An empirical Bayes method to estimate resource shares of household members

Denni Tommasi and Alexander Wolf*

ECARES, Université libre de Bruxelles

Current Draft: January 25, 2016

Abstract

In a recent and insightful paper, Dunbar et al. (2013) develop a collective model of the household that allows to identify resource shares, that is, how total household resources are divided up among household members. We show why, especially when the data exhibit flat(ish) Engel curves, the model induces high variability and an implausible pattern in least squares estimates. We propose an estimation strategy nested in their framework that greatly reduces this practical impediment to recovery of individual resource shares. To achieve this, we follow an empirical Bayes method that incorporates additional (or out-of-sample) information of singles and rely on mild assumptions on preferences. Our welfare analysis of the PROGRESA program in Mexico is the first to include separate poverty rates for men and women in a CCT program.

JEL Codes: D13, D11, D12, C31, I32

Keywords: collective model, sharing rule, resource shares, demand system, engel curves, bayes method, shrinkage, conditional cash transfers, PROGRESA.

*ECARES, Université Libre de Bruxelles. Avenue F. D. Roosevelt 50, CP 114, B-1050 Brussels, Belgium. E-mail: dtommasi@ulb.ac.be and awolf@ulb.ac.be. We are grateful to Arthur Lewbel for his invaluable comments on this project which has been partially conducted at Boston College. We thank our supervisor Bram De Rock for constant advice, guidance and encouragement. We thank Bet Caeyers, Laurens Cherchye, Donald Cox, Yves Dominicy, Frederic Vermeulen, Mark Watson, participants at the NEUDC Conference at Brown University, and seminar participants at Boston College, ECARES, Paris School of Economics, University of Bologna and University of Leuven for useful comments and suggestions. Denni acknowledges financial support from the Fonds National de la Recherche Scientifique (FNRS). All remaining errors are ours.

1 Introduction

The study of intra-household inequality in consumption faces a fundamental data problem: information on consumption at the individual level are commonly absent in survey data. Advances in the literature of collective models of the household have offered a way to circumvent this problem: by imposing empirically supportable restrictions on preferences of the household members, it is possible to derive a system of demand equations that allows to estimate a quantity known as *resource shares*. This is a useful measure of individual welfare, defined as each member's share of total household consumption, and can be estimated directly from household level data.

The collective model of the household, pioneered by Chiappori (1988, 1992) and Apps and Rees (1988), and subsequently elaborated by Browning et al. (1994), Browning and Chiappori (1998), Blundell et al. (2005) and Chiappori and Ekeland (2006), is a growing and popular framework for analyzing household behavior[1]. It recognizes that households consist of individuals with own rational preferences, and uses the assumption that the intra-household decision process, whatever it may turn out to be, produces Pareto-efficient outcomes. Recently, collective models have been adapted in order to estimate resource shares. The pioneering work is Browning et al. (2013) (hereafter BCL) who provide a model that non-parametrically identifies the levels of resource shares of adult's household members. BCL makes the identification restriction that the preference structure of singles is identical to individuals in (child-less) couples. In practice, they observe the demand functions of single men and single women living alone, and combine those demand functions with data on the demands of men and women living together as couples. Recovery of the parameters relating to sharing and efficiencies of scale is then much easier since preferences are already given[2]. Although the degree of restriction on the preferences structure can be mitigated by the choice of an "appropriate" dataset[3], BCL still remains a highly complex model and extraordinarily hard to estimate.

An alternative and attractive approach for the practitioner to estimate resource shares can be found in Dunbar et al. (2013) (hereafter DLP). The attractiveness of their strategy is due to the ability to combine a more general theoretical structure than BCL-type models with a lower data requirement and estimation complexity. The DLP model uses

[1]Donni and Chiappori (2011) and Browning et al. (2014) provide a comprehensive review of the theoretical and empirical advances of collective models.

[2]Papers that use this identification strategy include Lewbel and Pendakur (2008), Lise and Seitz (2011), Bargain and Donni (2012) and Cherchye et al. (2012). A completely different approach is taken by Cherchye et al. (2011) and Cherchye et al. (2015), who provide set identification of the resource shares on the basis of revealed preference theory. These papers are more general than BCL-type of approach, but are not yet standard in the literature.

[3]For example, in the case of Cherchye et al. (2012), the authors estimate the BCL model on a sample of elderly couples and elderly widow(er)s. For this particular sample selection, the somewhat restrictive assumption that preference of the singles are equivalent to preference of the couples becomes less stringent and more easily acceptable. The same sample selection is followed by Bütikofer et al. (2011) in their application of the Lewbel and Pendakur (2008) model.

information on private assignable goods and derive a structural model of demand to estimate the shares of household resources allocated to each family member. Resource shares are semi-parametrically identified by observing how household expenditures on private assignable goods of each member (e.g. clothing) varies by total expenditure. They impose two identifying restrictions: that resource shares do not vary with total expenditures[4] and one of two semi-parametric restrictions on individual preferences. With the first semi-parametric restriction, they assume that preferences for a particular assignable good are similar (in certain limited ways) across members within the same type of household, where a "type" is defined by the number of children present, and use this similarity to identify resource shares. They call this restriction "Similar Across People" (SAP). With the second, they assume that a person's preferences for a particular assignable good are similar across household types, and compare the consumption choices of people across households with varying number of children. They call this restriction "Similar Across Types" (SAT). Importantly, both these restrictions are much weaker than the above-cited identity between singles and couples of BCL-type of approach. Together with its ease of implementation, the relatively light data requirement and the fact that the strategy allows to recover shares for more than two household members, the DLP model currently offers the most attractive framework for the practitioner that aims to study intra-household dynamics[5].

In the present paper, we take the DLP framework and investigate its strengths and weaknesses in the recovery of resource shares. Our contribution is threefold. First we show that although the DLP model is identified, it is prone to a severe multicollinearity problem in the form of trade-offs between the estimates of resource shares and of tastes. This induces an implausible pattern in least squares estimates as well as large variability in estimated resource shares unless the sample of households exhibits large and systematic variability in the expenditure of the private assignable goods, like clothing. This creates a practical impediment to empirical applications because, regardless of the quality of the surveys collected, there is no guarantee that the population of interest has the required necessary variability in the consumption of these particular goods. This issue is much more detrimental than what can be thought at first because, in practice, it means that resource shares with the DLP framework cannot be estimated for a (potentially large) variety of datasets. Inefficiency of the estimates arises in two forms: large variability of the estimates of resource shares, which implies uncertainty about the location of the

[4]This is an important point discussed in the Appendix of DLP. This assumption requires that resource shares are independent of household expenditure y (at low level of y) but they do not require them to be constant. They are allowed to vary arbitrarily with prices p and other household characteristics. A class of utility functions that satisfies this condition, which is used in DLP, is the PIGLOG indirect utility functions of Muellbauer (1976). It is also worth noting that although resource shares must not depend on y, they are still permitted to depend on other variables closely related to y, such as household income, wealth, or member's wage.

[5]See Calvi (2015) for a recent application of the model to study female poverty within Indian families.

sharing rule, and a distortion of heterogenous parameter estimates resulting in an implausible negative correlation between the parameters of shares and taste difference between spouses.

Second, we offer a simple way to address these issues which requires mild preference restrictions of the DLP framework. The core of our approach lies in the inclusion of a shrinkage term in our estimation, following an empirical Bayes method a-la James and Stein (1961), and on imposing a non-correlation restriction on some parameters. The shrinkage term relies on a minimal aspect of single's behavior that we argue is likely to be similar in single and married individuals. By the very nature of this approach, we do not need to impose that any part of singles' preferences be invariant to marriage. Instead we nudge the estimator toward the quantity obtained from singles' data. This approach has recently received renewed attention (Fessler and Kasy, 2015), as it provides a middle ground between nested models which is particularly useful in our context. The non-correlation restriction between parameters of the resource shares and of the taste function is argued to be also mild. The effectiveness of our empirical strategy is tested through a series of Monte Carlo simulations where we show that our approach is robust even where the true data generating process (slightly) violates these restrictions.

Third, we apply our methodology to a sample drawn from the surveys collected to evaluate the impact of a conditional cash transfers program in a developing context. PROGRESA is a welfare program implemented in rural Mexico in the late 1990s, whose main objective was to fight poverty among marginalized households. The choice of this dataset for the empirical application is motivated by a scientific and a practical reason. First, it is generally of interest to study the short-term poverty effects of one of the largest welfare programs in a developing country context. There is a large body of evidence that PROGRESA has been beneficial to eligible households but there are no quantified estimates available in terms of effects on individual consumption and poverty. Second, from a practical point of view, the available surveys contain all the necessary information to implement the DLP framework and moreover, although they are generally considered of high quality, we show that this is no guarantee that the sample of households at hand has a sufficient variability in the expenditure of the private assignable goods, which is necessary to estimate individual resource shares with high precision. This is to indicate that, even with the best datasets, applying the DLP model may still be problematic. Using our approach, compared to the standard DLP strategy, we obtain much reduced standard deviations for virtually all parameters and a reduction in magnitude for the effect sizes of the treatment, which for the simple least square estimator were often quite strong. In contrast to least square estimates, in which the treatment effect on men's resource share was negative but very insignificant, we obtain a clear positive effect. The results of our welfare analysis on the effects of PROGRESA indicate that men benefited more from the treatment than did women as measured by individual poverty rates.

This paper lies at the intersection of three strands of literature. The first is the literature on collective models and intra-household allocations. Few papers provide identification of the levels of resource shares. With respect to this literature, our contribution is built on one of these by DLP, which provides a powerful framework for practitioners. With two mild restrictions on preferences we redress a multicollinearity issue of the model that can make identification of the resource shares very hard. In other words, we make the DLP framework applicable even in contexts where the dataset used for estimation exhibit flat(ish) Engel curves in the private assignable goods used for identification. This comes at the cost of needing some information on consumption of private assignable goods by singles, which is not required for identification. The use of a shrinkage term allows us to govern the strength of this restriction. The less weight is placed upon the prior information taken from singles, the closer we get to a modified DLP model that does not take singles into account. Interestingly, the more important the shrinkage becomes, the farther away we move from DLP and the closer we get to a model in which similarities between singles and married peoples preferences is assumed. This suggests that one possible way to interpret our contribution to the collective model literature is to see the empirical Bayes method that we introduce as a way to "shrink" the assumed individual's preference restrictions to one of the two existing (parametric) approaches to recover resource shares of household members. On one side (BCL) we have a strong assumption for identification but stable estimates; on the other side (DLP) we have a clever idea that relaxes the strong assumption but does not produce robust results yet. Our contribution is to find a way to combine the best of the two approaches and obtain more stable results.

The second strand is the statistical literature on empirical Bayes methods and shrinkages. Robbins (1956) was the first to consider the empirical Bayes approach to construct estimators, followed by the seminal contribution by James and Stein (1961) and further developments by Efron and Morris (1972). Morris (1983) was the first to discuss the parametric version of the empirical Bayes approach and provides the natural framework for many recent economic applications. A recent and up to date introduction to empirical Bayes estimation can be found in Efron (2010). Our approach is most closely related to a recent contribution by Fessler and Kasy (2015). Similar to these authors, we use economic theory to motivate and construct an empirical Bayes estimator within our DLP framework. Ideas related to this approach, in a fully Bayesian setting, have been used, for example, in the literature of macroeconomic forecasting, where DSGE models can be used to inform priors for the parameters of statistical VAR models fit to the data. For recent contributions using this approach see Del Negro and Schorfheide (2004) and Del Negro et al. (2007). With respect to this literature we extend the possible range of applications of common Bayesian tools, such as informative priors and shrinkage terms, and show how they may be employed to solve specific issues, such as our multicollinearity feature of the DLP model, proper to this emerging literature of collective models of the household.

