Inefficient Collective Households: Cooperation and Consumption

Arthur Lewbel and Krishna Pendakur*

Boston College and Simon Fraser University

February 2020

Abstract

We propose a model of inefficiency in collective households. Inefficiency depends on a "cooperation factor," which can also affect both the allocation of resources within a household and the utility of household members. Households are conditionally efficient, conditioning on the value of the cooperation factor. This lets us exploit convenient modeling features of efficient households (like not needing to specify the bargaining process), while still accounting for, and measuring the dollar cost of, inefficiency. An example of a cooperation factor is domestic violence, which we find, in Bangladeshi data, reduces consumption efficiency by 5%, and shifts 1.5% of household resources towards men.

JEL codes: D13, D11, D12, C31, I32. Keywords: Collective Household Model, Inefficiency, Bargaining Power, Sharing Rule, Demand Systems, Engel Curve

*Corresponding address: Arthur Lewbel, Dept of Economics, Maloney Hall, Boston College, 140 Commonwealth Ave., Chestnut Hill, MA, 02467, USA. (617)-552-3678, lewbel@bc.edu, https://www2.bc.edu/arthur-lewbel/. Pendakur acknowledges the financial support of the Social Sciences and Humanities Research Council of Canada. We are also grateful for the valuable research assistance provided by Shirleen Manzur, and helpful comment and feedback from our colleagues: Dave Freeman, Valerie Lechene, and Chris Muris.

1 Introduction

Collective household models of consumption often assume that the allocation and use of household resources is Pareto efficient. As observed by Chiappori (1988, 1992) and many later authors, the efficiency assumption greatly simplifies construction and estimation of such models. In particular, efficiency allows models to be estimated without specifying and solving for the specific bargaining model that is used by household members to allocate resources. Efficiency also means that households satisfy decentralization rules analogous to the first and second welfare theorems, in which the consumption behavior of the household as a whole is equivalent to each household member maximizing their own utility function, subject to a shadow budget constraint. The shadow prices in this constraint embody scale economies associated with the sharing and joint consumption of goods, while the shadow budget incorporates the allocation of resources to each member. This decentralization leads to many modeling simplifications.

However, a common objection to the use of these efficient household models in the development literature is that very prominent examples exist of inefficient household behavior. An example is household members concealing money from each other, even to the point of paying outside money holders or using low- (or negative) return savings instruments (e.g. Schaner 2015, 2017). Another example is actual or threatened domestic violence, which is widespread in some cultures and countries (e.g., Bloch and Rao 2002, Koç and Erkin 2011, Ramos 2016, Hughes, et. al. 2015, and Hidrobo, et. al. 2016).

We propose a collective household model that allows for the presence of some types of inefficiencies, but still maintains all the modeling properties and simplifications, such as decentralization theorems, that are associated with efficient household models. This model allows us to identify the resource share of each household member, defined as the fraction of the overall household budget consumed by that member (see Dunbar, Lewbel, and Pendakur 2013, hereafter denoted DLP), despite the presence of inefficiency. As DLP show, these resource shares may be used to construct measures of within-household inequality and person-level poverty measures. In addition, we also identify a dollar measure of the costs to the household attributable to inefficient use of resources.

How can models that assume efficient allocations be applied to inefficient households? The intuition for our result is to consider two different perfectly competitive economies, one of which has access to superior production technology. Each economy can be *conditionally* Pareto efficient, conditioning on the technology they have access to, even though the one with inferior technology is *unconditionally* inefficient relative to the superior economy. This conditional efficiency (conditioning on the available technology) is sufficient to obtain the decentralization simplifications associated with efficient collective households.[1]

We start with the collective household model of Browning, Chiappori, and Lewbel (2013, hereafter denoted BCL), which includes what they call a "consumption technology function" that summarizes a household's ability to share and jointly consume goods, or more generally to cooperate and thereby attain economies of scale in consumption. A household that has an inferior consumption technology relative to another is a household that has lower economies of scale to consumption, and as a result cannot attain as high a level of utility from goods for each of its members as could a household with a superior consumption technology. Nevertheless, households with each technology efficiently uses their consumption technology, and so models of efficient household behavior can be applied to each.

We first derive this conditional efficiency result in the context of the BCL model. We then extend this model to allow for unconditional inefficiency, where a given household has access to the superior consumption technology but could still choose an inefficient level of sharing. An example could be where a husband shirks responsibility for household activities to increase his own utility, even if that results in inefficiency in the household's consumption of goods.

We define the notion of a *cooperation factor* which, like a distribution factor (see Browning and Chiappori 1998), affects how resources are divided amongst household members and does not affect each member's indifference curves over goods. But unlike distribution factors, cooperation factors may also directly affect the household's consumption technology, and may affect the utility levels of individual household members.

Examples of cooperation factors could be direct indicators cooperation (e.g., measures of

[1] Models that are efficient within periods but not necessarily dynamically, like Abraham and Laczo (2017), are similar in spirit.

time spent together on household chores), or more generally be behaviors that correlate with cooperation or failures to cooperate, such as money hiding or domestic abuse. Still other examples could variables relating to private ownership of durables, or almost any variable that in the previous literature was considered to be distribution factor.

Most models in the collective household literature assume all goods are either purely private or purely public within the household (i.e., are either not shared at all, or are completely shared). Such models cannot capture our notion of efficiency, or the concept of a cooperation factor, because the definition of goods as purely private or purely public rules out any variation in how much goods are shared for a given good. This is why we start from the BCL model; it is general enough to allow for variation in sharing both between and within goods, and hence it allows for variation in consumption efficiency across households.

The BCL model is a very general collective household model, but it correspondingly has very demanding data requirements for estimation. DLP propose a restricted version of the BCL model that has far lower data requirements and is much simpler to estimate. In the present paper we first generalize BCL to allow for inefficiency in consumption, and then we add assumptions similar to those of DLP to obtain a practical empirical model that can be readily estimated with generally available household-level consumer expenditure data.

In our model, the cooperation factor affects efficiency and resource shares (which determine shadow prices and shadow budgets, respectively), and it directly affects member's utility functions. We show that variables that only influence these direct effects can serve as instruments for the endogenously chosen cooperation factor in the household's consumption demand equations. We therefore do not require a randomized instrument for identification. Instead, we only need a variable that correlates with the direct utility associated with the cooperation factor, and is conditionally uncorrelated with a consumption allocation decision.

We apply our model to data from Bangladesh. Like DLP, we use the model to construct separate measures of men's, women's, and children's resource shares to evaluate the within-household distribution of consumption. Unlike previous applications, our model allows for possible inefficiencies in shared consumption.

The cooperation factor in our application is an indicator of whether the husband threatens or commits domestic abuse. We recognize that the determinants of domestic violence are complicated, and we do *not* claim to fully model the decision to commit violence. Rather, we simply observe that one component of that decision is the impact it has on the consumption behavior of household members. What our model requires empirically is an instrument that correlates with the decision to threaten or commit abuse. This instrument does not need to be randomly assigned. Instead, it only needs to satisfy some separability conditions based on the properties of our model. We consider a couple of potential instruments, the strongest of which empirically is village level average rates of abuse (we also apply tests of instrument validity).

We find that domestic abuse reduces the efficiency of household consumption by roughly 5 per cent[2], and it shifts roughly 1.5 per cent of household resources towards men and away from women and children. Everyone loses in the face of lower efficiency, but men gain in terms of resource shares. The overall effect of this is to leave the money-metric measure of men's consumption roughly unaffected and to reduce that of women and children by roughly 2.5 per cent. Note that the model allows for multiple reasons why men might commit abuse, e.g., they could gain utility from avoiding the effort required to efficiently cooperate.

1.1 Resource Shares

Expenditure surveys generally collect consumption data at the household level. Standard poverty and welfare measurements based on such data are also typically calculated at the household level. But well-being and utility apply to individuals, not households. When household resources are distributed unequally across household members, official household level measures of income or consumption can seriously mischaracterize the prevalence of poverty and inequality in a country. For example, using the types of methods that we will extend here, DLP find poverty rates for children in Malawi that are much higher than those of men. Another example is Calvi (2019), who finds that older women have poorer access to consumption within households than do younger women. This results in poverty rates among older women that are much higher than those constructed by official household level

[2] This 5 percent means that, if such households had cooperated at the same level as housholds without abuse, then the household members could have increased their utility by an amount equivalent to having 5 percent more household spending.

surveys, which may partly explain the unexpectedly high mortality rate among older women in India.

A key component of collective household models are *resource shares*. Resource shares are defined as the fraction of a household's total resources or budget (spent on consumption goods) that are allocated to each household member. Resource shares are useful for several reasons. First, they are closely related to Pareto weights, and so are often interpreted as measures of the bargaining power of each household member. Second, they provide a measure of consumption inequality within households: if one member has a larger resource share than another member, then they have more consumption. Third, multiplying the resource share by the household budget gives each person's shadow budget. When this shadow budget is appropriately scaled to reflect scale economies, we can compare it to a poverty line and assess whether or not any (or all) household members are poor.

Our primary goal will be identification and estimation of resource shares allowing for inefficiency, and on measuring the economic costs of inefficiency.

1.2 Literature Review

Resource shares are in general difficult to identify, because consumption is typically measured at the household level, and many goods are jointly consumed and/or shareable. The earlier literature on collective household models, which assumes that households reach a Pareto efficient allocation of resources, includes Becker (1965, 1981), Chiappori (1988, 1992), Browning, Bourguignon, Chiappori, and Lechene (1994), Browning and Chiappori (1998), Vermeulen (2002), and Chiappori and Ekeland (2009). This literature showed that, even if one knew the entire demand vector-function of a household (that is, how much the household buys of every good and service as a function of prices, income, and other observed covariates), one still could not identify the level of each household member's resource. However, this earlier work also shows that one can generally identify how these resource shares would change in response to a change in observed covariates such as distribution factors. Other papers that make use of this result include Bourguignon and Chiappori (1994), Chiappori, Fortin and Lacroix (2002), and Blundell, Chiappori and Meghir (2005). Most of this earlier work also constrains goods to be either purely private or purely public within a household, whereas we work with the more general model based on BCL, which also allows some or all goods to be partly shared.

