *Economica*, 48: 1-15

Blundell, R., X. Chen and D. Kristensen (2007), "Semi-Nonparametric IV Estimation of Shape Invariant Engel Curves," *Econometrica* 75: 1613-1669.

Blundell, R., P.-A. Chiappori and C. Meghir (2005), "Collective Labor Supply with Children," *Journal of Political Economy*, 113: 1277-1306.

Bourguignon, F., M. Browning, and P.-A. Chiappori (2009), "Efficient Intra-household Allocations and Distribution Factors: Implications and Identification," *Review of Economic Studies*, 76: 503-528.

Bourguignon, F. and P.-A. Chiappori, (1994) "The Collective Approach to Household Behavior," in *The Measurement of Household Behavior*, R. Blundell, I. Preston, and I. Walker, Eds., Cambridge: Cambridge University Press.

Browning, M., (1992), "Children and Household Economic Behavior," *Journal of Economic Literature* 30: 1434-1475.

Browning, M., F. Bourguignon, P.-A. Chiappori, and V. Lechene (1994), "Incomes and Outcomes: a Structural Model of Intrahousehold Allocation," *Journal of Political Economy*, 102: 1067-1096.

Browning, M., and P.-A. Chiappori (1998), "Efficient Intra-household Allocations: A General Characterization and Empirical Tests," *Econometrica*, 66: 1241-1278.

Browning, M., P.-A. Chiappori and A. Lewbel (2008), "Estimating Consumption Economies of Scale, Adult Equivalence Scales, and Household Bargaining power," mimeo.

Cherchye, L., B. De Rock and F. Vermeulen (2008), "Opening the black box of intra-household decision-making: Theory and non-parametric empirical tests of general collective consumption models," *Journal of Political Economy*.

Cherchye, L., B. De Rock and F. Vermeulen (2010), "The revealed preference approach to collective consumption behavior: testing and sharing rule recovery", forthcoming in *Review of Economic Studies*

Chiappori, P.-A., (1988), "Nash-Bargained Households Decisions: A Comment," *International Economic Review*, 29: 791-796.

Chiappori, P.-A., (1992), "Collective Labor Supply and Welfare," *Journal of Political Economy*, 100: 437-467.

Chiappori, P.-A. and I. Ekeland (2009), "The Micro Economics of Efficient Group Behavior: Identification," *Econometrica*, 77: 763-99.

Dauphin, A., B. Fortin, G. Lacroix and A-R. el Lahga (2008), "Are Children Decision-Makers Within the Household", IZA Discussion Paper number 3728.

Deaton, A. (1989), "Looking for Boy-Girl Discrimination in Household Expenditure Data," *World Bank Economic Review*, 3(1): 1-15.

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

Donni, O. (2009), "A Simple Approach to Investigate Intrahousehold Allocation of Private and Public Goods," *The Review of Economics and Statistics*, 91: 617-628.

Duflo, E. (2003), "Grandmothers and Granddaughters: Old Age Pension and Intra-Household Allocation in South Africa," *World Bank Economic Review*, 17(1), 2003:

1-25.

Duflo, E. (2005), "Gender Equality in Development," unpublished notes.

Kooreman, P. (2000), "The Labeling Effect of a Child Benefit System," *American Economic Review*, 90: 571-583.

Lazear, E. P. and R. T. Michael (1988), *Allocation of Income Within the Household*, University Chicago Press.

Lewbel, A. (2003), "Equivalence Scales Based on Collective Household Models," in *Household Behavior, Equivalence Scales, Welfare and Poverty*; Camilo Dagum and Guido Ferrari, eds., Physica Verlag: New York: 1-9.

Lewbel, A. and K. Pendakur (2008), "Estimation of Collective Household Models With Engel Curves," *Journal of Econometrics*, 147: 350-358.

Lewbel, A. (2010), "Shape Invariant Demand Functions," *Review of Economics and Statistics*, 92: 549-556.

Lindahl, E., (1919), "Just Taxation - A Positive Solution," in R. Musgrave and A. Peacock, eds., *Classics in the Theory of Public Finance*, London: MacMillan.

Lise, J. and S. Seitz, (2004). "Consumption Inequality and Intra-Household Allocations," Working Papers, 1019, Queen's University, Department of Economics.

Logan, T. D. (2008), "Economies of Scale in the Household: Puzzles and Patterns From the American Past," NBER working paper 13869.

Lundberg, S.J., R. A. Pollak and T. J. Wales (1997). "Do Husbands and Wives Pool Their Resources? Evidence from the United Kingdom Child Benefit," *Journal of Human Resources*, 32(3): 463-480.

Muellbauer, J. (1974), "The Political Economy of Price indices," Birbeck Discussion Paper no 22.

Muellbauer, J. (1976). "Community Preference Fields and the Representative Consumer." *Econometrica*, 44: 979-999.

National Statistics Office of Malawi (NSOM). (2005). "Integrated Household Survey Wave 2". <http://www.nso.malawi.net> (obtained June 2007)

Paul, S. (1989). "A model of constructing the poverty line," *Journal of Development Economics*, 30(1): 129-144.

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

Pendakur, K. (2002). "Taking Prices Seriously in the Measurement of Inequality," *Journal of Public Economics*, 86: 47-69.

Rose, E. (1999), "Consumption Smoothing and Excess Female Mortality in Rural India," *Review of Economics and Statistics*, 81(1): 41-49.

Samuelson, P. (1956), "Social Indifference Curves," *The Quarterly Journal of Economics*, 70: 1-22.

Vermeulen, F., (2002), "Collective Household Models: Principles and Main Results," *Journal of Economic Surveys*, 16: 533-564.