Finally, the paper is linked to the treatment effects literature on CCT programs, which aims to identify causal effects of these programs. One of the big questions this literature has investigated is whether and to what extent the targeting of conditional cash transfers to women is effective (see Yoong et al. (2012) for a systematic review). In the empirical part of this paper, we study the welfare effects of the well known PROGRESA program. Several papers show that PROGRESA has been beneficial to eligible households in a number of important dimensions. The program has been found to increase educational attainment of children (Schultz, 2004; Attanasio et al., 2012), to improve health of all household members (Gertler, 2004; Behrman and Parker, 2011), to increase investments in business and small livestock (Davis et al. (2002) and Rubalcava et al. (2009)), to increase children's clothing expenses and household food consumption, and to decrease alcohol consumption by the man (Attanasio and Lechene, 2002, 2014). Differenty from all these papers, we are the first to quantify directly to what extend a CCT program has increased short-term individual consumption and has reduced poverty rates of men and women in eligible households. Our results with respect to this application are interesting in its own right. Indeed, our estimates of the sharing rule in treatment and control households indicate that individual consumption has increased and poverty rates have decreased for both men and women, but that men have benefited the most from the policy. To the best of our knowledge this is the first paper to estimate directly how beneficial a CCT program was in terms of individual, overall, consumption level and poverty rates. We show that the DLP framework, together with our estimation strategy, can open up new venues for estimating causal effects of these policy tools on individual material well-being.

The remainder of this paper is organized as follows. In Section 2 we describe the DLP model and discuss how the multicollinearity issue arises and the implications that this has on the parameter estimates. Section 3 describes the modifications thatwe make in order to address the problem and showcases our methodology's performance in simulation. Sections 4 is devoted to our empirical application on PROGRESA data and Section 5 concludes.

2 Theoretical Framework

This section is divided into three parts. The first subsection summarizes briefly how the concept of resource shares fits into the collective model of DLP and how they arrive at the Engel curve system that allows their identification. The DLP is a very general model because it allows to recover, in principle, the resource shares of all household members. For simplicity and necessity of our empirical strategy, we focus on a simpler version of their model, with only two decision makers and where children enter as a public good for the parents. This will allow us to recover resource shares of mother and father only. Section 2.2 illustrates the issue that generates a trade-off potential in the DLP framework, that is, the reason the model has a hard time distinguishing between differences in resource shares and differences in tastes. Finally, Section 2.3 outlines the set-up of a simulation study that we conduct to investigate the features of DLP and shows how the trade-off potential that arises in their framework affects least squares estimates of the parameters. We emphasizes the large variability encountered in the parameter estimates, which arises in two forms: inefficient estimates of resource shares, that is, uncertainty of the location of the sharing rule, and distortion of heterogeneous parameter estimates, resulting in an implausible negative correlation between the parameters of shares and taste difference between spouses.

2.1 The collective model and resource shares

Let us consider the simplest case of two decision makers $j \in \{1, 2\}$ of a household. These may be mother and father, while children enter the model only in the form of public goods, which are goods both decision makers benefit from. Households differ according to a set of observable attributes, such as number of children, age, location, and other socio-economic characteristics. Members 1 and 2 may have different preferences but must jointly decide on the purchase of K types of goods with prices $\mathbf{p} = (p^1, \ldots, p^K)$. $\mathbf{z} = (z^1, \ldots, z^K)$ is the vector of quantity of goods purchased by the household, $\mathbf{x}_j = (x_j^1, \ldots, x_j^K)$ is the vector of quantity of goods consumed by member j of the household and y is household's total expenditure. The DLP framework allows for economies of scale in consumption through a linear consumption technology, which takes the form of a $K \times K$ matrix A. This permits to convert the quantity of goods $\mathbf{z}$ purchased by the household into private good equivalent $\mathbf{x}_j$[6]. The shadow price vector does not equal the market price vector $\mathbf{p}$ because sharable goods have lower shadow prices than market prices. The difference between shadow and market prices accounts for scale economies in household consumption.

Let $U_j(\mathbf{x}_j)$ be a monotonically increasing, twice continuously differentiable and strictly quasi-concave, utility function of member j over a bundle of K goods. In principle this may depend also on the utility of her partner, but for simplicity we assume that they are weakly separable over the sub-utility functions of goods. The key assumption in the literature of collective models is that, even if household members may have different pref-

[6]Formally: $\mathbf{z} = A(\mathbf{x}_1 + \mathbf{x}_2)$. A practical example commonly reported in the literature is the following. Suppose a household is composed of 2 adults. They ride their car together half of the time, in which case they share the cost of gasoline 50:50. When one of them rides alone, he or she pays alone. Then the consumption of gasoline, in private good equivalents, is 1.5 times larger than the purchased quantity of gasoline at the household level. If we assume that the consumption of gasoline does not depend on the consumption of other goods, then the k^{th} diagonal element of matrix A would read 2/3 such that: $z^K = 2/3 * (\mathbf{x}_1 + \mathbf{x}_2)$. In this example, 2/3 represents the degree of publicness of good K within the household.

erences, they make consumption decisions efficiently, that is, their joint choices maximize the following (Bergson-Samuelson) social welfare function:

$$\tilde{U}(U_1, U_2) = \mu U_1(\mathbf{x}_1) + (1-\mu)U_2(\mathbf{x}_2) \tag{1}$$

where the Pareto weight μ depends on prices, individual characteristics and household expenditure. The form of the household's utility function (1) is in contrast to what has been called the unitary model of the household, where choices stem from the maximization of a single well-behaved utility function (Becker, 1991). Under this set of assumptions, the household's program reads as follows:

$$\begin{aligned} &\max_{\mathbf{x}_1,\mathbf{x}_2,\mathbf{z}} \tilde{U}(U_1, U_2) \\ &\text{s.t. } \mathbf{z} = A(\mathbf{x}_1 + \mathbf{x}_2) \\ &\quad y = z'\mathbf{p} \end{aligned} \tag{2}$$

Note that an implication of the efficiency assumption is that the collective allocation process (2) can be equivalently represented as a two-stage process (Chiappori, 1992). First members divide up non-labor income, then each makes choices according to individual preferences. Each member's optimization problem is to maximize her utility subject to a budget constraint characterized by a shadow price vector, which is the same for all household members, and a shadow budget, which is specific to that member. The solutions to (2) yield the optimal household demand functions for each good k:

$$z^k = A^k\big(h_1^k(A'\mathbf{p}, \eta_{1,h}y) + h_2^k(A'\mathbf{p}, \eta_{1,h}y)\big) \tag{3}$$

where h_j^k are the individual demand functions, $\eta_{1,h}$ and $\eta_{2,h} = (1 - \eta_{1,h}$ are the resource shares attributed to each member j of household h, that is, the member's share of total household consumption[7]. The latter are obtained after pricing the quantity consumed by each individual $\mathbf{x}_j$ at the shadow prices $A'\mathbf{p}$.

The task of identifying the resource shares is accomplished in DLP by focusing on the consumption of *private assignable goods* for each household member. These are goods that do not have any economies of scale in consumption and are consumed exclusively by one member. The typical case is men's and women's clothing because it can be

[7]Note that, in general, the share of resources is a better measure of inequality in the household than the Pareto weight μ, because the latter suffers from an important shortcoming: It is dependent on the cardinalization of preferences in the form of $U_j(\mathbf{x}_j)$. However, it is interesting to point out that Proposition 2 of BCL shows that for any given cardinalization of preferences, an exact monotonic relationship exists between μ and $\eta_{j,h}$. Hence, by estimating $\eta_{j,h}$, we are automatically saying something about the bargaining power of the decision makers within the household. Information on the resource shares can be used for various purposes. As it will become clearer in the empirical application, in the present paper we wish to estimate $\eta_{j,h}$ in order to inform policy about the effects of a certain public intervention aimed at improving individual welfare.

assumed that women do not consume men's clothing and vice versa[8]. Furthermore, DLP make the restriction that $\eta_{j,h}$ does not depend on household expenditure y (at low level of y). The Engel curve setting does not generally allow for the testing of this assumption directly. However, in the literature there are empirical evidence supporting the identification of resource shares based on this assumption (e.g. Menon et al. (2012))[9]. Given this strategy, the household demand functions (3) simplify considerably, because the shadow price of a private assignable good is equal to its market price. By using a set of preference restrictions that are discussed below, DLP provide a model that identifies resource shares without needing the identity restriction between preference of singles and married individuals.

When the household's problem is to choose the optimal amount of private assignable goods, the demand system (3) can be represented in the form of Engel curves, one for each household member for whom we have information on the assignable consumption. These Engel curves relate budget shares of the assignable good to the budget allocated to that member for the consumption of the assignable good. In our case, we look at two individuals, and the household budget share for man and woman's clothing reads as follows:

$$\begin{aligned} W_{1,h} &= \frac{z_1}{y} = \eta_{1,h} \cdot w_1(\eta_h y) \\ W_{2,h} &= \frac{z_2}{y} = \eta_{2,h} \cdot w_2(\eta_h y) \end{aligned} \tag{4}$$

where $W_{j,h}$ is the share of total household expenditure spent by member j on her private assignable good, $\eta_{j,h}$ is the resource share attributed to that member and $w_j(\eta_{j,h}y)$ is the hypothetical share of y that member j would spend on her private good when maximizing her own utility function subject to the shadow price $A'\mathbf{p}$ and a budget $\eta_{j,h}y$. The function $w_j(\eta_{j,h}y)$ can be thought in terms of "desired budget share", which takes the shape of a (standard) Engel curve in j's resources.

In System (4), $W_{j,h}$ and y are observable, and the goal is to identify the resource shares $\eta_{j,h}$. The challenge in identifying it is that for every observable $W_{j,h}$ in the left hand side, there are two unknown functions on the right hand side: $\eta_{j,h}$ and $w_j(\eta_{j,h}y)$. This is when the preference restrictions proposed by DLP become important. The authors impose that the functions $w_j(\eta_{j,h}y)$ have similar shapes, essentially fixed curvatures, either across household sizes (number of children) or across household members. Under this structure, resource shares are identified without further restrictions on the shape of the preference function $w_j(\eta_h y)$. This is because, given that j's assignable good is not consumed by her partner, the other member's desired budget share for this good is zero. To save on

[8]Goods that are consumed by only one member are also sometimes called exclusive goods. The distinction lies in the availability of separate prices. Where the goods for men and women have the same price, we consider them the same good and call it assignable. The distinction is irrelevant here because in DLP there is no need of price variation for identification purposes.

[9]See the online Appendix of DLP for further discussions on this restriction.

notation, in what follows we refer to $\eta_h := \eta_{1,h}$, member 1's share, as the principal object of investigation.

Assume that each household member has PIGLOG utility functions at all levels of expenditure (Muellbauer, 1976). Then, the Engel curve for the private assignable good of each household member (4) becomes linear in the logarithm of own expenditure and the system takes the following form:

$$\begin{aligned} W_{1,h} &= \eta_h && \big(\delta_h + \Delta_h && + \beta_h \ln(\eta_h y_h)\big) \\ W_{2,h} &= (1-\eta_h) && \big(\delta_h && + \beta_h \ln((1-\eta_h) y_h)\big) \end{aligned} \tag{5}$$

where $\delta_h + \Delta_h$ and δ_h are respectively the intercepts for the husband's and wife's Engel curves. The parameter Δ_h captures man's taste difference in household h with respect to the woman, whereas β_h is the slope of these Engel curves in the log resources of each spouse, which are equal to each spouse's share in total household resources. The crucial difference between these parameters is that the intercepts are allowed to vary across spouses, whereas β_h is assumed to be the same. DLP show that the system is identified exactly because the parameter β_h must be the same across spouses. This is the preference restriction that the authors call "Similar Across People" (SAP). Notice that in the empirical application the set of parameters η_h, δ_h, Δ_h and β_h are going to be linear functions of characteristics.