Some interesting policy questions can be addressed without identifying levels of resource shares. However, many fundamental policy questions, such as identifying the prevalence of women's poverty, requires identifying resource share levels. A number of previous approaches exist to address the fundamental nonidentification of resource share levels just from household demand data. One direct approach, taken e.g. by Menon, Perali and Pendakur (2012) and Cherchye, De Rock and Vermeulen (2012), is to collect as much detailed data as possible on the separate consumption of each household member, rather than just observing household-level consumption. To the extent that such data can be collected, resource shares may be observed directly, by taking each individual's observed consumption as a fraction of the total. However, this method requires detailed and difficult data collection, and is likely to suffer from considerable measurement errors, particularly in the allocation of public and shared goods to individual household members.

A second approach is taken by Cherchye, De Rock and Vermeulen (2011). While the levels of resource shares cannot be identified without additional assumptions, these authors show that it is possible to obtain bounds on resource shares, using revealed preference inequalities. Cherchye, De Rock, Lewbel, and Vermeulen (2015) considerably tighten these bounds by combining Slutsky symmetry restrictions with revealed preference inequalities. Bounds can be further tightened, sometimes leading to point identification of shares, by further combining revealed preference inequalities with assumed restrictions regarding marriage markets. See, e.g., Cherchye, De Rock, Demuynck, and Vermeulen (2014).

A third method is to point-identify the level of resource shares from household level data by imposing additional restrictions either on preferences, or on the household's allocation process, or both. Papers that use these methods include Lewbel (2003), Lewbel and Pendakur (2008), Couprie, Peluso and Trannoy (2010), Bargain and Donni (2009, 2012), Lise and Seitz (2011), BCL, DLP, and Dunbar, Lewbel, and Pendakur (2019).

One feature that all of the above cited works have in common is that they assume the household is efficient, in that it reaches the Pareto frontier. While many of the above papers cite evidence supporting these efficient collective models (see, e.g., Bobonis 2009), other papers reject Pareto efficiency within the household, including Udry (1996) and Dercon and Krishnan (2003).

A number of models of noncooperative household behavior exist. Gutierrez (2018) proposes a model that nests both cooperative and noncooperative behavior. Castilla and Walker (2013) provide a model and associated empirical evidence of inefficiency based on information assymmetry, that is, hiding income. Other evidence of income hiding includes Vogley and Pahl (1994) and Ashraf (2009). Similar in spirit to our conditional efficiency frameworks are models that have static efficiency but are intertemporally inefficient. Examples include Mazzocco (2007), Lise and Yamada (2014), and Chiappori and Mazzocco (2017).

Ramos (2016) proposes a model wherein domestic violence affects the efficiency of household production, and estimates the model with Mexican data on household Engel curves. In her model violence is exogenous, does not directly affect men's utility, and only reduces production of a single home produced good. In contrast, we provide a general framework for modeling inefficiency in a collective household. In our model, the cooperation factor is endogenously determined, can have direct effects on the utilities of each household member, affects the allocation of resources within the household, and affects, in different ways, the consumption efficiency of each good, by differentially affecting their the economies of scale. Other relevant papers on the roles and causes of domestic violence include Rao (1997), Angelucci (2008), Boyle, et. al. (2009), Babu and Kar (2010), Krishnan et al (2010), and Anderberg and Rainer (2013).

More generally, our model can be interpreted as a two stage process, where one stage (the determination of the cooperation factor) is non-cooperative and the other stage (consumption decisions) is conditionally cooperative and conditionally efficient. Examples of other models with analogous stages are Lundberg and Pollak (1993), Gobbi (2018), and Doepke and Kindermann (2019).

2 A Class of Inefficient Collective Household Models

In this section we first summarize the BCL model, and generalize it to allow for household inefficiency. We then further generalize the model by allowing the source of inefficiency, the cooperation factor, to be endogenous. This general model could be estimated with sufficiently rich consumption and price data. We next provide simplifying assumptions that allow the model to be estimated with readily available cross section household survey data on household budgets, consumption levels of a few goods, and demographic data. For ease of exposition, derivations here are presented somewhat informally. Formal assumptions and proofs regarding the derivation of the model are provided in the Appendix.

2.1 Collective Households with Varying Consumption Technologies

For simplicity, start with a household consisting of just two members, a husband and a wife, indexed by $j = 1, 2$. Let g denote the vector of continuous quantities of goods purchased by the household. Let p denote the vector of market prices of the goods in g. Let y be the household's budget, that is, total expenditures, which is the total amount of money the household spends on goods. We begin with a simplified version of the BCL model. Given prices p and a budget y, the purchased quantity vector g is determined by

$$\max_{g_1, g_2} U_1(g_1)\,\omega_1(p, y) + U_2(g_2)\,\omega_2(p, y) \tag{1}$$

$$\text{such that} \quad p'g = y, \quad g = A(g_1 + g_2)$$

Here $p'g = y$ is the usual linear budget constraint, g_1 and g_2 are private good equivalents (described below) for person 1 and 2, and A is a matrix. The functions U_1 and U_2 are the utility functions of members 1 and 2, respectively, while ω_1 and ω_2 are the so-called "Pareto Weights" of each member. Each member's Pareto weight summarizes that member's relative bargaining power in a bargaining model, or the relative weight of their utility function in a household social welfare function. The fact that these weight functions can depend on prices p and the budget y is what makes the collective household model more general than a unitary model[3].

Each utility function $U_j(g_j)$ depends on a quantity vector of goods g_j that member j

[3] A unitary model is one that is observationally equivalent to the behavior of a single utility maximizing individual. See, e.g., Chiappori (1988, 1992)

consumes. Goods can be partly shared, and so are not constrained to be purely privately consumed or purely publicly consumed within the household. In equation (1), $g = A(g_1 + g_2)$ is the "consumption technology function," which describes the extent to which each good is shared by the household members. Each household member j consumes (and gets utility from) the quantity vector g_j, which BCL call, "private good equivalents".

The square matrix A summarizes how much goods are shared. Suppose A were diagonal (it need not be, but this case is useful for understanding sharing). The extent to which each element of $g_1 + g_2$ exceeds the corresponding element of g is the extent to which that good is shared by household members. For example, suppose that g^1, the first element of g, was the quantity of gasoline consumed by a couple. If both household members shared their car (riding together) 1/2 of the time, then, in terms of the total distance traveled by each member, it is as if member 1 consumed a quantity g_1^1 of gasoline and member 2 consumed a quantity g_2^1 where $g^1 = (3/4)(g_1^1 + g_2^1)$. For example, Person 1 drives 100km and person 2 drives 100km, but because 50km are driven together, the vehicle only drives 150km. Here, the upper left corner of the matrix A would be 3/4 (which summarizes the extent to which gasoline is shared).

Non-zero off-diagonal elements of A allow the sharing of one good to depend on the purchases of other goods, e.g., more gasoline might be shared by households that purchase less public transportation. As a result, the model is also equivalent to some restricted forms of home production, e.g., a household that wastes less food by cooperating and coordinating on the production of meals could be represented by having a lower value of the k'th element on the diagonal of the matrix A, where g^k is the quantity of purchased food.

Given some regularity conditions (see the Appendix for details), there exist resource share functions $\eta_j(p, y)$ such that the household's behavior is equivalent to each member j solving the program

$$\max_{g_j} U_j(g_j) \quad \text{such that} \quad p' A g_j = \eta_j(p, y)\, y \tag{2}$$

Each η_j is the fraction of the household's total resources y that are claimed by member j. Resource shares sum to one, so that with two household members we have $\eta_1 + \eta_2 = 1$. Equation (2) is the key decentralization result: the couple's behavior is observationally equivalent to a model where each member j chooses a consumption vector g_j to maximize their own utility function, subject to their own personal budget constraint, which has shadow price vector $A'p$ and shadow budget $\eta_j(p, y) y$.

Let $g_j = h_j(p, y)$ be the demand equations that would be obtained from maximizing the utility function $U_j(g_j)$ under the linear budget constraint $p'g_j = y$. Each member j faces the constraint in equation (2), so

$$g_j = h_j(p'A, \eta_j(p, y) y) \tag{3}$$

and $g = A(g_1 + g_2)$ so the household's demand equations are

$$g = A(h_1(p'A, \eta_1(p, y) y) + h_2(p'A, \eta_2(p, y) y)). \tag{4}$$

BCL show that if the demand functions h_j are known, then the consumption technology matrix A and the resource share functions $\eta_j(p, y)$ are generically identified from household demand data. They suggest that the h_j demand functions could be identified from observing the demands of people living alone.

A feature of this model is that, the more that goods are shared, the lower is the shadow price vector $A'p$ relative to market prices p, reflecting the greater efficiency associated with increased sharing. In the gasoline example above, the shadow price of gasoline is 3/4 that of the market price. This means that the household's actual expenditures on gasoline, g^1p^1,,is equal to the cost of buying the sum of what the couple consumed, $g_1^1 + g_2^1$, at the shadow price of $(3/4)p^1$. If the couple had consumed the total quantity of gasoline $g_1^1 + g_2^1$ without any sharing, it would have cost $(g_1^1 + g_2^1) p_1$ dollars instead of what they actually spent, $g^1p^1 = (3/4)(g_1^1 + g_2^1)p_1$. The difference between these two is the dollar savings they obtained by sharing gasoline.

Analogous gains are obtained with all goods that are shared. The more efficient the household is, (i.e., the more they share consumption), the greater is the difference between what they would have had to spend on all goods if they hadn't shared, which is $p'(g_1 + g_2) = p'A^{-1}g$, relative to what they actually spent, which is $y = p'g$. Thus, the matrix A embodies the scale economies due to sharing that are available to the household.