Before concluding this section, two final remarks on the DLP framework are in order. First, as we said earlier, the authors consider also an alternative semi-parametric restriction called "Similar Across Types" (SAT), in which the slope parameter is unaffected by the number of children in the household. This is only a slight departure from the more standard idea of imposing equality across household's configurations and also turns out to have less bite than SAP as the authors themselves acknowledge. This is also the reason why we place emphasis on the SAP and conduct our empirical investigation using this semi-parametric restriction rather than the other one. Second, a further reason why the DLP has a great appeal is because it allows to include children as an additional (unique) household member and to estimate the corresponding share of resources. This is possible also in practice because data on the consumption of children's clothing and footwear is commonly available in survey data. However, in the present paper, we use a 2-person version of the model, both for ease of exposition of the trade-off issue, and to be able to obtain information from singles which is part of our proposed remedy. As it will become clearer later, children cannot be observed as singles, so it is not possible to collect further, prior (or out-of-sample), information on their behavior as singles.

2.2 Trade-off potential in DLP: resource shares versus tastes

In System (5), η_h is the husband's resource share and its recovery is the principal objective of the investigation. β_h is the only parameter fixed across the two equations due to the SAP restriction, whereas η_h and Δ_h are free to move. Clearly, given the *multiplicative feature* of the system, there are two ways in which this model can account for differences between men's and women's expenditure on the assignable goods within a household: Husband and wife can differ in their resource shares, or they can differ in how much they desire clothing. In terms of parameters, this means that variation between spouses is accounted for either by η_h or by Δ_h. Since the model has been shown to be identified, in principle it can distinguish between the two. This is most fortunate, since recovery of η_h is the main goal of the analysis. Less fortunately, while the identification result proves that recovery of η_h is possible in principle, there is no guarantee that this will be easy. In fact it is not.

To see why, consider the objective function of the least squares estimator of System (5) for identical households:

$$\begin{aligned} RSS(\eta, \delta, \Delta, \beta) = & \sum_h \left(W_{1,h} - \eta\left(\delta + \Delta + \beta \ln(\eta y_h)\right)\right)^2 + \\ & + \sum_h \left(W_{2,h} - (1-\eta)\left(\delta + \beta \ln((1-\eta) y_h)\right)\right)^2 \end{aligned} \tag{6}$$

where RSS stands for Residual Sum of Square. Figures 1a-1b show a profile of the RSS over the sub-parameter space spanned by η and Δ, that is, the function shown gives the minimal values the RSS can attain for each pair (η, Δ). These two figures crucially differ in how large the parameter β is being set, which (implicitly) indicates how much "bite" the SAP assumption has. The size of the parameter β is an empirical and practical matter of the dataset at hand. "Small" values of the slope of the Engel curves indicates that the sample of households under consideration exhibits little variation in the consumption of the private assignable good (clothing) in log of total expenditure. This is a practical concern that may be faced by the practitioner more often than not, even if the dataset used is in principle of high quality. This problem is also acknowledged by DLP.

In our error-free toy example the true values are 0.4 for η and 0 for Δ, at which point a minimum of zero is attained. As expected from the identification result, this minimum is indeed unique. Unique though it may be, the minimum is situated in a long, narrow, integral sign-shaped canyon, the floor of which barely rises as one moves away from the minimum. The slope along the valley's base reflects the degree to which the values of $W_{j,h}$ depend on log individual resources as can be seen in Figure 1b where a large value of the slope parameter β was chosen. Though the shape of the problem stemming from the trade-off potential remains, the minimum becomes more prominent. This is in line with the initial intuition about the beauty of the SAP assumption: The more important

Figure 1: RSS minima as a function of (fixed) η and Δ

(a) Small β (b) Large β

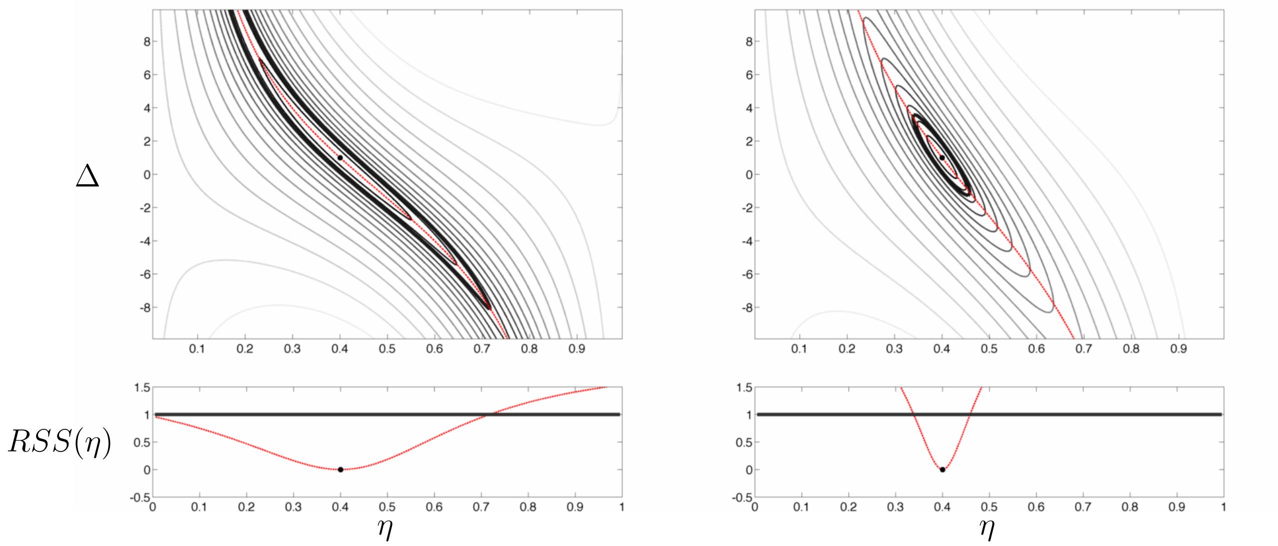

Notes: These figures provide a graphical illustration of the objective function of a least square estimator of System (5). Both figures show a profile of Equation (6) over the sub-parameter space spanned by η and Δ. The function shown gives the minimal values the RSS can attain for each pair (η, Δ). They crucially differ in how large β is being set, which (implicitly) indicates how much "bite" the SAP assumption has: The more important β becomes, the stronger is the "bite" of this very restriction and the easier becomes the recovery of the resource share η. Whereas if β is small, the resource share η and the taste parameter Δ can be traded off against one another at very little cost in terms of the objective function of the estimator.

β becomes, the stronger is the "bite" of this very restriction and the easier becomes the recovery of the resource share η. Whereas if β is small, the resource share η and the taste parameter Δ can be traded off against one another at very little cost in terms of the objective function of the estimator. It also means that as noise is added to the model, lifting and warping the RSS function, the unique minimum is prone to move swiftly up or down the canyon even as it will hardly venture away from the red line that marks its floor.

In Appendix **??** we develop a simple example to show the difference between a non-identified system and a practically (or weakly) non-identified one using the DLP framework. This feature of the model is similar to what Raue et al. (2009) have called practical "non-identifiability", a situation where the shape of the problem makes recovery of the parameters hard even if it is identified in theory. If our purpose here were only to obtain a good fit and to predict consumption of the exclusive goods, the issue described above would be marginal. But it is precisely the point of this model to recover the resource shares η_h and in this respect, we can do better.

2.3 The implications: trade-offs in a simulation study

Before describing our own estimation strategy in Section 3 and its application in Section 4, it is useful to see how the trade-off potential in DLP's model translates into features of least squares estimates. In order to do so, we first choose an empirical specification of our model and then implement a Monte Carlo study.

We augment (5) by replacing each parameter with a linear index in p demographic characteristics which are collected in a design matrix Z with rows z_h. The columns in Z are normalized to have zero mean and unit variance. In a blatant act of abuse of notation, we now understand η, δ and Δ to be vectors of parameters of length p such that $\eta_h = z_h\eta$ and so forth. Since the columns of Z are normalized, the sample mean of η_h is now determined only by the first entry in the parameter vector η, which we call η_0 This yields the following data generating process used in the simulations:

$$\begin{aligned} W_{1,h} &= z_h\eta \qquad &&\big(z_h\delta + z_h\Delta \quad &&+ \beta\ln(z_h\eta y_h)\big) + \epsilon_{h,1} \\ W_{2,h} &= (1 - z_h\eta) \quad &&\big(z_h\delta &&+ \beta\ln((1 - z_h\eta)y_h)\big) + \epsilon_{h,2} \end{aligned} \tag{7}$$

Two things should be noted in the setup of System (7). First, the slope parameter β is left as a scalar, even though in principle this could be a linear function of characteristics. This is done for ease of exposition because it will make it easier to focus on the relationship between $z_h\eta$ and $z_h\Delta$. Second, the set of variables that are allowed to influence the heterogenous parameters is the same for each of them. Again, this is for simplicity rather than necessity, but the setup is also realistic: In practice the number of variables is limited and researchers will want to allow for as much observed heterogeneity as possible. Also, any exclusion of one variable from one side but not the other is a modeling decision, adding to researcher degrees of freedom. Inclusion of the same set of variables everywhere can be argued to be an honest and agnostic choice.

The Monte Carlo study is then set up as follows. First, we generate a sample of households whose demographic characteristics match the first and second moment of the households in the dataset that we use in the empirical application. Second, we set the true value of the parameters that we wish to recover. The parameters $\eta_0, \delta_0, \Delta_0$ and β are set to estimates obtained from the full sample[10], meaning that the simulated true means of the heterogenous parameters $\eta_h, \delta_h, \Delta_h$ and β_h will be equal to those estimated because all covariates are standardized. The remaining parameters govern the effects of the covariates and are drawn from normal distributions. The only exception to this last rule is that a large effect on the resource share was set for the first covariate (a dummy variable) in order to simulate the effect of a treatment that affects only the resource share. Table 5 shows the true values set for the entries of the parameter vector η. As for the

[10] These estimates were obtained using the DLP system and a simple nonlinear least squares estimator, that is, without the adjustments proposed in this paper.

parameter β, note that in the application of Section 4 it turns out to be quite small and often non-significant. Since we use this value in our simulations, we are in the case of a "small" β as depicted in Figure 1a. Third, we generate budget shares $W_{j,h}$ according to a data generating process described by an augmented version of System (5) that allows for heterogeneous parameters in the same manner as DLP[11]. Finally, we draw 1.000 samples of the same size as the dataset for our application (2.628 observations) by generating appropriate (gamma distributed) error terms that result in a realistic distribution of observed budget shares in the real sample. Before the outline of the data generating process of private assignable goods and the simulations of the system, in what follows we explain the exact structure and rationale for the choice of the error terms.

To show that our strategy works even in a setting of non-normal errors, we generate these in a way that mimics measurement error from an infrequent purchase problem. This results in an asymmetric and heteroscedastic error structure. We start by generating "true" budget shares $W^t_{1,h} = W_{1,h} - \epsilon_{1,h}$ and $W^t_{2,h} = W_{2,h} - \epsilon_{2,h}$ according to System (7). These can be interpreted as the true consumption of clothing and footwear of our simulated households but they are not equal to the purchased quantities. To obtain households' observed budget shares $W_{1,h}$ and $W_{2,h}$ we then multiply these by a random factor $\nu_{j,h}$ that indicates how much of this consumption was replaced during the recall period. $\nu_{j,h}$ is zero with probability 0.3, which is chosen to mimic the characteristics of the real dataset, and otherwise follows a gamma distribution with shape parameter $k = 4$ and scale parameter $\theta = 5/14$. The latter choice of values results in $\mathbb{E}[\nu_{j,h}] = 1$. $\nu_{j,h}$ is i.i.d. across both households and individuals. Since this multiplicative deviation defines our observed budget shares, the additive errors $\epsilon_{j,h}$ are simply defined as the difference between observed and true budget shares:

$$\begin{aligned}\epsilon_{j,h} &= W_{j,h} - W^t_{j,h} \\ &= W_{j,h} - \nu_{h_j} W^t_{j,h}\end{aligned}$$

These additive errors have an expected value of zero by construction but are clearly not themselves identically distributed. Instead they depend on the true budget shares $W^t_{j,h}$ and are so heteroscedastic. This specific structure will be ignored by our estimation strategy in Section 4 as the true distribution is unknown in practice.