Now consider two couples that differ in how much they are able to share consumption goods, and so have different consumption technology matrices A_0 and A_1. Suppose the couple with A_0 is more efficient in their consumption, meaning that they share more. Then, even if both couples bought the same market quantity of goods g, the first would have higher consumption of private good equivalents. By the above logic, this increased efficiency in dollar terms equals the difference between $p'A_0^{-1}g$ and $p'A_1^{-1}g$.

Even though the second of these couples is inefficient relative to the first, each is conditionally efficient, conditioning on each couple's ability to share and cooperate. Equivalently, each is conditionally efficient, conditioning on the consumption technology matrix that they possess (either A_0 or A_1). Because they are each conditionally efficient, each household's decision problem can be represented by the decentralized program (2).

Now suppose we have two sets of households. One set has consumption technology matrix A_0 and the other set has A_1. Even though the latter households are inefficient, we can still apply and estimate the collective household model to each set of households separately. In particular, we can treat inefficient households as if they were Pareto efficient, satisfying decentralization and other properties of efficiency, because they are conditionally efficient, conditioning on their particular consumption technology matrix A_1.

2.2 Collective Households With Endogenous Inefficiency

In the previous subsection, each household possessed an ability to cooperate (in terms of sharing consumption) given by a matrix A_f. We call the f index a "cooperation factor." A cooperation factor is an observable behavior f that affects the household's level of cooperation and hence their level of sharing. As noted earlier, examples of cooperation factors could be direct indicators of cooperation (like the relative time each member spent on household chores), or behaviors associated with likely cooperation or failures to cooperate, such as money hiding or domestic abuse. We will now let f be an endogenous choice. Again derivations here are presented informally for ease of exposition. Formal assumptions, theorems and proofs are in the appendix.

Here we generalize the model of the previous section. First, we allow for an arbitrary number of household members instead of two. Second, we incorporate f into the model, reflecting all of its potential impacts on the household. Third, we let f be a choice variable. The resulting model of the household is now

$$\max_{g_1,g_1,\dots g_J,g_J} \sum\nolimits_{j=1}^{J} \left(U_j\left(g_j\right) + u_j\left(f,v\right)\right) \omega_j\left(p,y,f\right) \tag{5}$$

$$\text{such that } \; p'g = y, \quad g = A_f \sum\nolimits_{j=1}^{J} g_j$$

The new variable v is discussed below. As before, assume the most efficient value for A is A_0. The factor f appears in three places in this model. First, f affects sharing through A_f. Second, f appears in the Pareto weight functions ω_j, showing its potential impact on relative power, and the associated allocation of resources, among household members. Third, f directly affects member utilities through the u_j functions.

The term $u_j\left(f,v\right)$ is the utility member j directly experiences (not including his or her utility over goods) from living in a household with cooperation factor f. The variable v is any observed covariate (or vector of covariates) that affects this utility associated with f, but does not affect the rest of the model. The role of the variable v for identification and estimation is discussed below.

To illustrate, if cooperating at the level A_0 instead of A_1 requires more effort, $u_j\left(1,v\right)$ may be negative, reflecting member j's disutility from expending that extra effort. Or, if f is an indicator of domestic abuse committed by member 1, $u_1\left(1,v\right)$ could include the disutility (from guilt) or utility (from feeling powerful) that member 1 directly experiences by committing the abuse, while $u_2\left(1,v\right)$ could be the direct disutility member 2 experiences from being subjected to abuse.

Generalizing the model to equation (5) means that the resource share functions η_j now depend on f, and the demand equations (3) and (4) become

$$g_j = h_j\left(p'A_f, \eta_j\left(p,y,f\right)y\right) \tag{6}$$

and

$$g = A_f \sum\nolimits_{j=1}^{J} h_j\left(p'A_f, \eta_j\left(p,y,f\right)y\right) \tag{7}$$

Now consider what happens to this model when f becomes a choice variable. First, as discussed in the previous section, the household remains conditionally efficient, conditioning on the chosen level of f, so equations (6) and (7) continue to hold. Second, we must consider how f is chosen. Suppose that member 1 is the husband, and he chooses whether to cooperate or not (e.g., he chooses whether to help with chores, or to commit abuse, or to hide money). We assume he chooses f to maximize his own attainable utility level, which means

$$f = \arg\max U_1 \left(h_1 \left(p' A_f, \eta_1 \left(p, y, f \right) y \right) \right) + u_1 \left(f, v \right). \tag{8}$$

Equation (8) is obtained by plugging the household's maximized value of g_1, given by equation (6), into member 1's utility function.

We will *not* need to actually specify or estimate equation (8). This is important because we may know very little about u_1, the husband's direct utility from choosing f. For example, the decisions to hide money or commit violence likely have many determinants we cannot observe in our consumption data. In short, while our theoretical model determines f by equation (8), our empirical model does *not* require trying to estimate this equation, because there are likely to be too many relevant factors (specifically, factors affecting u_1) we cannot observe.

Relative to the efficient $f = 0$, equation (8) shows that choosing $f = 1$ has three effects on member 1's utility. First, it raises shadow prices $p' A_f$, reflecting that fact that, by reducing cooperation, the total effective quantities for consumption by the household, $\sum_{j=1}^{J} g_j$, are reduced. All else equal, this would lower member 1's utility. Second, the inefficient f could also raise or lower member 1's resource share η_1. For example, he might engage in abuse, or hide money, specifically to raise his resource share. Finally, $u_1 \left(1, v \right)$ could be positive if he likes engaging the behaviors that raise f, or if he dislikes the effort required to cooperate on consumption. Similarly $u_1 \left(1, v \right)$ could be negative if he feels guilt or remorse from behaviors that raise f, or if he enjoys cooperating on consumption.[4] In this model,

[4]Since member 2 loses utility if $f = 1$ is chosen, she would have an incentive to pay member 1 to choose $f = 0$ instead. Any such payment would correspond to deceasing η_2 and increasing η_1, which is another reason why these resource share function could depend on f. Even in the presence of such side payments, the household could still end up choosing $f = 1$. This will occur if the magnitude of the required side payment would lower member 2's utility more than her loss in utility due to having $f = 1$.

inefficiency is possible because member 1's choice of the cooperation factor f can impose negative externalities on the utility of other household members through all of the above channels.

Given sufficient data, the household's demand equations (7) could be estimated as described by BCL. The main additional complication here is that f would be an endogenous regressor. However, this is where the role of the covariate v comes in. As can be seen in equation (8), the variable v affects the direct utility u_1 member 1 gets from his choice of f, and so correlates with his choice of f. However, by equations (6), (7), v does not affect the household's demand functions for goods (except to the extent that v helps to determine f). This is because v only affects the u_j functions, not utility from goods consumption U_j or Pareto weights ω_j. This means that v is a valid instrument for the endogenous f in the demand equations. To illustrate, suppose f indicates domestic abuse. Then v must be a variable that doesn't directly affect the demands for goods, but does affect the likelihood of committing abuse. Examples of v could be measures of the prevalence or acceptability of abuse in the household's village.

An important feature of our model is that we do *not* need to actually estimate equation (8). In particular, the function $u_1(f, v)$ would in general depend on many other variables that we do not observe. So, e.g., in our empirical application where f is domestic abuse, we do not claim to model all of the many complicated determinents of household violence. Instead, what we observe and make use of is just some variable v that affects the decision to threaten or commit abuse.

Given estimates of the model, particularly of A_f, we could calculate dollar costs of inefficiency, such as the difference between $p'A_0^{-1}g$ and $p'A_1^{-1}g$. However, we would not be able to estimate the direct costs of violence in terms of the functions $u_j(1, v)$. For example, we would not be able to put a price on the sadness felt by household members as a result of the abuse.

A final complication is that the BCL model, like many earlier collective household models, does not identify the resource shares of children, and without additional assumptions, our extension here would similarly not identify impacts on children. We address this issue, and overcome other data limitations, in the next subsection.

2.3 Empirically Practical Identification and Estimation

Estimation of the above model is complicated. DLP propose a restricted version of the BCL model that has many convenient features for empirical work. Here we propose restrictions, similar to those used in DLP, to obtain a version of our collective model that has many advantages for empirical work, including: 1) the model can be estimated using readily available "Engel curve" data, that is, cross sectional data on expenditures without price variation; 2) the model identifies resource shares for children as well as adult household members, and 3) despite lacking price variation, the model still identifies the economic cost of inefficiency. Further, we extend the model of the previous section to allow for both observed and unobserved preference heterogeneity, and to households having more than one member of each type (in particular, multiple children). As in the previous subsections, we summarize our main results in the text here, while providing formal assumptions, derivations, and identification proofs in Appendix A.

As in DLP, our estimating equations are based on private, assignable goods. A good j is *private* if it is consumed by a single member and its diagonal element of the matrix A equals one, meaning it cannot be jointly consumed at all. A good j is *assignable* if the researcher knows which household member consumes it. Assume now that each household member j consumes a quantity q_j of a good that is private and assignable to member j. This is another good, in addition to the K vector of quantities of goods g_j that each household member j consumes. Let $\pi = (\pi_1, ..., \pi_J)$ denote the vector of prices of these private assignable goods.[5]

We further generalize the model by allowing prices to affect u_j (since there is no *a priori* economic reason for excluding them, and like v, prices appearing in u_j only affect the determination of f, not the demand functions for goods). We also now include additional observed household-level demographic variables z (which can affect both tastes and Pareto weights) to allow for observable heterogeneity across households. Taking all this into account,

[5] In practice, the private assignable goods will often all have the same price, making $\pi_1 = ... = \pi_J$. For example, the private assignable good could be rice if we observed how much rice each household member eats, and rice has the same market price for all household members. As with DLP, some of the formal assumptions of our model are easier to satisfy when the private assignable goods all have the same price.

the model of equation (5) becomes

$$\max_{g_1,q_1,\dots g_J,q_J} \sum\nolimits_{j=1}^{J} \left[U_j\left(q_j,g_j,z\right)+u_j\left(f,v,z,p,\pi,y\right)\right]\omega_j\left(f,z,p,\pi,y\right) \tag{9}$$

$$\text{such that } \; p'g+\sum\nolimits_{j=1}^{J}\pi_J q_j = y \;\text{ and }\; g = A_f\sum\nolimits_{j=1}^{J} g_j$$

Note the budget constraint is comprised of spending on private assignables q_j and market purchases of the shared unassignable goods g, which are converted to the sum of private equivalents $\sum_{j=1}^{J} g_j$ by the matrix A_f.