Figures 2a and 2b illustrate the results of our simulation study and summarize the two implications of the trade-off potential implicit in the DLP system. Figure 2a shows the means of estimates of the heterogenous parameters $\hat{\eta}_0 = \bar{z}\hat{\eta}_1$ and $\hat{\Delta}_0 = \bar{z}\hat{\Delta}$ from a our Monte Carlo simulation. These means are highly correlated across draws and clearly follow a similar pattern as shown in Figure 1a. In particular, the graph shows that

[11]Thanks to our setup, the theoretical budget shares were non-negative for all households, while still being similar in magnitude to the observed budget shares from our empirical application.

Figure 2: Estimates of resource shares and taste differences

(a) Means of $\hat{\eta}_h$ vs $\hat{\Delta}_h$ (b) $\hat{\eta}_h$ vs $\hat{\Delta}_h$ in one draw

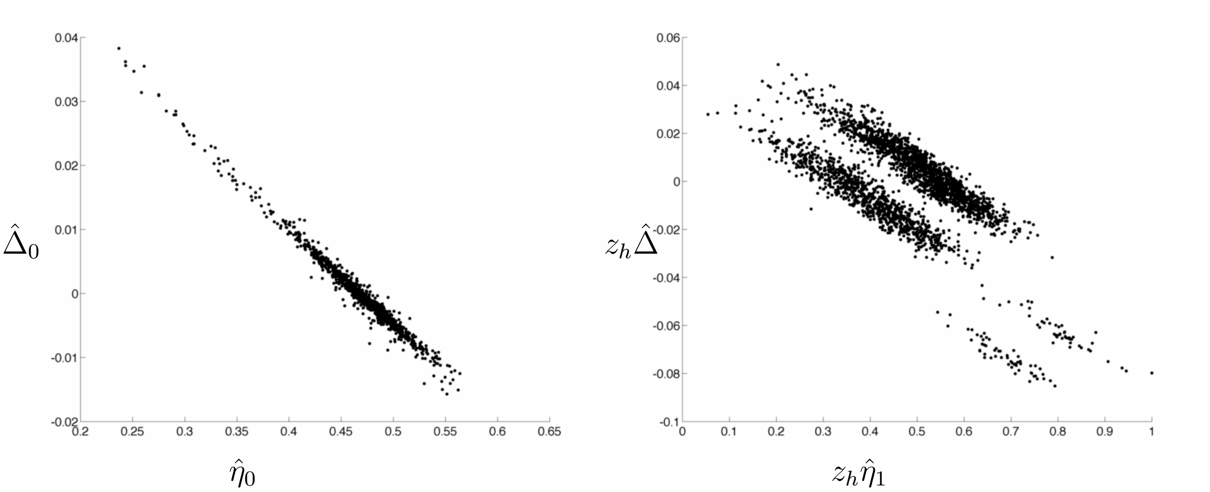

Notes: Figure 2a shows the estimates of η_0 and Δ_0 from our Monte Carlo study. Clearly the mean of the resource share $\hat{\eta}_h$ is highly uncertain and correlated with the mean of $\hat{\Delta}_h$ across draws. Figure 2b plots $z_h\hat{\eta}_1$ and $z_h\hat{\Delta}$ for all 2628 households in just one draw of the simulation. According to these estimates, individuals with large resource shares have systematically low taste difference with respect to the spouse, and vice versa.

the location of the sharing rule is very uncertain. While this was to be expected from the outline of the model in the previous subsection, a perhaps less obvious phenomenon occurs in the estimates of just one estimation round. Figure 2b plots $z_h\hat{\eta}$ and $z_h\hat{\Delta}$ for all 2628 households in the simulated dataset for just one of our 1000 draws. The picture is strikingly similar. If these estimates are to be believed, there is great observed variability among the households in the data. Yet this variability has little effect on predicted expenditures since those husbands who command only a small share of resources reliably also happen to like clothing much more than their wives and vice versa. We argue instead that this negative correlation is not to be believed. Indeed, if we generate a new dataset where we impose a positive correlation between shares and taste differences, the estimator would still yield the negative correlation observed in this graph. At the same time, the demographic characteristics that are estimated to strongly reduce shares η_h but increase Δ_h in one round of the simulation, will often be estimated to do just the opposite in the next, leading to high variability in these parameters. We believe that it is reasonable to think that true resource shares and tastes should display little correlation in a large enough data set. In consequence, we will use a restriction on this correlation to discipline the estimates of heterogenous parameters.

3 Estimation Strategy

Our aim in the present paper is to reduce the large variability obtained in the parameter estimates of the DLP model. The importance of this objective lies in the need for the practitioner to obtain tight and plausible confidence intervals around parameter estimates, in order to be able to give precise information about the location of the sharing rule. To achieve this, we follow an empirical Bayes method that incorporates a minimal amount of information of singles in a shrinkage term, and rely on the assumption that resource shares and differences in preferences between spouses are not correlated. Both restrictions imply mild preference assumptions of the underlying collective model. In this section, we first outline and motivate our strategy to achieve this objective, and then assess its effectiveness through a simulation study.

3.1 Stabilization of mean resource shares

We wish to limit the variability of $\hat{\eta}_0 = \bar{z}\hat{\eta}$, the intercept of the resource shares function, which identifies the location of the sharing rule. This large variability is caused by the trade-off potential in mean shares and tastes shown in Figure 2a. In order to achieve this objective, we incorporate what we consider a safe aspect of singles' behavior as prior (or out-of-sample) information on tastes of men and women. **We do so by computing a ratio between the budget shares that single men and single women are estimated to allocate to clothing for some fixed income and demographic characteristics.** This information is then introduced into the estimator using the empirical Bayes method of shrinkage. We therefore assume that an equivalent ratio between married men's and women's "desired budget shares" will be similar:

Restriction 1. *For some 'anchor' value y_a of individual resources and at means of the demographic variables, the ratios*

$$R_c = \frac{\delta_h + \Delta_h + \beta_h \ln(y_a)}{\delta_h + \beta_h \ln(y_a)} \text{ and } R_s = \frac{w_{1,s}(\ln(y_a))}{w_{2,s}(\ln(y_a))}$$

are similar.

We obtain the ratio R_s for singles by estimating their demands $w_{j,s}(\ln(y_a))$ for the assignable goods separately for single men and women, then evaluating these at the reference level y_a of single income. This step, which yields the scalar $\hat{R}_s$, does not in principle need to be done using linear Engel curves. We would be equally justified in estimating the single's demand for the assignable good non-parametrically as a function of log income.

To get a sense of the implications of Restriction 1, note first that we cannot simply estimate the couple's desired budget share functions on singles and substitute the results into the couple's system because there are good reasons to expect these to change when individuals form a household. The main such reason is that even though our private assignable good is not shared by assumption, the rest of the households' consumption is shared to some degree, leading to cost savings and effectively relative intra-household shadow prices that differ from relative market prices. In fact, the parameters δ, Δ and β, which determine the desired budget share functions, hide a consideration of economies of scale. This consideration is invisible in the model seen so far, but in Lewbel and Pendakur (2008) (hereafter LP), which is very much linked to the DLP model, the relationship between the desired budget share functions of members and the Engel curves of singles is clear.

Following LP, assume that the demand for the assignable good by singles of gender j is determined by the (linear, for simplicity) Engel curve:

$$w_{j,s}(\ln y_s) = a_j + b_j \ln y_s$$

Then, the demand by the same individual, for the same good, in a two-person household, can be written as:

$$\begin{aligned} w_{j,1}(\ln y_h) &= \eta_j(w_{j,s}(\ln y_h - \ln D_j + \ln \eta_j) + d_j) \\ &= \eta_j(a_j + b_j(\ln y_h - \ln D_j + \ln \eta_j) + d_j) \end{aligned}$$

where D_j is the scale economy index for member j of the couple. This parameter is an overall measure of the cost savings for j from the household's consumption technology. $D_j \in [0.5, 1]$ where 1 is no sharing and 0.5 is perfect sharing. The parameter d_j also concerns the consumption technology. It is the elasticity of D_j with respect to the price of the assignable good: $d_j = \dfrac{\partial \ln D_j}{\partial \ln p_k}$. This elasticity should be positive since an increase in the price of clothing will normally lead to an increase in expenditure on this non-shared good and so to a decrease in the portion of expenditure on sharable goods.

Our "similarity" Restriction 1 then becomes, in the terminology of LP:

$$\frac{a_1 + b_1(\ln y_a - \ln D_1 + \ln \eta) + d_1}{a_2 + b_2(\ln y_a - \ln D_2 + \ln(1-\eta)) + d_2} \approx \frac{w_{1,s}(\ln(y_a))}{w_{2,s}(\ln(y_a))}$$

Note that the word "similar" in our restriction is quite flexible. We will show in simulation that we can achieve solid improvements even when the true couples' ratio differs from that used in the shrinkage term by a large factor, meaning that we can get solid improvements even when we have relatively poor prior information.

The new objective function is the sum of the old RSS and a shrinkage term, the weight of which is regulated by a shrinkage parameter λ:

$$\min_{\eta,\delta,\Delta,\beta} RSS(\eta,\delta,\Delta,\beta) + \lambda n \left(\exp \left| \ln \left(\frac{\hat{R}_s}{R_c} \right) \right| - 1 \right)^2 \tag{8}$$

where n indicates the sample size. As λ becomes very small, this problem approaches the problem without shrinkage from Equation (6). As λ becomes very large, the ratio R_c in the estimates is assured to match $\hat{R}_s$, the quantity estimated from singles' data. Choosing a large λ is thus equivalent to imposing the restriction $R_s = R_c$. By choosing a shrinkage parameter that is in between these extremes, we seek a middle ground: We do not impose the restriction fully but do give some credence to it (Fessler and Kasy, 2015). In the process, we reap the rewards of reduced variance in our estimate of the mean resource share η, our parameter of interest. But, given that we do not strictly believe in the truth of the statement $R_s = R_c$, we also pay a cost in the form of a bias. In the following sub-section we outline our proposed strategy for the choice of an appropriate λ.