This model yields household demand functions for vectors of goods g and q, analogous to those of equation (7). But for the private assignable goods q, these demand functions greatly simplify, because for each private assignable good the quantity q_j that is consumed by member j is the same as the quantity purchased by the household. For these private assignable goods, the household demand equations arising from the household model of equation (9) have the form

$$q_j = H_j\left(p'A_f,\pi,z,\eta_j\left(p,\pi,y,f,z\right)y\right) \tag{10}$$

where H_j is the Marshallian demand function for the good q_j that comes from the utility function $U_j\left(q_j,g_j,z\right)$. Note the resource share functions η_j may now depend on the additional variables π and z that we've introduced into the model. But importantly, as a result of the household's consumption optimizing behavior, v does not appear in this equation. This is what makes v be a valid instrument for f (see the Appendix for details).

We now make some simplifying assumptions (again, details are in the Appendix) to transform this model of price-dependent demand equations into a model of Engel curves giving demands at fixed prices. First, we assume that the resource share function η_j does not depend on y. This assumption is also made by DLP, who provide a range of theoretical and empirical arguments in support of this assumption.

Let $V_j\left(\pi_j,p,y,z\right)$ denote the indirect utility function corresponding to the maximization of the direct utility function $U_j\left(q_j,g_j,z\right)$ under the hypothetical linear budget constraint $q_j\pi_j+g_jp=y$. The actual utility level over goods attained by member j in the household (which does not include the u_j component of utility) equals this indirect utility function V_j evaluated at the household's shadow prices $A_f p$ and member $j's$ shadow budget $\eta_j(\pi, A_f p, f, z)\, y$.

The second main simplifying assumption we make is that this attained level of indirect utility is semiparametrically restricted to have the form

$$U_j = \left[\ln \eta_j(\pi, A_f p, f, z) + \ln y - \ln s_j(\pi_j, p, z) + \varepsilon_j^*(\pi_j, p) + \ln \zeta(A_f p, z)\right]\left[m_j(A_f p, z) - \beta(z)\ln \pi\right] \tag{11}$$

for some functions s_j, ζ, m_j, and β, where, without loss of generality $\ln \zeta(A_0 p, z) = 0$. Here $\varepsilon_j^*(\pi_j, p)$ is an unobserved taste shifter, i.e., a random utility parameter.

The restrictions imposed by equation (11) have empirical support, e.g., the popular Deaton and Muellbauer (1980) Almost Ideal Demand System model is a special case of equation (11). This equation also satisfies the SAP (similar across people) restriction used by DLP, which they show also has empirical support.[6]

The decentralization described in the previous subsections carries over to this model. As shown in the appendix, this allows us to apply Roys identity to equation (11) to obtain the household's demand functions for each private assignable good j. The resulting demand functions are most conveniently represented in budget share form. Let $w_j = \pi_j q_j / y$ be the budget share for each private assignable good j, giving the fraction of the household's budget y that is spent on buying member j's private assignable good. The demand functions determining each q_j can therefore be multiplied by π_j / y to give corresponding demands expressed in terms of budget shares w_j.

We will estimate our model using data from a single price regime, so both p and π are treated as constants, which can then be absorbed into the functions that comprise the budget share demand equations. After introducing the random utility parameters, deriving the budget share demand functions from equation (11) using Roy's identity, and treating all prices as constants, we obtain budget share demand functions that we show in the Appendix

[6] Equation (11) also implies restrictions on A_f relative to the range of possible vectors p. These restrictions are comparable to those imposed by other empirical consumer demand models. See Lewbel and Pendakur (2008) and Appendix A for details.

take the Engel curve form

$$w_j = \eta_j(f,z)\left[\gamma_j(z) - \beta(z)\left(\ln y + \ln \eta_j(f,z) + \ln \delta(f,z)\right) + \varepsilon_j\right] \tag{12}$$

Here $\eta_j(f,z)$ is member j's resource share function, $\gamma_j(z)$ and $\beta(z)$ are functions representing variation in tastes, ε_j is an error term that comes from the random utility parameters, and $\delta(f,z)$, which equals $\zeta(A_f p, z)$ at the fixed level of prices, is a function that reveals the dollar costs of inefficiency as described below. In the Appendix we add random utility parameters to the model, including an unobserved taste shifter that is added to the function $\ln s_j(\pi_j, p)$ in equation (11).

We prove in Appendix A that the functions in equation (12) are each nonparametrically identified. This includes showing that the levels of the resource shares, $\eta_j(f,z)$, are identified. The function $\delta(f,z)$, which is also identified, has an important interpretation in our model.

Recall our assumption that member 1 chooses f to maximize his own utility. That means f is chosen to maximize $U_1 + u_1$. We show in the Appendix that in general the resulting value of f is endogenous (i.e., it is correlated with ε_j), but also that v is a valid instrument for f.

Inspection of equation (12) shows that the cooperation factor f has two effects on the budget shares of private assignable goods. One is that it affects resource shares η_j. The second effect, which is on A_f, affects the Engel curve demands through the function $\delta(f,z)$. Inspection of equations (11) and (12) shows that a change in $\ln \delta(f,z)$ has the same effect on utility and on budget shares as the same change in $\ln y$. This then provides a dollar measure of the unconditional efficiency loss to the household resulting from choosing $f > 0$.

Since $\ln \delta(0,z) = 0$, a change from $f = 0$ to a level of $f > 0$ is equivalent, in terms of consumption of goods, to a change in the household's budget from y to $y\delta(f,z)$. The reduction in sharing from an increase in f has the same effect on demands, and on the member's attained utility levels over goods, as a reduction in total expenditures y. The term $\delta(f,z)$ measures the size of this reduction. Note that although we identify and estimate $\delta(f,z)$ using just the private assignable goods, this function actually measures the impact of f on the efficiency of consumption of *all* goods, because it is equivalent in everyone's utility function to a change in the total budget y.

All of the derivations in this section go through allowing the cooperation factor f to take many different values (where we have normalized the most efficient case to be $f = 0$). However, in our empirical application we will just let f take on two values zero and 1, as in our earlier discussions.

The model we estimate is based on equation (12) for each private assignable good $j \in \{1, ..., J\}$. Recall that f is endogenous. In our data, the observed y is partly constructed and so may contain measurement error. The budget y could also be endogenous, because it's a choice variable. That is, if one considers the dynamic optimization problem of the household, given the household's income and assets, we are assuming it first decides how much to spend on consumption this period (that is, if first chooses y), and then decides what fraction of y to spend on buying each good. This latter decision is what the household's program, equation (9), determines.

Let r be a vector of observed variables that may affect the determination of y, such as functions of the household's income or wealth. We assume ε_j is uncorrelated with r, either because the measurement error in y is unrelated to r, or (if y is endogenous) because ε_j is only based on random utility associated with the within period budget allocation, not the utility of saving vs spending. Elements of v might also be correlated with or include these measures of income, assets, or wealth, so to be as general as possible, we let r also include v.

Equation (12) then yields conditional moments of the form

$$E\left(\frac{w_j}{\eta_j(f,z)} - \gamma_j(z) - \beta(z)\left(\ln y + \ln \eta_j(f,z) + \ln \delta(f,z)\right) \mid r, z\right) = 0$$

We show in the Appendix that the model can be nonparametrically identified from these conditional moments.

However, given limitations on the size of the data set and complexity of the model, it is more practical to estimate the model parametrically. Another data limitation is that the household may have more than one member of each type j, and we may not observe an assignable good for each. In particular, most households have multiple children. Let N_j be the number of members in the household of type j. Equation (12) is the budget share demand function of each such member. Since w_j is the budget share of food that is assignable to all household members of type j, the resource share of any one member of type j is $\eta_j(f,z)/N_j$.

Letting θ be a vector of parameters, we parameterize each of the functions in the above equation, and incorporate N_j, to obtain unconditional moments

$$E\left[\left(\frac{w_j}{\eta_j(f,z,\theta)}-\gamma_j(z,\theta)-\beta(z,\theta)(\ln y-\ln N_{jh}+\ln\eta_j(f,z,\theta)+\ln\delta(f,z,\theta))\right)\phi(r,z)\right]=0 \tag{13}$$

Equation (13) holds for any vector of bounded functions $\phi(r,z)$. We construct an estimator for θ by choosing functions $\phi(r,z)$ as discussed in the Appendix, and applying Hansen's (1982) Generalized Method of Moments (GMM).

3 Application to households in Rural Bangladesh

3.1 Data

We use data from the 2015 Bangladesh Integrated Household Survey. This dataset is based on a household survey panel conducted jointly by the International Food Policy Research Institute and the World Bank. In this survey, a detailed questionnaire was administered to a sample of rural Bangladeshi households. This data set has two useful features for our model: 1) it includes person-level data on food consumption as well as total household expenditures; and 2) it includes recall data on the exposure of the primary female spouse to both physical and verbal abuse. The former allows us to use food, a large and important element of consumption, as an assignable good to identify our collective household model parameters. The latter allows us to divide households into those with and without reported violence, which we treat as a cooperation factor.

The questionnaire was initially administered to 6503 households in 2012, drawn from a representative sample frame of all Bangladeshi rural households. 6436 households remained in the sample in 2015, and of these, 6142 households reported data on household expenditures. We drop 11 households with a discrepancy between people reported present in the household and the personal food consumption record, and we drop 8 households with no daily food diary data.