Figure 3: RSS minima as a function of (fixed) η and Δ with shrinkage

(a) $\hat{R}_s = R_c$ (b) $\hat{R}_s = 1.4R_c$

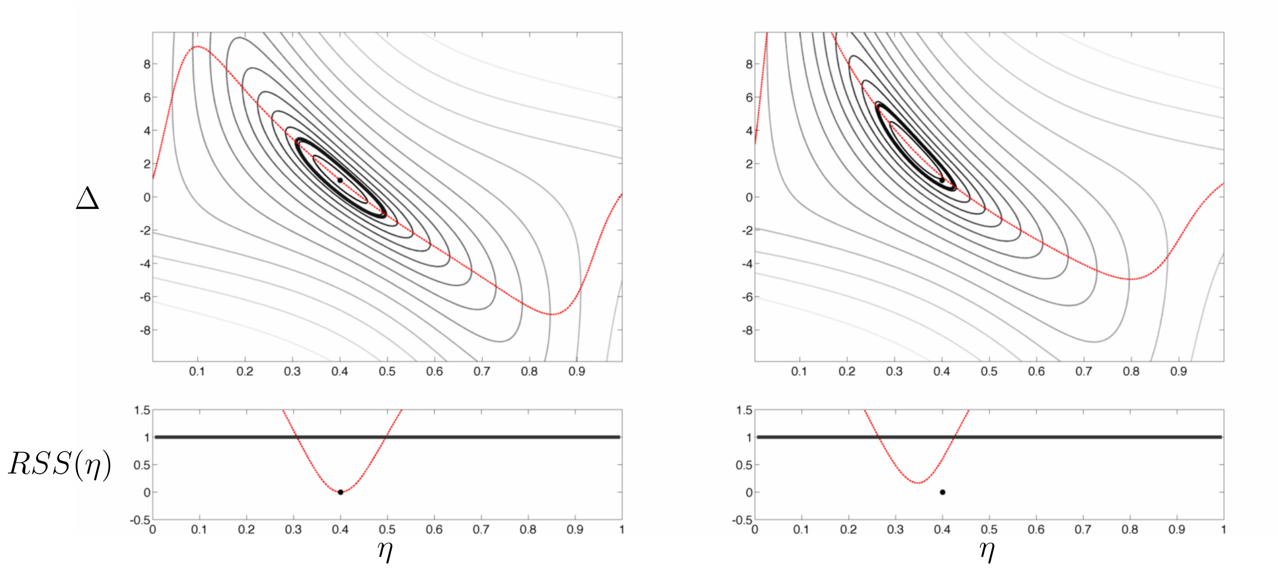

Notes: These figures provide a graphical illustration of the objective function of a modified least square estimator with shrinkage terms with different prior information. Both figures show a profile of Equation (8) over the sub-parameter space spanned by η and Δ. As before, the function shown gives the minimal values of RSS can attain for each pair (η, Δ). In both cases the shrinkage parameter is the same but the two differ in the quality of the prior information that was employed in the form of $\hat{R}_s$. While in Figure 3a, the value of $\hat{R}_s$ was simply set to the true R_c from our data generating process, this is not the case in Figure 3b, where the value was intentionally overstated by 40%.

Before concluding this sub-section, an illustration of these two sides of our shrinkage approach is given in Figures 3a-3b. In these images, the setup is the same as in Figure 1a, except that a shrinkage estimator has been used. In both cases the shrinkage parameter is the same but the two differ in the quality of the prior information that was employed in the form of $\hat{R}_s$. While in Figure 3a the value of $\hat{R}_s$ was simply set to the true R_c from our data generating process, this is not the case in Figure 3b, where the value was intentionally overstated by 40%. The effect of shrinkage is similar to that of having steeper Engel curves as shown in Figure 1a. Here, the valley floor is raised left and right of the objective function's minimum, making identification of this minimum much easier. At the same time, this minimum does not necessarily coincide with the true value any more.

3.2 Choosing the shrinkage parameter λ

The choice of the shrinkage parameter λ reflects the degree of certainty we attach to Restriction 1 in our empirical application of Section 4 and governs the bias-variance trade-off. The empirical Bayes literature offers a wide range of possible strategies that one can employ. Here, we use an agnostic approach that is easy to implement[12].

We first estimate our complete model in Program (10) with $\lambda = 0$ on the data at hand. Then, a Monte Carlo simulation of our model is performed, where we set the estimated parameters as our true values. Since the "true" R_c is now known, we can center the prior on a "wrong" $\hat{R}_s = 1.4R_c$ to study the performance of our strategy for different values of λ[13]. The results of these trials are reported in Table 1 for a grid of values of λ, which clearly shows the bias-variance trade-off. The first column can be considered a measure of the estimated bias that we induce, where small values of λ result in very little bias and large values result in a strong bias as the taste ratio among couples $\hat{R}_c$ is effectively forced to agree with our prior value. The second column shows how the mean squared error (MSE) on our parameter of interest, η_0, comes to be dominated by this bias as λ becomes large. While both of these statistics should be taken into account in the choice of the appropriate λ, we select the minimizer of the (root) mean squared error on η_0. Note that for this value of $\lambda = 10^{-7}$, the estimated $\hat{R}_c$ remains close to the truth on average and the loss (as measured by MSE) remains dominated by variance. We therefore conclude that even for a fairly "wrong" $\hat{R}_s$, we are only applying a light nudge and can expect a relatively small bias.

[12]We would like to thank Arthur Lewbel for suggesting this strategy.

[13]This choice of $\hat{R}_s$ is subjective but we judge the choice to be conservative in that we expect the true deviation of our estimate of $\hat{R}_s$ from the true $\hat{R}_c$ to be smaller than 40%

Table 1: Performance statistics from MCS for values of λ

λ	Median estimated $\hat{R}_c$	Variance share in MSE for η_0	RMSE for η_0
0	1.0607	0.9628	0.0690
10^{-8}	1.0621	0.9567	0.0678
10^{-7}	1.0840	0.9041	0.0619
10^{-6}	**1.1798**	**0.5958**	**0.0541**
10^{-5}	1.2838	0.3017	0.0577
10^{-4}	1.3099	0.2554	0.0600

In this exercise $\hat{R}_s$ was set to $1.3R_c$. 100 draws were used for each trial. The loss as measured by RMSE first decreases as the variance is reduced, then increases as the bias becomes large. The simulations are done using the covariance restricted estimator with shrinkage. The bold line corresponds to the λ chosen for the application.

3.3 Covariance Restriction

The inclusion of our shrinkage term in the objective function of the estimator is not enough to eliminate the implausible negative correlation observed in the estimates of shares and taste difference in Figure 2b. A negative correlation between resource shares η_h and taste differences Δ_h will continue to be induced in our estimates because the shrinkage term only disciplines their means, as determined by the parameters η_0 and Δ_0. The parameters that govern how covariates affect η_h and Δ_h remain free and continue to be traded off against one another. In order to address this issue, we require that $\eta_h = z_h\eta$ and $\Delta_h = z_h\Delta$ be uncorrelated:

Restriction 2. *The vectors $Z\eta$ and $Z\Delta$ are uncorrelated:*

$$\rho(Z\eta, Z\Delta) = 0 \Leftrightarrow cov(Z\eta, Z\Delta) = 0 \Leftrightarrow \langle X\eta_+, X\Delta_+\rangle = 0 \tag{9}$$

where X and η_+, Δ_+ are the constant-free equivalents of Z and η, Δ respectively.

Here X is the matrix of standardized demographic characteristics, equal to Z without the first column, while η_+ and Δ_+ are the parameter vectors η and Δ without the first entries η_0 and Δ_0, which gave the constant terms. Because $\Delta_h = \Delta_0 + X_h\Delta_+$ gives the difference in *desired budget share functions* between the household members of household h, Restriction 2 implies that this difference may not depend on the resource share. This still allows for differences in *desired budget shares* to depend on the resource share. In fact, since the desired budget share is a linear function of log individual resources $\ln \eta_{j,h} + \ln y_h$, they will. Nonetheless, Restriction 2 is a restriction on household behavior. We will show below that our estimation strategy is highly successful even when this is violated in the true data generating process. It should also be noted that we are making an assumption on a quantity about which we cannot learn anything from the unconstrained approach:

The strong negative correlation seen in Figure 2b is induced in the unconstrained OLS estimator even when the true data generating process features a positive correlation.

Finally, notice that Equation (9) can be further simplified using the covariance matrix Σ_X of the variables in X:

$$\begin{aligned} 0 &= \frac{1}{n-1}\langle X\eta_+, X\Delta_+\rangle \\ &= \frac{1}{n-1}\sum_{h=1}^{n}(x_h\eta_+ x_h\Delta_+) \\ &= \frac{1}{n-1}\sum_{t=1}^{p}\sum_{s=1}^{p}\left(\eta_{+,s}\Delta_{+,t}\sum_{h=1}^{n}x_{hs}x_{ht}\right) \\ &= \eta_+'\Sigma_X\Delta_+ \end{aligned}$$

Though this constraint is nonlinear it is relatively simple and can be used as in the following equality constrained program:

$$\begin{aligned} &\min_{\eta,\delta,\Delta,\beta} RSS(\eta,\delta,\Delta,\beta) + \lambda n\left(\exp\left|\ln\left(\frac{\hat{R}_s}{R_c}\right)\right| - 1\right)^2 \\ &\text{s.t. } \eta_+'\Sigma_X\Delta_+ = 0 \end{aligned} \tag{10}$$

which summarizes the empirical strategy proposed in the present paper to address both issues of the DLP framework.

3.4 Performance in simulations

We illustrate the effect of the combined estimator described in Program (10) by placing its results in the graphs of Figures 2a and 2b. This is done in red in Figures 4a and 4b which show the success of the shrinkage term in reducing variability of the mean share across bootstrap iterations and of the covariance constraint in eliminating the trade-off between effect sizes in resource shares η_h and tasted differences Δ_h. The random nature of all but one of the effect sizes guarantees that the true correlations between η_h and Δ_h are almost surely not equal to zero and therefore do not satisfy Restriction 2. At the same time, the correlations are unlikely to be strongly negative or positive. In the instance shown here, the true correlation between η_h and Δ_h happens to be 0.07. Nonetheless the unrestricted OLS estimator consistently displays strong negative correlations. These averaged around -0.6 and were negative in every single instance. Restriction 1 is also not satisfied: In the Monte Carlo study, we shrink toward a "singles' ratio" $\hat{R}_s$ that overstates the true ratio R_c in our data generating process by a full 40%.

Table 6 shows parameter estimates of Program (10) of the linear function η for this Monte Carlo simulation with the true parameters in the first column. Whereas Table 5

Figure 4: Estimates of resource shares and taste differences in the MC study.

(a) Means of $\hat{\eta}_h$ vs $\hat{\Delta}_h$ (b) $\hat{\eta}_h$ vs $\hat{\Delta}_h$ in one draw

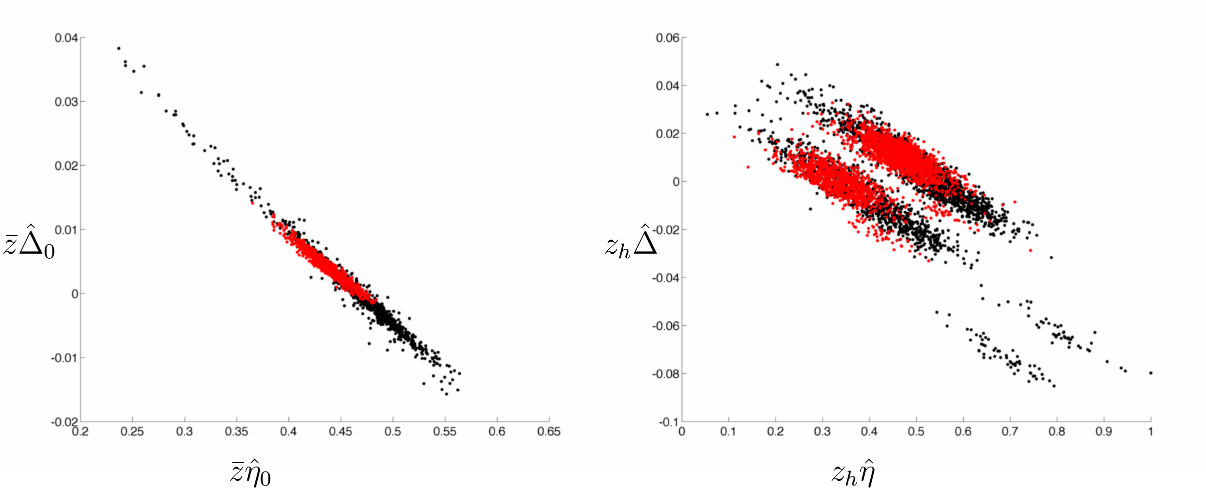

Notes: These figures provide a graphical illustration of the success of the shrinkage term in reducing variability of the mean share across Monte Carlo iterations and of the covariance constraint in eliminating the trade-off between effect sizes in resource shares η and tasted differences Δ. The meaning of these figures is explained in Figures 2a-2b.

shows results of η in case of a good prior (effectively $\hat{R}_s = R_c$). As expected, our combined strategy succeeds in reducing the standard error on our estimate of the mean resource share parameter η_0. At the same time, the bias introduced into this estimate by our shrinkage factor was by all appearances relatively mild, resulting in a marked reduction in estimated root mean squared error. The same pattern appears in the parameters η_1 to η_{13} associated with the covariates. Though the reductions in loss are not entirely constant, we do see for instance that the estimate of our artificial treatment effect η_1 is much improved by a much smaller variance.

4 Application using PROGRESA data

We apply our methodology to a sample drawn from the surveys collected to evaluate the impact of PROGRESA. This is a conditional cash transfer (CCT) program implemented in rural Mexico in the late 1990s. The dataset is suitable for a variety of reasons. First, the consumption module includes, in the six-month recall period, household expenditures on clothing and shoes for the household head, spouse and children. This is the crucial information necessary to apply the DLP model. In our empirical implementation, we use a single private assignable good for each adult member which is equal to the sum of clothing and footwear expenditures for that person. Second, the sample is large enough to have a sufficient number of single-headed households to allow us to gather a meaningful prior (or out-of-sample) information on private assignable consumption. This is needed in order to implement our strategy aimed at reducing the variability of the parameter estimates in the DLP framework. Third, the surveys have been collected to evaluate the impact of a welfare program whose objective was to fight poverty among marginalized households. By applying our strategy on this dataset, we are able to quantify the welfare effects of PROGRESA in terms of change in individual consumption and poverty of the adult members. Finally, although the surveys collected to evaluate the effect of the program are generally considered of high quality, we show that this is no guarantee that the sample of households at hand has a sufficient variability in the expenditure of the private assignable goods, which is necessary to estimate individual resource shares with high precision. Hence the choice of using this particular dataset is precisely to motivate the need to use our strategy in order to obtain tight(er) parameter estimates.

This section is divided in four parts. First, we provide some background information on the program, then we present descriptive statistics of the sample adopted in our empirical analysis, next we show the results of the parameter estimates with both least squares and our augmented estimator, and finally we conduct a welfare analysis of the policy by computing poverty rates before and after the treatment.

4.1 Program design

PROGRESA is the first conditional cash transfers (CCT) program of a new generation of welfare interventions, launched by the Mexican government in late 1990s to help poor people in marginalized rural areas. It was implemented based on a phase-in approach starting in 1997. Of 10.000 localities included in the first expansion phase, 506 localities were selected in the evaluation sample, 320 of them were randomly chosen to have an early start of the program, whereas the remaining 186 formed the control group. In practice, households in these latter villages were excluded from the program until late 1999 and became eligible for the grant only afterwards. This means that households in treatment villages, who were qualified as "eligible", started receiving cash transfers subject to the appropriate conditionalities in April 1998. Whereas "eligible" households in control villages received no payment until after November 1999.

The stated objectives were to introduce incentives to improve the accumulation of human capital of children and at the same time to alleviate short-term poverty. To achieve these objectives, the government was providing poor households with cash transfers conditional on the fulfillment of certain behaviors. The first set of conditions were related to education. Eligible households could receive a (large) portion of the grant conditional on their child school enrollment and attendance. Given that school attendance in primary school was nearly universal, whereas only about 60% of children continue to secondary education, the conditions were binding, in practice, only for households with older children. The second set of conditions were related to health seeking behavior. A further portion of the grant was conditional on woman taking their young children to health centers and attending a number of courses organized by the program.

Three aspects of the design were crucial. First of all, mothers were the primary recipient of the cash which was received bi-monthly. Woman role and involvement in the program was decided under the assumption that this would allow them to gain bargaining power in the decision making process of the household. Second, price subsidies and transfers in kind were replaced by monetary transfers which directly affected total household expenditure. Third, the amount of transfers available for each family varies with the school-level, gender and age of the child, in order to match the different opportunity cost faced by the family.

4.2 Sample selection and descriptive statistics

Throughout the observational period, extensive surveys were administered roughly every six months from August 1997 to November 2000 and the surveys collected in each village was a survey of the population. The original evaluation sample contains 24.077 households, of which 61.5% are couples with any number of children and no other adult individual living in the household, 6.5% are female single-headed households with any number of children and 4% are male single-headed households with any number of children. The remaining 28% of households are extended families with more than two adult members living in it.

Table 2: Sample means for demographics by treatment status

	Control	Treatment	Difference
Parents' characteristics			
Age of the father	32.37	32.77	-0.39
Age of the mother	28.50	29.17	-0,669**
Education of the father	4.05	4.14	-0.09
Education of the mother	4.12	4.11	0.01
Household composition			
Household size	4.22	4.24	-0.02
Total number of children	2.22	2.24	-0.02
Number of children in sec school	0.11	0.14	-0,033**
Number of children aged 0-5	1.14	1.14	0.00
Number of children aged 6-16	1.08	1.11	-0.03
Mean age of all children	5.78	5.92	-0.14
Proportion of female children	0.49	0.46	0,025*
Clothing budget shares			
Men	0.01	0.01	-0,001**
Women	0.01	0.01	-0,001*
Children	0.03	0.03	-0,006***
Observations	991	1637	

Notes: Summary statistics of the first post-treatment wave of the Mexican PROGRESA dataset (October 1998). The sample includes only natural couples from control and treatment villages.

In the present paper, we use two waves from the beginning of the first trial: A wave of surveys from October 1998 and one from May 1999. Our sample consists of 2628 households comprised of married couples with one to three children all under 15 years of age. 1673 of them reside in treatment villages, the rest belong to the control. We exclude households that were deemed non-poor (in the program sense) and therefore illegible for the grant. All are nuclear families such that the only adults present are the one man and one woman who are parents of the children present in the household. We also exclude households with no children and those with more than three to obtain a degree of homogeneity. We compute total expenditure as the sum of all non durable expenditure including food, which makes up around 70% of all expenditure. Our assignable good expenditure is the sum of expenses for clothing and footwear. These are available separately for men and women. Separate information on such expenditure for girls and boys in the household are not used in the model, though these naturally count toward total expenditure. Table 2 gives summary statistics of our assignable goods and some demographic data for our sample of rural families with 2 parents and 1-3 children.

Since the PROGRESA data are very rich, in our empirical implementation we include several demographic variables, which may affect preferences and resource shares. Among these is a dummy indicating treatment status, meaning that the household was eligible for the cash transfer. The eligibility for the PROGRESA grant, which benefits mothers and was randomly assigned, has been considered an example of a distribution factor (see Attanasio and Lechene, 2002, 2014; Bobonis, 2009).

4.3 Parameter estimates

We estimate the model on the sample described above both by SURE and by using our modified shrinkage estimator (Program (10)). We use a total of 13 control variables plus dummies for the seven Mexican states in which the households were located. Five of these are characteristics of the parents: their education level, age in years as well as a dummy for whether the husband speaks an indigenous language. Seven relate to the children in the household: two dummies for the number of children present, number of school-age children, number of children enrolled in secondary school, ratio of female children in the household, number of children 5 years old or younger, and mean age of the children. One variable of special interest is the treatment status of the household, i.e. whether the household is eligible to receive the grant. All demographic variables are allowed to affect both the allocation of resources across individuals (i.e. they enter the term η_h), and the preferences of all individuals in the household (i.e. the terms δ_h and Δ_h). This includes the dummy variable indicating eligibility for the PROGRESA grant.

In order to implement our empirical strategy, we proceed in three steps. The first step is to obtain an estimate of the ratio between men's and women's "desired budget shares" from our sample of single fathers and single mothers. Singles' demands for the assignable good can be represented by linear Engel curves:

$$w_{1,s} = z_s\delta_{1,s} + \beta_s \ln(y_s) + \epsilon_{1,s}$$
$$w_{2,s} = z_s\delta_{2,s} + \beta_s \ln(y_s) + \epsilon_{2,s}$$

where $\delta_{j,s}$ and $\beta_{j,s}$ may differ from their couples' counterparts and z_s is the vector containing the single's demographic characteristics. Unfortunately, due to the small number of single fathers and mothers in our rural Mexican sample, the slope coefficients are very imprecisely estimated and we turn instead to means of the observed budget shares: $\hat{R}_s = \frac{\bar{w}_{1,s}}{\bar{w}_{2,s}}$. This is our prior, or out-of-sample, information on singles' behavior that we introduce into the DLP framework. In our dataset, this value turns out to be 0.98, hence very close to 1.

Second, we set the shrinkage parameter λ. This parameter reflects the degree to which we place trust in our Restriction 1. We set it to what we believe to be a conservative value of 10^{-7}, for the reasons outlined in Section 3.4. The chosen value has shown itself to be large enough to yield significant improvements in the simulation and small enough to improve the estimator in a mean squared error sense even for a substantial overestimate of R_c, which is the true ratio of married men's to married women's taste for clothing and footwear at mean characteristics.

Table 3: Estimates and standard deviations for selected parameters

	LS Estimator				**Shrinkage Estimator**			
	est	*sd*	*t-stat*	*p-value*	*est*	*sd*	*t-stat*	*p-value*
$\eta_{1,0}$	0.521	0.108	4.824	0.000	0.509	0.020	26.067	0.000
$\eta_{1,treat}$	-0.020	0.036	-0.556	0.289	0.009	0.004	2.234	0.013
δ_0	-0.071	0.013	-5.486	0.000	-0.082	0.011	-7.650	0.000
δ_{treat}	-0.001	0.002	-0.586	0.279	0.001	0.001	1.318	0.094
Δ_0	-0.001	0.020	-0.041	0.484	0.001	0.003	0.402	0.344
Δ_{treat}	0.003	0.006	0.558	0.289	-0.001	0.000	-1.632	0.051
β	0.011	0.002	6.979	0.000	0.012	0.001	9.431	0.000

Notes: Estimates of the main parameters of the vectors η_1, δ, Δ and β are from least squares and from modified estimator. Standard errors are bootstrapped.

Finally, we take these values and solve the constrained minimization problem in (10). Estimates of selected parameters on the full sample are shown in Table 3, whereas the full set of coefficients is given in the appendix. In the differences between the least squares estimator and our shrinkage estimator with covariance constraint, the effects of our two modifications can be seen quite clearly. The shrinkage term yields a strong reduction in variance especially for the parameters η_0 and Δ_0. As was already the case in the simulation, the estimate of β changes very little. This is as it should be, since this parameter was not affected by the trade-off problem we seek to redress. Due to the covariance restriction from Restriction 2, the effect sizes of any covariate on η and Δ can no longer be traded off against one another. This has lead to a substantial reduction in the variability of the associated estimates. Only the effect sizes related to treatment are shown in Table 3. For both η_{treat} and Δ_{treat} the precision of the estimates increases greatly.

4.4 Handling endogeneity

It should be noted that in the results we present here, we do not apply any instrumental variables techniques to address a potential endogeneity of total expenditures in this model. Given the non-linearity of the setup as well as the paper's focus on the trade-off issue and a reduction in estimator variance, an application of the usual techniques is not trivial. This is exacerbated by a lack of strong instruments[14]. A control function caused additional instability in the system and a substitution approach considerably lowered the variation in total expenditures available for estimation. This is especially problematic given that the model is entirely based upon Engel curve estimation. We also argue that any reverse causality problem in which tastes for the assignable goods determine income is unlikely to be at the top of the list of problems in resource share inference, especially in light of the considerations outlined in this paper.