Some of these households have many adult members, and such households may have more complicated interactions regarding the choice of f than our model assumes. However, nuclear families, i.e., those with just one adult man and one adult woman, make up less than 30 per cent of all households in the data. We therefore will include non-nuclear families in our sample, but will also report estimates based on just nuclear households. Define the *composition* of a household to be its number of men, number of women, and number of children (in these data, children are defined as members aged 14 or less). To eliminate households with unusual compositions, we select households that have at least 1 man, 1 woman and 1 child, and for which there are at least 100 households with the given composition in our data. The resulting sample consists of households with 1 or 2 men, 1 or 2 women, and 1 or 2 children, plus additional nuclear households with 1 man, 1 woman and 3 or 4 children. This leaves us with 3087 households. Of these, we drop all households that report zero food consumption for men, women or children, leaving us with 2866 households in our final estimation sample. Households are indexed by $h = 1, ..., H$, so $H = 2866$ in our main estimation sample.

We use food consumption as our assignable goods. The survey contain 7-day recall data on household-level quantities (measured in kilograms) of food consumption in 7 categories: Cereals, Pulses, Oils; Vegetables; Fruits; Proteins; Drinks and Others. These consumption quantities include home-produced food and purchased food and gifts. They include both food consumed in the home (both cooked at home and prepared ready-to-eat food), as well as food consumed outside the home (at food carts or restaurants). Thus, we have the widest possible definition of food consumption. For each of these food categories, we construct village-level average prices as equal to the village-level total spending in the food category divided by the village-level total quantity in the food category.

Note that all references to the "village level" in this paper actually refer to data collected at the Upazila level, which are official administrative units in Bangladesh, one level below the district. There were 492 Upazilas in Bangladesh in 2015, of which 281 are represented in this exclusively rural dataset.

To calculate the fraction of these household level food expenditures that are separately consumed by men, women, and children within the household, we make use of an additional one-day recall diary of individual-level quantities of food in the 7 categories. These are the quantities of food that are consumed by each individual in the household, and so do not include leftovers or food served to guests. We multiply each individual's share of the household's one-day quantities in each category by household-level weekly quantity to get individual-level weekly quantity by category. These are summed over the 7 categories and multiplied by village-level prices to get total individual-level weekly expenditure on food, and are multiplied by 52 to get individual-level annual food spending. Finally, we aggregate individuals by type to yield adult male food spending, s_{mh}, adult female food spending, s_{fh}, and children's food spending, s_{ch}.[7]

The model uses assignable good budget-shares of household-level total expenditure. Our household-level total expenditure measure is equal to the twelve times the sum of household-level monthly spending, including imputed consumption of home produced goods. These spending levels derive from one-month duration recall data in the questionnaire. Specifically, this includes monthly-level recall data on purchases and home-produced values of: rent, food, clothing, footwear, bedding, nonrent housing expense, medical expenses, education, remittances, devotional/sacrificial goods[8], entertainment, fines and legal expenses, utensils, furniture, personal items, lights, fuel and lighting energy, personal care, cleaning, transport and telecommunication, use-value from assets, and other miscellaneous items. This constructed total expenditures variable, denoted y_h, represents the total flow of consumption of goods and services into the household, which includes purchases, home produced goods and consumption flows from assets. The budget-shares of each type of person, $j = m, f, c$, are denoted w_{jh} and are given by $w_{jh} = s_{jh}/y_h$.

Our models are also conditioned on a set of demographic variables z_h. We include several types of observed covariates in z_h. We condition on household size and structure, defined as

[7] Specifically, let Q_{ph} be the observed quantity of category p, $p = 1, ...7$, for household h and let S_{ph} be the observed spending for the weekly food recall data. Following Deaton (1993), let π_{ph} be the village-level unit value equal to village-level aggregate spending, S_{ph}, divided by village-level agregate quantity, Q_{ph}. We take π_{ph} to be the local price of food category p.

Let $\widetilde{q}_{jph}$ be the observed quantity of category p for all people of type j in household h from the one-day diary data. One-day diary data do not include spending data. We take shares of each category, $\left(\widetilde{q}_{jph}/\sum_j \widetilde{q}_{ph}\right)$, and attribute to each type of person their share of weekly quantities in each category, multiply these by the local price of that category, multiply by 52 to generate food spending by type: $s_{jh} = 52 * \sum_p \pi_{ph} \left(\widetilde{q}_{jph}/\sum_j \widetilde{q}_{ph}\right) Q_{ph}$

[8] These are: jakat, fitra, daan, sodka, kurbani, milad, and other religious offerings.

a set of 10 dummy variables covering all combinations of 1 or 2 men, 1 or 2 women, and 1 or 2 children plus the additional nuclear families consisting of 1 man, 1 woman, and 3 or 4 children. The left-out dummy variable is the indicator for a household with 1 man, 1 woman and 2 children (the largest single composition). We call this particular nuclear household type the reference composition.

We also include other variables in z_h that may affect both preferences and resource shares: 1) the average age of adult males divided by 10; 2) the average age of adult females divided by 10; 3) the average age of children divided by 10; 4) the average education in years of adult males; 5) the average education in years of adult females ; 6) the fraction of children that are girls minus 0.5; (7) the log of marital wealth (aka: dowry); (8) the log of household wealth.

We do not normalize dichotomous composition variables or the fraction of girl children. However, we normalize all other elements of z to be mean-zero for households with the reference composition. These normalizations give $z_h = 0$ for a *reference household* defined by reference composition and all covariates equal to the mean values for the reference composition. We also normalize the log of household expenditure, $\ln y_h$, to be mean 0 for the reference composition. These normalizations simplify the economic interpretation of our estimated coefficients, since they then either provide estimates of the behavior of the reference household type, or (in the case of coefficients of z_h) they describe departures from the reference household's behavior.

In our empirical application, we take the cooperation factor for household h, f_h, to be an indicator of threatened or actual abuse. Specifically, our recall survey includes the following questions asked of the head female in the household: "Has any of the following happened to you in the past year? Your husband threatened you with divorce? Your husband, another family member, or household resident verbally abused you? Your husband, another family member, or household resident physically abused you?" We define f_h to equal to 1 if the head female responds yes to any of these questions, and no otherwise.

Other variables that relate to consumption and to f_h, comprising the vector of instruments r_h, are log of household income, the square of log household income, a dummy variable that equals one if the household's home is built of either mud or bamboo (as opposed to stronger materials like concrete or wood), and the average number of households in the village of residence (excluding household h) that report abuse.

Table 1a: Distribution of Household Structures

men	women	children	variable name	mean
1	1	1	m1_f1_c1	0.180
		2	constant	0.257
		3	m1_f1_c3	0.103
		4	m1_f1_c4	0.030
1	2	1	m1_f2_c1	0.086
		2	m1_f2_c2	0.086
2	1	1	m2_f1_c1	0.081
		2	m2_f1_c2	0.055
2	2	1	m2_f2_c1	0.072
		2	m2_f2_c2	0.048

Table 1a gives summary statistics regarding household structures. The 10 summarized household structures each correspond to a dummy variable included in the list of demographic shifters z_h (except for the omitted reference household). Nuclear households (with only 1 adult male and 1 adult female) account for roughly half of the households in our sample. Roughly 30 per cent of households have 3 adults.

Table 1b: Summary Statistics

	Mean	Std Dev	Min	Max
ln(total consumption), $\ln y$	0.105	0.554	-1.676	2.769
men's food, w_m	0.161	0.070	0.014	0.514
women's food, w_f	0.145	0.065	0.013	0.534
children's food, w_c	0.131	0.080	0.001	0.488
age men	0.165	1.179	-2.274	6.026
age women	0.374	0.924	-1.336	5.864
education men	0.417	3.538	-3.387	6.613
education women	-0.342	3.185	-4.311	5.689
age children	0.047	0.355	-0.718	0.682
fraction girls	-0.028	0.412	-0.500	0.500
ln(dowry)	-0.416	3.369	-8.705	5.667
ln(wealth)	0.077	2.681	-9.409	4.351
f, abuse	0.420	0.494	0.000	1.000
f, village average	0.420	0.265	0.000	1.000
Building Mat: Mud, Bamboo	0.156	0.363	0.000	1.000
ln(income)	0.083	1.440	-8.378	3.157

Table 1b gives summary statistics on the log of household expenditures $\ln y_h$, assignable food budget shares w_{jh}, additional demographic shifters (the elements of z_h other than household structure dummies), the abuse indicator f_h, and our instrumental variables (household income, village-average abuse, and building materials). Recall that all continuous regressors (except the fraction of girls) and instruments are normalized to average zero for households with 1 man, 1 woman and 2 children. However, they do not average zero for the entire sample. Dummy variable covariates are not normalized. Village-average abuse $\overline{f}_h$ is the leave-out average for each household, and it is also unnormalized. We measure age in decades, education in years and total consumption, dowry, wealth and income in Taka, the local currency of Bangladesh. These units are chosen to keep the standard deviations of dependent variables, covariates and instruments roughly comparable.

We note a couple of important features of these data. First, the assignable good budget shares (w_{mh}, w_{fh} and w_{ch}) are large; roughly 10 per cent of the household budget goes to each of these assignable food aggregates. This is in sharp contrast to other research (e.g., Calvi 2019) that uses clothing instead of food as the assignable good, where clothing shares may be less than 1 per cent of the household budget. Second, the abuse indicator f_h has a mean of 0.42, suggesting that verbal and physical abuse are high incidence phenomena in Bangladesh. The village-level leave-out average of abuse has a standard deviation of 0.265, which suggests that much of the variation in abuse is at the village level.