4.5 Welfare effects of PROGRESA

The difference between the two estimation strategies, particularly in the parameter η_{treat}, has consequences on our best guess of the effect of the treatment on separate poverty rates for men and women. Because we now see married men's resource shares rise slightly with treatment, their welfare is improved to a greater degree than that of their wives. This can be seen in Table 4, which displays poverty rates. The poverty rates displayed were calculated using the OECD-modified equivalence scale and the World Bank's 2 USD per day threshold. For the overall household poverty rates this means that an index is first computed based on the household's composition that gives the number of single-adult equivalents. The parents count as 1.5 rather than two because the scale assumes a degree of sharing in consumption of public goods. Each child is assigned a value of 0.3, indicating further economies of scale from sharing as well as reduced needs (Oecd, 2013). The poverty threshold is then applied to this index. The fact that these rates are high overall comes as no surprise since the sample only contained households that were deemed poor for the purposes of the cash transfer program.

To compute comparable separate poverty rates for married men and women using our resource shares we have to engage in some mild trickery. This is firstly because the model does not take children into account as separate members of the household but rather treats their welfare as public goods which both parents consume and secondly because our model considers only assignable (non shared) goods and so is mute about the overall degree of sharing. In order to overcome this we consider the adults to jointly command a share of household resources equal to $\frac{1.5}{1.5+0.3nkids}$ where $nkids$ is the number of children in the household. In so doing we implicitly assume that the parents benefit from the children's welfare equally such that the resource share η remains valid when only considering adult resources. Before computing the individual adults' resources using the resource share we then translate the adult resources into private good equivalents by multiplying by $\frac{2}{1.5}$. This is the degree of economies of scale from sharing that is implied

[14]Following Attanasio and Lechene (2014), we used average hourly wages measured at the village level for trials. However, even if we use the same PROGRESA dataset, it turns out that the same instruments that they use are not strong in our sample.

Table 4: Household and individual poverty rates according to point estimates

	LS Estimator			**Shrinkage Estimator**		
Type	*Control*	*Treatment*	*Difference*	*Control*	*Treatment*	*Difference*
1 child						
Household (OECD)	0.672	0.584	-0.088	0.672	0.584	-0.088
Man	0.672	0.640	-0.032	0.678	0.535	-0.143
Woman	0.623	0.505	-0.118	0.678	0.657	-0.021
2 children						
Household (OECD)	0.693	0.633	-0.060	0.693	0.633	-0.060
Man	0.624	0.586	-0.038	0.718	0.622	-0.095
Woman	0.745	0.684	-0.061	0.673	0.657	-0.016
3 children						
Household (OECD)	0.748	0.688	-0.060	0.748	0.688	-0.060
Man	0.681	0.615	-0.066	0.765	0.645	-0.120
Woman	0.800	0.769	-0.031	0.752	0.728	-0.025
Total						
Household (OECD)	0.711	0.648	-0.064	0.711	0.648	-0.064
Man	0.656	0.608	-0.047	0.730	0.616	-0.114
Woman	0.745	0.687	-0.057	0.706	0.687	-0.019

Notes: The numbers correspond to the percentage of households considered poor by the World Bank which has set a poverty threshold at 2$ per day. The OECD index corresponds to the standard index used. The rest is the poverty rate computed using our estimated resource shares.

by the OECD equivalence scale.

5 Conclusion

Recent advances in the literature of collective models of the household have allowed practitioners to estimate an useful measure of individual welfare known as resource share. This is defined as each member's share of total household consumption, and can be estimated in a collective model framework directly from household level data. Among the (few) collective models that allow to identify resource shares, the one proposed by Dunbar et al. (2013) (DLP) has several attractive features that make it more likely to be applied by practitioners in the near future. Indeed, the authors are able to combine a more general theoretical structure of the household with a lower data requirement and estimation complexity.

In the present paper, we take the DLP framework and investigate the sensitivity of their empirical strategy in recovering the parameters of the resource shares. Our contribution is threefold. First, we show that the DLP model suffers from highly variable estimates in some specific situations. This happens because a potential for trade-offs exists in the estimator that makes it hard to distinguish the roles of resource shares and differences in preferences. This affects the precision of both the mean level of resource shares and of effect sizes of covariates on these resource shares. Since the model achieves identification, by assuming that the preferences of persons within a household are similar (in certain limited ways), the degree of this trade-off potential depends on the importance of the restricted part of preferences.

Second, we show that the trade-off potential can be mitigated, and the variance of the estimator greatly reduced, by making two modifications which address the problem on two levels. First, we bring information on singles' preferences back into the picture but focus on what we believe to be a safe aspect of their behavior. We compute a ratio of single men's and single women's preferred budget share of clothing and footwear for the average singles in the data and call it $\hat{R}_s$. We then include a shrinkage term in the couples' objective function which penalizes large deviations in the couples' estimate from this ratio at a midpoint in the household income distribution. In this way, prior information on the ratio between men's and women's tastes (but not on their levels) is used to nudge the estimator toward what we believe to be a more reasonable point. Second, we impose a zero-correlation restriction such that resource shares and spousal differences in tastes cannot be correlated in the estimates. This makes the variability-inducing trade-off between groups of households impossible. The modified estimator is shown to perform better in the mean squared error sense on simulated data, despite the fact that our true data generating process both violated the zero-correlation assumption and featured a true ratio of men's and women's preferences at the mean that was considerably different from that of singles.

Third, we apply the methodology to conduct an individual welfare analysis of the effects of the PROGRESA program in Mexico, the first such analysis in the context of a cash transfer program. The PROGRESA program was implemented to reduce poverty among marginalized rural households. We obtain much reduced standard deviations for virtually all parameters and a reduction in magnitude for the effect sizes of the treatment, which for the simple least square estimator were often very large. In contrast to least square estimates, in which the **treatment effect on men's resource share was negative but very insignificant**, we obtain a positive effect that varies little in the bootstrap. The results of our welfare analysis on the effects of PROGRESA indicate that men benefited more from the treatment than did women as measured by individual poverty rates.

Beyond the efficacy of the methodology itself, our work shows that empirical collective models can be used to obtain precise estimates of resource shares without hard restrictions that would render the exercise trivial. This approach may help to make the estimation of collective models more useful in practice and could be extended to versions which draw on different or richer available data or make alternative identifying assumptions.

References

Apps, P. and Rees, R. (1988). Taxation and the Household. *Journal of Public Economics*, 35.

Attanasio, O. and Lechene, V. (2002). Tests of Income Pooling in Household Decisions. *Review of Economic Dynamics*, 5(4):720–748.

Attanasio, O., Meghir, C., and Santiago, A. (2012). Education choices in Mexico: using a structural model and a randomized experiment to evaluate Progresa. *The Review of Economic Studies*, 79(1):37–66.

Attanasio, O. P. and Lechene, V. (2014). Efficient responses to targeted cash transfers. *Journal of Political Economy*, 122(1):178–222.

Bargain, O. and Donni, O. (2012). Expenditure on children: A Rothbarth-type method consistent with scale economies and parents' bargaining. *European Economic Review*, 56(4):792–813.

Becker, G. S. (1991). *A Treatise on the Family*. Harvard university press.

Behrman, J. and Parker, S. (2011). The impact of the progresa/oportunidades conditional cash transfer program on health and related outcomes for the aging in mexico.

Blundell, R., Chiappori, P., and Meghir, C. (2005). Collective labor supply with children. *Journal of Political Economy*, 113(6):1277–1306.

Bobonis, G. J. (2009). Is the allocation of resources within the household efficient? new evidence from a randomized experiment. *Journal of political Economy*, 117(3):453–503.

Browning, M., Bourguignon, F., Chiappori, P.-A., and Lechene, V. (1994). Income and outcomes: A structural model of intrahousehold allocation. *Journal of Political ...*, 102(6):1067–1096.

Browning, M. and Chiappori, P. (1998). Efficient intra-household allocations: A general characterization and empirical tests. *Econometrica*, 66(6):1241–1278.

Browning, M., Chiappori, P.-a., and Lewbel, A. (2013). Estimating Consumption Economies of Scale, Adult Equivalence Scales, and Household Bargaining Power. *The Review of Economic Studies*, 80(4):1267–1303.

Browning, M., Chiappori, P.-A., and Weiss, Y. (2014). *Economics of the Family*. Cambridge University Press.

Bütikofer, A., Lewbel, A., and Seitz, S. (2011). Health and retirement effects in a collective consumption model of older households. Technical report, Boston College Department of Economics.

Calvi, R. (2015). Why are older women missing in india? the age profile of bargaining power and poverty.

Cherchye, L., De Rock, B., Lewbel, A., and Vermeulen, F. (2015). Sharing rule identification for general collective consumption models. *Econometrica*.

Cherchye, L., De Rock, B., and Vermeulen, F. (2011). The Revealed Preference Approach to Collective Consumption Behaviour: Testing and Sharing Rule Recovery. *Review of Economic Studies*, pages 176–198.

Cherchye, L., De Rock, B., and Vermeulen, F. (2012). Economic well-being and poverty among the elderly: An analysis based on a collective consumption model. *European Economic Review*, 56(6):985–1000.

Chiappori, P. (1988). Rational household labor supply. *Econometrica: Journal of the Econometric Society*, 56(1):63–90.

Chiappori, P. (1992). Collective labor supply and welfare. *Journal of Political Economy*, 100(3):437–467.

Chiappori, P. a. and Ekeland, I. (2006). The micro economics of group behavior: General characterization. *Journal of Economic Theory*, 130(1):1–26.

Davis, B., Handa, S., Ruiz, M., Stampini, M., and Winters, P. (2002). Conditionality and the impact of program design on household welfare: Comparing two diverse cash transfer programs in rural mexico. *ESA/FAO WP07-2002*.

Del Negro, M. and Schorfheide, F. (2004). Priors from general equilibrium models for vars*. *International Economic Review*, 45(2):643–673.

Del Negro, M., Schorfheide, F., Smets, F., and Wouters, R. (2007). On the fit of new keynesian models. *Journal of Business & Economic Statistics*, 25(2):123–143.

Donni, O. and Chiappori, P. (2011). Nonunitary Models of Household Behavior: A Survey of the Literature. *Household economic behaviors*.

Dunbar, G. R., Lewbel, A., and Pendakur, K. (2013). Children's Resources in Collective Households: Identification, Estimation, and an Application to Child Poverty in Malawi. *American Economic Review*, 103(1):438–471.

Efron, B. (2010). *Large-scale inference: empirical Bayes methods for estimation, testing, and prediction*, volume 1. Cambridge University Press.

Efron, B. and Morris, C. (1972). Limiting the risk of bayes and empirical bayes estimators—part ii: The empirical bayes case. *Journal of the American Statistical Association*, 67(337):130–139.

Fessler, P. and Kasy, M. (2015). How to use economic theory to improve estimators, with an application to labor demand and wage inequality in europe. Technical report.

Gertler, P. (2004). Do conditional cash transfers improve child health? evidence from progresa's control randomized experiment. *American economic review*, pages 336–341.

James, W. and Stein, C. (1961). Estimation with quadratic loss. In *Proceedings of the fourth Berkeley symposium on mathematical statistics and probability*, volume 1, pages 361–379.

Lewbel, A. and Pendakur, K. (2008). Estimation of collective household models with Engel curves. *Journal of Econometrics*, 147(2):350–358.

Lise, J. and Seitz, S. (2011). Consumption Inequality and Intra-household Allocations. *The Review of Economic Studies*, 78(1):328–355.

Menon, M., Pendakur, K., and Perali, F. (2012). On the expenditure-dependence of children's resource shares. *Economics Letters*, 117(3):739–742.