3.2 Instruments

Our model has two endogeneous regressors: the log of household total expenditures, ln y_h, and the cooperation factor f_h. As discussed earlier, if we assume that the consumption allocation decision in our model is separable from the decision of how to allocate household income between total consumption and savings, then functions of household income are valid instruments for ln y_h. This is a standard assumption in the consumer demand literature, including in collective household models (see, e.g., Lewbel and Pendakur 2008). We discuss time separability formally in the Appendix. We therefore include log household income and its square in our vector of instruments r_h. Another reason y_h could potentially be endogeneous is measurement error, stemming from, e.g., purchase mismeasurement, or infrequency of expenditures on some consumption items. Our income measures are also valid instruments for dealing with expenditure measurement issues (see, e.g., Banks, Blundell, and Lewbel 1997), since income is measured at the annual level (as opposed to the one month consumption recall), and so smooths out noise due to infrequency.

Now consider instruments for the domestic violence indicator f_h. In the derivation of our model, the cooperation factor f_h is determined by Equation (8) which, after adding in additional covariates and unobserved heterogeneity, becomes equation (25) in the Appendix. We do not attempt to specify and estimate this equation. Among other obstacles to doing so, the function u_1, which summarizes the direct utility the husband gets from committing violence, likely contains very many determinents that we cannot observe.

Without estimating a model for f_h, what we require is an instrument v_h for f_h. This instrument does not need to be randomly assigned, but it does need to correlate with the choice of f_h, while not (after conditioning on other covariates) directly affecting the household's food consumption decisions (in terms of the model, v_h must appear in the function u_1 but not in the functions U_j and ω_j for $j = 1, ..., J$).

A literature on family violence in India has shown that three primary correlates of domestic violence are: alcohol consumption; insufficient dowry (more generally associated with wealth); and local social acceptance of violence (see, e.g., Rao 1997, Boyle, et. al. 2009, Babu and Kar 2010, or Krishnan et al 2010). Our data do not contain information on alcohol consumption or purchases. We do have data on dowry and wealth, however, these correlate with many consumption related household decisions, and so are not suitable as instruments, but instead appear among the z_h covariates. This leaves us looking for instruments that relate to local acceptability of violence. We consider two such potential instruments. Since only one is needed for identification, this gives us overidentifying restrictions that we can use to test instrument validity.

The first potential instrument we consider a dummy variable that equals one if the household's home is built of either mud or bamboo (as opposed to stronger materials like concrete or wood).[9] The relevance of building materials is that some materials are less

[9] We have data on whether or not the household building material is: 1) Concrete/Brick; 2) Tin/CI Sheet; 3) Wood; 4) Mud; 5) Bamboo; 6) Jute/Straw; 7) Plastic; 8) Golpaata/Leaf; 9) Grass/Straw; and 10) Other. Categories 1-3 represent 73% of observations; categories 4-5 represent 22% and categories 6-9 represent the

sound-permeable than others, and therefore offer more privacy than others. This could in turn affect the incentive to engage in domestic abuse, to the extent that abuse is detectable and censured by neighbours. Building materials could be invalid as instruments if they correlate with consumption decisions. The likeliest source of such correlation would be through wealth, which correlates with both consumption and building materials. We control for this correlation by including both log wealth and dowry as covariates in the model, allowing them to affect preferences, resource shares, and f (this is in addition to how the model includes total expenditures y).

Even if valid after appropriate conditioning, there are reasons to think building material might be a relatively weak instrument, e.g., we don't have data on the distance between dwellings or on the extent to which disutility from violence is affected by what neighbors hear. We therefore also consider a direct measure of local social acceptance of violence as instrument, namely, the average number of households in the village of residence (excluding household h) that report abuse.

The idea here is that variation in the local prevalence of domestic abuse is likely to correlate with local tolerance or acceptance of abuse, which in turn affects an individual's utility or disutility from engaging in abuse. In particular, a village where domestic abuse is common is likely to be one where abuse is more tolerated by the community. In this case, the disincentive to engage in domestic abuse is weaker. If this variation in village-level abuse (and therefore tolerance) is uncorrelated with unobserved preference heterogeneity in the demand for food, ε_{jh}, then it is a valid instrument.

To more formally define conditions under which village-level abuse is a valid instrument, assume the household h random utility parameters $\widetilde{e}_{1fh}$ and $\widetilde{\varepsilon}_{jh}$ defined in Appendix A are independent across households. Let $\overline{f}_h$ equal the expected value of f_h conditional on being a household other than h in the village that h resides. Then $\overline{f}_h$ is the probability that a randomly chosen household in the village, other than household h, commits violence. Assume that we include $\overline{f}_h$ in the function R defined in Appendix A. Taking the conditional mean

remaining 5% of households. A priori, we suspected that categories 1-3 are least sound permeable and 4-9 more sound permeable. However, of 4-9, only 4 and 5 are common, and only the contrast between 1-3 and 4-5 is correlated with abuse (the contrast between 1-3 and 6-9 is not significantly correlated with abuse). We considered models with instruments for categories 4-5 and categories 6-9. These yielded very similar results, and somewhat smaller standard errors.

of Equation (25) across households other than household h in the village then shows that $\overline{f}_h$ equals a function of the joint distribution of $y_{h'}$, $r_{h'}$, $z_{h'}$, $\widetilde{e}_{1fh'}$ and $\widetilde{\varepsilon}_{1h'}$ across all households h' other than h in the village. It follows that $\overline{f}_h$ is a useful instrument in that it affects the choice of f (by being in R) and is a valid instrument in the quantity demand equations because $\overline{f}_h$ is independent of household h's specific value of $\widetilde{\varepsilon}_{jh}$ and hence of ε_{jh}.

As this derivation shows, validity of the village-level abuse instrument depends on some independence assumptions regarding the random utility parameters that are not otherwise required by the model. In particular, our estimates cluster standard errors at the village level, and so, other than to make use of this instrument, we would not need to assume independence of unobserved model heterogeneity across households within villages. We will therefore later, as a robustness check, report estimates that drop the village-level abuse measure as an instrument.

For estimation, we do not need to distinguish which elements of the instrument list r_h are intended to be specifically instruments for f_h vs for y_h (i.e., elements of v vs elements of $\widetilde{r}$ in the Appendix). In particular, though we argue that village-level abuse and building materials should primarily correlate with f_h and income primarily with y_h, either or both could affect both.

To assess the relevance of our instruments r_h, in Table 2 we give regression estimates and associated robust standard errors from a linear regression of our endogenous regressors, f_h and $\ln y_h$ on our 18 demographic variables z_h and our 4 instruments r_h (corresponding to what would be the first stage of two stage least squares if our model were linear). Standard errors are clustered at the village (i.e., the Upazila) level.

Table 2: " First Stage"							
		f, abuse			$\ln y$, ln(total consumption)		
		Est	*Std Err*	t	Est	*Std Err*	t
regressors	constant	0.203	*0.022*	9.04	0.035	*0.030*	1.18
	avg_age_men	-0.002	*0.008*	-0.28	0.000	*0.009*	-0.01
	avg_age_women	-0.040	*0.011*	-3.63	0.023	*0.011*	2.08
	avg_edu_men	-0.008	*0.003*	-2.6	0.031	*0.003*	9.96
	avg_edu_women	-0.005	*0.003*	-1.68	0.038	*0.003*	11.18
	children	-0.032	*0.026*	-1.25	0.125	*0.027*	4.67
	frac_girl	-0.026	*0.021*	-1.27	0.055	*0.019*	2.81
	ln_dowry	0.006	*0.003*	2.59	0.005	*0.003*	1.94
	ln_real_wealth	-0.008	*0.003*	-2.48	0.033	*0.006*	6.08
composition	m1_f1_c1	-0.045	*0.027*	-1.66	-0.108	*0.029*	-3.71
	m1_f1_c3	-0.013	*0.032*	-0.4	0.067	*0.029*	2.32
	m1_f1_c4	-0.034	*0.057*	-0.6	0.115	*0.041*	2.82
	m1_f2_c1	-0.101	*0.033*	-3.07	0.187	*0.035*	5.28
	m1_f2_c2	-0.035	*0.035*	-1.02	0.210	*0.035*	6.05
	m2_f1_c1	-0.033	*0.039*	-0.84	0.067	*0.036*	1.86
	m2_f1_c2	-0.024	*0.041*	-0.59	0.200	*0.037*	5.45
	m2_f2_c1	-0.072	*0.033*	-2.19	0.293	*0.042*	7.04
	m2_f2_c2	-0.002	*0.047*	-0.05	0.297	*0.039*	7.67
instruments	ln(income)	0.004	*0.011*	0.32	0.132	*0.014*	9.25
	ln(income)^2	-0.001	*0.002*	-0.79	0.023	*0.002*	10.25
	Building Material	-0.035	*0.019*	-1.86	-0.098	*0.025*	-3.89
	village-average f	0.666	*0.033*	20.02	-0.135	*0.048*	-2.82
R-squared		0.160			0.356		
F-stat		105			34		

Table 2 shows that the violence indicator f_h is difficult to predict, with an R^2 of just 0.16, but the instruments collectively appear strong, in that the F-statistic of the significance of the instruments (the log of income, its square, the indicator of mud or bamboo building material and village-average abuse $\overline{f}_h$) is 105. The low R^2 of this regression emphasizes the point that we can't (and don't try) to actually model the decision to commit violence; there are too many unobserved attributes of households that affect violence (and hence too many unobservables that affect the u_j functions in our behavioral model). All we need are sufficiently strong instruments, which our F-statistic indicates is the case.

Although we can't treat this regression as a formal model of abuse, it is still suggestive regarding covariates. The regression suggests that wealth is negatively correlated with abuse, and that both income and building material predict abuse, conditional on wealth. In particular, the thinner building materials are correlated with less abuse. However, this is not a very strong instrument; its t-statistic is 1.9. In contrast, village-level average abuse is a very strong instrument, with a t-statistic of 20. Some of the strength of the building materials instrument is masked by village-level average abuse. The t-statistic for building materials becomes 2.7 if village-level average abuse is dropped as an instrument. This correlation is not surprising; thin walls may be less of a deterent to abuse if abuse is widely accepted in one's neighborhood.

The household log budget $\ln y_h$ is fitted with an R^2 of 0.36 and an F-statistic of the instruments of 34. The log budget is highly correlated with our income instruments, but is also correlated with our building materials dummy and village-level average abuse.

Although the GMM estimator uses all instruments for all moments, we may think of the above results as empirically motivating the relevance of income for the endogenous budget and motivating the relevance of building materials and village-level abuse for the endogenous abuse indicator. For further reassurance that the instruments are valid, we later examine overidentification test statistics.

3.3 Model Specification

To estimate our model we divide household members into $J = 3$ types, indexed by j equaling m, f, or c, referring to adult males, adult females, and children, respectively. By equation (13), our estimator applies GMM to estimate the parameter vector θ using moments of the form $E\left(\varepsilon_{jh}\phi\left(r_h, z_h\right)\right) = 0$ where the errors ε_{jh} are given by

$$\varepsilon_{jh} = \frac{w_{jh}}{\eta_j\left(f_h, z_h, \theta\right)} - \gamma_j\left(z_h, \theta\right) - \beta\left(z_h, \theta\right)\left(\ln y_h - \ln N_{jh} + \ln \eta_j\left(f_h, z_h, \theta\right) + \ln \delta\left(f_h, z_h, \theta\right)\right). \tag{14}$$

The functions η_j, γ_j, δ and β are specified as

$$\eta_j\left(f_h, z_h, \theta\right) = k_{j0} + k_j' z_h + c_j f_h,$$

$$\gamma_j\left(z_h, \theta\right) = l_{j0} + l_j' z_h,$$

$$\ln \delta\left(f_h, z_h, \theta\right) = \left(a_0 + a_1' z_h\right) f_h,$$

and

$$\beta\left(z_h, \theta\right) = b_0 + b_1' z_h.$$

The vector θ is therefore defined as all the coefficients in $a_0, a_1, b_0, c_j, k_{j0}, k_j, l_0$, and l_j' for $j \in \{m, f, c\}$. Note the definition of δ enforces the restriction that $\ln \delta = 0$ when f_h is zero. To impose the constraint that resource shares sum to one, we impose $\sum_{j \in \{m,f,c\}} k_{j0} = 1$, $\sum_{j \in \{m,f,c\}} k_j = 0$, and $\sum_{j \in \{m,f,c\}} c_j = 0$.

Due to multicollinearity (which we empirically document later) in our baseline specification we take $a_1 = 0$. We are particularly interested in the estimates of c_j, which gives the response of the resource shares to abuse f_h, and the estimate of a_0 which gives the response of the household scale economies to f_h.

The 18 demographic variables comprising z_h were given in Tables 1a and 1b. As noted earlier, our instruments r_h are the continuous log-income measure, the indicator that the housing building material is either mud or bamboo, and the continuous village leave-out average for each household (summarized at the bottom of Table 1b).

Our moment equations (13) require a vector of functions $\phi\left(r_h, z_h\right)$. In theory, any vector of functions satisfying the rank condition for identification would suffice. Ideally, one wants to choose functions that highly correlate with the components of the model. Inspection of equation (12) shows that budget shares w_j are linear in $\eta_j(f_h, z_h)\beta(z_h) \ln y_h$ and $\eta_j(f_h, z_h)\beta(z_h) \ln \delta(f_h, z_h)$. This suggests that $(f_h, z_h) \times z_h$ times $\ln y_h$ and $(f_h, z_h) \times z_h$ times f_h would be informative. Since both $\ln y_h$ and f_h are endogenous, we create functions $\phi\left(r_h, z_h\right)$ that are analagous to these products, by replacing these endogenous variables with instruments. So, our functions $\phi\left(r_h, z_h\right)$ are

$$\phi\left(r_h, z_h\right) = \left(1, z_h, r_h\right) \times \left(1, z_h, r_h\right) \times \left(1, r_h\right),$$

where $\times$ indicates element-wise multiplication, deleting redundant elements. This yields a vector $\phi\left(r_h, z_h\right)$ with 601 elements (including the constant) if cubic terms are included and with 185 elements if they are excluded. Our baseline model has 111 parameters, so our model is highly overidentified (having far more moments than parameters). This can lead to small-sample efficiency issues that we will need to investigate due to imprecision in the estimation of the GMM weighting matrix. The estimated standard errors we report are clustered at the village level.

3.4 Model Estimates

Our main results are given in Tables 3 to 5. In these tables we focus on a subset of the most relevant coefficients. The full set of baseline model parameter estimates are reported after the Appendix in Table 6.

Identification of many of the parameters in the model requires $\beta \neq 0$ and hence either $b_0 \neq 0$ or $b_1 \neq 0$. As can be seen in Table 6, the estimates of b_0 and b_1 are statistically significantly different from zero. These estimated coefficients show that, as usual, food budget share Engel curves slope downwards.

Identification also requires exogeneity of the instrument vector $\phi(r, z)$. The bottom rows of Tables 3 to 5 present estimated J test statistics to assess this exogeneity restriction. In particular, the $J-$tests are tests of the hypothesis that the elements of $\phi(r, z)$ are all uncorrelated with the error ε_j. None have a p-value less than 0.05, so we fail to reject the null of instrument validity.

In Tables 3, 4, and 5 we present parameter estimates that are readily interpreted as applying to the reference household type z_0 (1 man, 1 woman and 2 children, with $z = 0$). In the first row in each of these tables, we provide estimates of a_0, which equals $\ln \delta(1, z_0, \theta)$ for the reference household, i.e., the response of log-efficiency to abuse (more precisely, the percent change in total budget y that would be equivalent to the loss in efficiency associated with abuse). The next rows provide k_{j0} and c_j for each type j in the household. These equal the reference household resource share without abuse, $\eta_j(0, z_0)$, and the change in resource share when abuse is present, $\eta_j(1, z) - \eta_j(0, z)$, respectively.

The next block of rows in Tables 3, 4, and 5 report, for each type j, the proportional difference in type j's shadow budget from having abuse present vs not. This is the effect of abuse on type j's money metric consumption utility. This affect has two components. First, the effect of abuse on efficiency is equivalent to a reduction in the household's budget from y to $\delta(1, z)y$. Second, abuse changes the resource share going to members of type j from $\eta_j(0, z)$ to $\eta_j(1, z)$. Together, these effects of abuse are equivalent to changing type j's shadow budget from $\eta_j(0,z)y$ to $\eta_j(1,z)\delta(1,z)y$. Expressed as a fraction of y, we define the proportionate change in the money metric as

$$\Delta \text{ money metric} = \eta_j(0,z) - \eta_j(1,z)\delta(1,z)$$

which, in our baseline model, equals $k_{j0} - (k_{j0} + c_j f)\exp(a_0)$. This is reported in the third block of rows in Tables 3, 4, and 5. Finally, as noted above, the bottom row of each of these tables gives tests of instrument validity based on overidentification of the model.

Table 3 presents results from estimation of our baseline model, with GMM standard errors clustered at the village level. In this baseline model, all demographic variables z are included in $\gamma_j(z)$, $\beta(z)$, and $\eta_j(f,z)$ but $a_1 = 0$ so that δ takes the simplest possible form, $\ln\delta(f,z,\theta) = a_0 f_h$. Table 3 has 3 blocks of columns, corresponding to estimates of the same baseline model but varying the instruments somewhat. For now we will focus on the first set of columns, which uses our complete set of instruments.

Table 3: GMM Estimates, Selected Coefficients, Varying Instruments

			(1) Baseline		(2) No Village_f		(3) No Cubed Insts	
function	person	variable	Estimate	*Std Err*	Estimate	*Std Err*	Estimate	*Std Err*
$\ln\delta$	all	constant, a_0	-0.0534	*0.0179*	-0.0609	*0.0351*	0.0088	*0.0666*
η	men	constant, k_{m0}	0.3476	*0.0053*	0.3232	*0.0080*	0.3205	*0.0147*
		f , c_m	0.0141	*0.0020*	0.0250	*0.0039*	0.0146	*0.0076*
	women	constant, k_{f0}	0.3047	*0.0048*	0.2989	*0.0076*	0.3548	*0.0168*
		f , c_f	-0.0086	*0.0018*	-0.0039	*0.0035*	-0.0123	*0.0072*
	children	constant, k_{c0}	0.3477	*0.0067*	0.3779	*0.0104*	0.3248	*0.0208*
		f , c_c	-0.0054	*0.0022*	-0.0211	*0.0045*	-0.0024	*0.0058*
$\triangle$money metric	men	m1_f1_c2	-0.0047	*0.0064*	0.0044	*0.0122*	0.0176	*0.0228*
	women	m1_f1_c2	-0.0240	*0.0056*	-0.0214	*0.0101*	-0.0092	*0.0261*
	children	m1_f1_c2	-0.0232	*0.0060*	-0.0422	*0.0127*	0.0005	*0.0213*
Number of Observations			2866		2866		2866	
J-statistic: value [df] p-value			1741 [1692]	0.1988	1198 [1137]	0.1019	484 [444]	0.0924

The top cell of column (1) in Table 3 gives the estimate of a_0 as -0.0534, suggesting that abuse reduces efficiency by an amount equivalent to reducing the household's total expenditures budget y by roughly 5 per cent.

The next block of column (1) gives estimates of reference household resource shares. These estimates suggest that the man gets 35 per cent of household resources, the woman gets 30 per cent, and the two children split the remaining 35 per cent. These estimates are similar to what Dunbar, Lewbel and Pendakur (2013) found in poor households in Malawi and to what Brown, Calvi and Penglase (2018) find in Bangladesh. The estimated values of c_j in this block gives the marginal effects of abuse on resource shares. These show that abuse increases men's resource shares by 1.41 percentage points, and lowers women's and children's shares by 0.86 and 0.54 percentage points, respectively. Although these estimated effects on resource shares are small, they're estimated very precisely, with z statistics of 7.1, 4.7 and 2.5 for men, women and children, respectively.