Morris, C. N. (1983). Parametric empirical bayes inference: theory and applications. *Journal of the American Statistical Association*, 78(381):47–55.

Muellbauer, J. (1976). Community preferences and the representative consumer. *Econometrica: Journal of the Econometric Society*, pages 979–999.

Oecd (2013). OECD Framework for Statistics on the Distribution of Household Income, Consumption and Wealth. page 232.

Raue, A., Kreutz, C., Maiwald, T., Bachmann, J., Schilling, M., Klingmüller, U., and Timmer, J. (2009). Structural and practical identifiability analysis of partially observed dynamical models by exploiting the profile likelihood. *Bioinformatics*, 25:1923–1929.

Robbins, H. (1956). An empirical bayes approach to statistics. In *Proceedings of the Third Berkeley Symposium on Mathematical Statistics and Probability, Volume 1: Contributions to the Theory of Statistics*, pages 157–163, Berkeley, Calif. University of California Press.

Rubalcava, L., Teruel, G., and Thomas, D. (2009). Investments, time preferences and public transfers paid to women. *Economic Development and cultural change*, 57(3):507.

Schultz, T. P. (2004). School subsidies for the poor: evaluating the mexican progresa poverty program. *Journal of Development Economics*, 74(1):199–250.

Yoong, J., Rabinovich, L., and Diepeveen, S. (2012). The impact of economic resource transfers to women versus men.

6 Tables

Simulations

Table 5: True values and MC means for the parameter vector η_1. Standard errors in parentheses; $\hat{R}_s$ set to true R_c

Parameter	True Value	**LS Estimator**		**Shrinkage Estimator**	
η_0	0.4652	0.4782	(0.0328)	0.4640	(0.0060)
η_1	-0.0150	-0.0039	(0.0290)	-0.0048	(0.0052)
η_2	0.0060	0.0008	(0.0219)	0.0011	(0.0045)
η_3	0.0019	0.0022	(0.0455)	0.0023	(0.0097)
η_4	-0.0040	0.0019	(0.0294)	-0.0075	(0.0053)
η_5	-0.0053	0.0040	(0.0486)	-0.0027	(0.0118)
η_6	0.0005	0.0032	(0.0312)	-0.0012	(0.0069)
η_7	0.0011	-0.0012	(0.0243)	0.0019	(0.0056)
η_8	-0.0026	-0.0017	(0.0316)	0.0036	(0.0070)
η_9	-0.0078	-0.0036	(0.0383)	0.0020	(0.0136)
η_{10}	0.0062	0.0004	(0.0258)	0.0016	(0.0070)
η_{11}	0.0128	-0.0003	(0.0355)	0.0067	(0.0143)
η_{12}	0.0150	0.0031	(0.0239)	0.0102	(0.0070)

Notes: Estimates of the parameters vector η_1 are from least squares and from modified estimator. Standard errors in parenthesis.

Table 6: True values and MC means for the parameter vector η_1. Standard errors in parentheses; $\hat{R}_s$ set 40% above true R_c

Parameter	True Value	**LS Estimator**		**Shrinkage Estimator**	
η_0	0.4652	0.4861	(0.0276)	0.4564	(0.0051)
η_1	-0.0105	-0.0094	(0.0293)	-0.0027	(0.0056)
η_2	-0.0006	-0.0009	(0.0218)	-0.0014	(0.0048)
η_3	-0.0007	0.0056	(0.0390)	0.0022	(0.0090)
η_4	-0.0052	0.0037	(0.0250)	-0.0033	(0.0052)
η_5	-0.0063	-0.0026	(0.0465)	-0.0016	(0.0113)
η_6	0.0005	0.0071	(0.0289)	-0.0022	(0.0071)
η_7	0.0031	-0.0056	(0.0237)	-0.0001	(0.0059)
η_8	-0.0054	-0.0012	(0.0323)	-0.0035	(0.0078)
η_9	0.0079	0.0029	(0.0355)	0.0060	(0.0118)
η_{10}	0.0078	-0.0021	(0.0251)	0.0039	(0.0072)
η_{11}	0.0092	0.0089	(0.0333)	0.0079	(0.0118)
η_{12}	-0.0120	-0.0041	(0.0261)	-0.0044	(0.0077)

Notes: Estimates of the parameters vector η_1 are from least squares and from modified estimator. Standard errors in parenthesis.

Application

Table 7: Individual resource shares in PROGRESA sample according to point estimates

	LS Estimator							
	Control				Treatment			
Type	mean	sd	min	max	mean	sd	min	max
1 child								
man	0.48	0.13	0.21	0.83	0.47	0.14	0.17	0.87
woman	0.52	0.13	0.17	0.79	0.53	0.14	0.13	0.83
2 children								
man	0.53	0.12	0.29	1.00	0.52	0.13	0.21	0.95
woman	0.47	0.12	-0.00	0.71	0.48	0.13	0.05	0.79
3 children								
man	0.55	0.13	0.30	0.91	0.53	0.14	0.22	0.94
woman	0.45	0.13	0.09	0.70	0.47	0.14	0.06	0.78
Total								
man	0.53	0.13	0.21	1.00	0.52	0.14	0.17	0.95
woman	0.47	0.13	-0.00	0.79	0.48	0.14	0.05	0.83
	Shrinkage Estimator							
	Control				Treatment			
Type	mean	sd	min	max	mean	sd	min	max
1 child								
man	0.51	0.02	0.46	0.55	0.52	0.02	0.46	0.58
woman	0.49	0.02	0.45	0.54	0.48	0.02	0.42	0.54
2 children								
man	0.49	0.02	0.43	0.54	0.51	0.02	0.45	0.56
woman	0.51	0.02	0.46	0.57	0.49	0.02	0.44	0.55
3 children								
man	0.50	0.02	0.41	0.56	0.52	0.02	0.44	0.58
woman	0.50	0.02	0.44	0.59	0.48	0.02	0.42	0.56
Total								
man	0.50	0.02	0.41	0.56	0.52	0.02	0.44	0.58
woman	0.50	0.02	0.44	0.59	0.48	0.02	0.42	0.56

Notes: Resource shares of the adults are estimated with least squares. Treatment and Control refer to households with and without the cash of the grant.

Table 8: Estimated Parameters of the Resource Shares: η_1

	LS Estimator				**Shrinkage Estimator**			
η_1 is a linear function of:	*est*	*sd*	*t-stat*	*p-value*	*est*	*sd*	*t-stat*	*p-value*
Constant	0.521	0.108	4.824	0.000	0.509	0.020	26.067	0.000
Treatment	-0.020	0.036	-0.556	0.289	0.009	0.004	2.234	0.013
No. kids in secondary	0.034	0.031	1.097	0.136	-0.011	0.006	-1.844	0.033
Two kids	0.017	0.032	0.542	0.294	-0.005	0.004	-1.151	0.125
Three kids	0.015	0.016	0.919	0.179	0.001	0.003	0.413	0.340
Indigenous language	0.015	0.022	0.684	0.247	0.001	0.005	0.188	0.425
Kids' mean age	-0.007	0.034	-0.213	0.416	0.000	0.008	-0.014	0.494
No. of young kids	0.023	0.026	0.892	0.186	0.000	0.004	-0.059	0.476
No. school age kids	0.004	0.020	0.191	0.424	-0.002	0.004	-0.409	0.341
No. of girls	-0.007	0.021	-0.323	0.373	-0.005	0.004	-1.259	0.104
Age man	0.011	0.039	0.274	0.392	-0.009	0.008	-1.128	0.130
Education man	-0.002	0.019	-0.097	0.462	0.003	0.005	0.575	0.283
Age woman	0.025	0.028	0.865	0.194	0.009	0.008	1.223	0.111
Education woman	-0.008	0.022	-0.338	0.368	-0.006	0.005	-1.232	0.109

Notes: Estimates of the parameters vector η are from least squares and from modified estimator. Standard errors are bootstrapped.

Table 9: Estimated Parameters of the Resource Shares: δ

	LS Estimator				**Shrinkage Estimator**			
δ is a linear function of:	*est*	*sd*	*t-stat*	*p-value*	*est*	*sd*	*t-stat*	*p-value*
Constant	-0.071	0.013	-5.486	0.000	-0.082	0.011	-7.650	0.000
Treatment	-0.001	0.002	-0.586	0.279	0.001	0.001	1.318	0.094
No. kids in secondary	0.006	0.003	2.022	0.022	0.001	0.001	1.322	0.093
Two kids	0.000	0.002	-0.016	0.493	-0.002	0.001	-2.328	0.010
Three kids	-0.001	0.001	-0.954	0.170	-0.003	0.001	-2.806	0.003
Indigenous language	0.002	0.002	0.910	0.181	0.001	0.001	1.248	0.106
Kids' mean age	0.001	0.003	0.336	0.368	0.002	0.001	1.489	0.068
No. of young kids	0.001	0.002	0.664	0.253	0.000	0.001	-0.331	0.371
No. school age kids	-0.003	0.002	-1.558	0.060	-0.003	0.001	-2.521	0.006
No. of girls	0.001	0.001	0.624	0.266	0.001	0.001	1.982	0.024
Age man	0.002	0.003	0.732	0.232	0.000	0.001	0.456	0.324
Education man	0.000	0.001	0.161	0.436	0.001	0.001	0.910	0.181
Age woman	0.000	0.002	-0.135	0.446	-0.002	0.001	-1.946	0.026
Education woman	-0.001	0.001	-0.377	0.353	0.000	0.001	-0.451	0.326

Notes: Estimates of the parameters vector δ are from least squares and from modified estimator. Standard errors are bootstrapped.

Table 10: Estimated Parameters of the Resource Shares: Δ

	LS Estimator				**Shrinkage Estimator**			
Δ is a linear function of:	*est*	*sd*	*t-stat*	*p-value*	*est*	*sd*	*t-stat*	*p-value*
Constant	-0.001	0.020	-0.041	0.484	0.001	0.003	0.402	0.344
Treatment	0.003	0.006	0.558	0.289	-0.001	0.000	-1.632	0.051
No. kids in secondary	-0.006	0.006	-1.025	0.153	0.001	0.001	1.809	0.035
Two kids	-0.003	0.006	-0.599	0.275	-0.001	0.001	-1.202	0.115
Three kids	-0.002	0.003	-0.659	0.255	-0.002	0.001	-2.103	0.018
Indigenous language	-0.002	0.004	-0.522	0.301	0.000	0.001	0.512	0.304
Kids' mean age	0.000	0.006	-0.031	0.488	-0.001	0.001	-0.953	0.170
No. of young kids	-0.003	0.004	-0.688	0.246	0.003	0.001	3.385	0.000
No. school age kids	-0.001	0.004	-0.316	0.376	0.001	0.001	1.285	0.099
No. of girls	0.000	0.003	0.094	0.463	0.000	0.000	0.031	0.487
Age man	-0.003	0.007	-0.491	0.312	-0.001	0.001	-0.549	0.291
Education man	0.001	0.003	0.376	0.353	0.000	0.001	0.747	0.228
Age woman	-0.001	0.005	-0.270	0.393	0.001	0.001	1.320	0.093
Education woman	-0.001	0.004	-0.241	0.405	-0.001	0.001	-1.974	0.024

Notes: Estimates of the parameters vector Δ are from least squares and from modified estimator. Standard errors are bootstrapped.

Table 11: Estimates and standard deviations for parameters of β_h

	LS Estimator				**Shrinkage Estimator**			
β is a linear function of:	*est*	*sd*	*t-stat*	*p-value*	*est*	*sd*	*t-stat*	*p-value*
Constant	0.011	0.002	6.979	0.000	0.012	0.001	9.431	0.000

Notes: Estimates of the parameters vector η are from least squares and from modified estimator. Standard errors are bootstrapped.