The third rows of estimates gives the net effect of abuse on the shadow income (money metric utility) of each household member type j. For men, the increase in their resource share from committing abuse is offset by the decrease in household's efficiency, resulting in a near zero (and statistically insignificant) -0.47 per cent change. For women, the estimate is -2.40 per cent, and for children it's -2.32 per cent, and both are statistically significant (with z- statistics around -4). This loss in money metric utility for women and children comes from both channels. About one third of the decline is from the loss of resource shares (which are gained by the men), while two thirds comes from the loss in consumption efficiency.

The middle column of Table 3 drops village-level average abuse as an instrument, since its validity depends on some additional assumptions as described in the previous section. Doing so greatly reduces estimation precision (standard errors roughly double), but the estimates show the same pattern of a large loss of efficiency, a gain in resource share for men yielding a net near zero effect on men's shadow income (this time a plus instead of a minus half a percent, but still insignificant), and substantial corresponding losses for both women and children. The overidentification tests presented in the bottom row of the Table show that we do not reject exogeneity of all instruments in either the baseline specification or in this one.

In finite samples, estimation of the GMM weighting matrix can be poor when the model has far more moments than parameters. To assess if this is an issue, in the third set of columns in Table 3 we omit the cubic functions of our instruments, which as noted in the previous section greatly decreases our number of estimating moments from 601 to 185. The instrument here are $\phi(r_h, z_h) = (1, z_h, r_h) \times (1, z_h, r_h)$, deleting redundant elements. Inspection of the standard errors in column (3) suggest that this is very costly in terms of precision: standard errors are roughly quadruple those in our baseline specification. The parameter a_0 now become insignificant, but we still see the same pattern of effects on resource shares.

In Table 4, we consider three alternatives regarding data construction in our baseline model. The sets of columns here are labeled (4) to (6), since Table 3 had estimates labeled (1) to (3). In the leftmost column, labeled (4), we use a simpler food definition than in the baseline. In the baseline, individual one-day quantity data are used to divide weekly food spending into individual weekly consumption levels. Here, we instead take each individual's share of one-day quantities and multiply by 365 to get annual quantities, and then convert this to individual spending via multiplication by prices. This method is simpler, but is more likely to suffer from measurement error due to the greater variability of daily intakes in comparison to weekly spending. The main difference with the baseline model is a higher estimate of the inefficiency due to abuse: 7.29 per cent vs 5.34 per cent.

Table 4: GMM Estimates, Selected Coefficients, Varying Data

			(4) Simpler Food Def		(5) Keep Zeroes		(6) Nuclear Families	
function	person	variable	Estimate	*Std Err*	Estimate	*Std Err*	Estimate	*Std Err*
$\ln \delta$	all	constant, a_0	-0.0729	*0.0196*	-0.0500	*0.0231*	-0.0730	*0.0237*
η	men	constant, k_{m0}	0.3287	*0.0052*	0.3378	*0.0061*	0.3437	*0.0040*
		f , c_m	0.0124	*0.0019*	0.0157	*0.0025*	0.0137	*0.0021*
	women	constant, k_{f0}	0.2993	*0.0040*	0.2995	*0.0060*	0.2904	*0.0042*
		f , c_f	-0.0077	*0.0018*	-0.0070	*0.0020*	-0.0065	*0.0024*
	children	constant, k_{c0}	0.3720	*0.0055*	0.3627	*0.0079*	0.3659	*0.0055*
		f , c_c	-0.0047	*0.0021*	-0.0087	*0.0028*	-0.0072	*0.0031*
$\triangle$money	men	m1_f1_c2	-0.0115	*0.0066*	-0.0015	*0.0082*	-0.0115	*0.0081*
metric	women	m1_f1_c2	-0.0282	*0.0058*	-0.0213	*0.0069*	-0.0265	*0.0067*
	children	m1_f1_c2	-0.0305	*0.0068*	-0.0260	*0.0079*	-0.0325	*0.0085*
Number of Observations			2866		3087		1630	
J-statistic: value [df] p-value			1687 [1692]	0.5297	1771 [1692]	0.0887	954 [990]	0.7893

In our baseline model, we drop 221 households (7% of observations) that report zero food intake in the one-day diary for any household member. In column (5) of Table 4, we retain these observations, and assign a food share of zero where the observed one-day food intake is zero. This is simpler and induces less potential selection bias, but it includes households where we know there must be measurement error in the annual individual-level food consumption. The resulting estimates differ little from our baseline, though with slightly larger standard errors.

The nuclear households in our data have 1 adult man and 1 adult woman and one to four children. We also have 1236 non-nuclear households, having either more than 1 adult man or more than 1 adult woman. Our model and data may be less appropriate for these non-nuclear households, both because multiple adult men might or might not coordinate on abuse and cooperation effort, and because abuse is only reported by "the main" adult female in the household. Column (6) in Table 4 reports result from estimating the model just with nuclear households, at a cost of considerable loss in sample size. The main difference with the baseline model is a higher estimate of the inefficiency due to abuse of 7.3 per cent.

Table 5: GMM Estimates, Selected Coefficients, Varying delta

			(7) $\delta = 1$		(8) δ(hhsize)		(9) δ(all)	
function	person	variable	Est	*Std Err*	Est	*Std Err*	Est	*Std Err*
$\ln \delta$	all	constant			-0.0720	*0.0214*	-0.0866	*0.0332*
		ln(HHsize/4)			0.1903	*0.0834*	0.1885	*0.0948*
η	men	constant	0.3473	*0.0052*	0.3471	*0.0053*	0.3473	*0.0053*
		f	0.0143	*0.0020*	0.0146	*0.0020*	0.0142	*0.0020*
	women	constant	0.3047	*0.0048*	0.3051	*0.0049*	0.3049	*0.0051*
		f	-0.0080	*0.0018*	-0.0087	*0.0019*	-0.0097	*0.0020*
	children	constant	0.3480	*0.0066*	0.3479	*0.0067*	0.3478	*0.0070*
		f	-0.0063	*0.0022*	-0.0059	*0.0022*	-0.0045	*0.0024*
$\triangle$money metric	men	m1_f1_c2	0.0143	*0.0020*	-0.0105	*0.0073*	-0.0158	*0.0109*
	women	m1_f1_c2	-0.0080	*0.0018*	-0.0293	*0.0064*	-0.0342	*0.0092*
	children	m1_f1_c2	-0.0063	*0.0022*	-0.0297	*0.0071*	-0.0330	*0.0110*
	men	m1_f1_c4	0.0143	*0.0020*	0.0161	*0.0086*	0.0113	*0.0122*
	women	m1_f1_c4	-0.0080	*0.0018*	-0.0074	*0.0074*	-0.0122	*0.0105*
	children	m1_f1_c4	-0.0063	*0.0022*	-0.0035	*0.0133*	-0.0092	*0.0196*
Number of Observations			2866		2866		2866	
J-statistic: value [df] p-value			1742 [1693]	0.1989	1739 [1691]	0.2035	1735 [1681]	0.1754

The function δ, which gives the percentage cost of inefficiency associated with abuse, is a novel feature of our model. In Table 5, we consider alternative specifications for this cost of inefficiency function. The leftmost block of Table 5, column (7), imposes the restriction $a_0 = a_1 = 0$, which makes $\ln \delta = 0$. This specification imposes the constraint that abuse does not affect efficiency. Column (8) allows the diseconomies of scale associated with abuse to vary by household size. In this specification, a_1 is a scalar which multiplies the log of household size divided by 4 (the size of the reference household type). Finally, in the third block, We let a_1 be a vector of coefficients of household size, an indicator that the household has 2 men, an indicator that it has 2 women, and all of elements of z except the household composition dummies. In these specifications, the coefficient a_0 gives the effect of domestic abuse on the value of household scale economies for the reference household type.

Consider first estimates (7) where we don't allow for inefficiency. Compared to our baseline specification (estimates (1)), the estimated values of the constant terms in resource shares are virtually identical. Similarly, the estimated marginal effect of abuse is the roughly the same in these two specifications. This suggests that leaving out the inefficiency channel does not bias our estimates of the levels of resource shares, or the response of resource shares to domestic abuse.

In estimates (8), we allow the inefficiency function to depend on the log of household size. For households with 4 members, this model makes $\ln \delta \left(f_h, z_h, \theta\right) = a_0 f_h$, and for larger households the marginal effect of a proportionate increase in household size is given by the scalar a_1. It is plausible to think that scale economies are larger for larger households, so there is more at stake when we consider efficiency loss for these households. Efficiency loss due to domestic abuse could therefore be greater in large households than in small households because non-cooperation costs more. Alternatively, the fact that there is more at stake for large households could manifest as reduced efficiency loss in large households: members of larger households might find a way to cooperate even in the presence of abuse because the costs of non-cooperation are so much higher.

The estimated value of a_0 in column (8) indicates that the reference household type (nuclear family with 4 members) would face an efficiency loss of 7.2 per cent in the presence of domestic abuse, while the estimate of the scalar a_1 is 0.19, implying that efficiency loss is smaller in larger households. For the largest households in our sample, which have 6 members, the predicted efficiency loss due to domestic abuse is 0.52 per cent, which is very close to (and statistically indistinguishable from) zero.

In estimates (9), we see a similar pattern to estimates (8) (and to our baseline estimates (1)), but with larger standard errors due to the inclusion of many additional parameters.

The bottom panel of Table 5 gives estimates of the change in the money metric of consumption utility for each type of person in response to domestic abuse. The upper rows give an estimate of this welfare loss of people in the reference household type; the lower row give an estimate of this welfare loss for the largest households (nuclear households with 4 children). The estimates of the proportionate changes in money metric utility due to domestic abuse for the reference household are similar to those reported in the baseline, with men losing roughly 1 percent (insignificantly different from zero) and women and children losing roughly 3 percent (statistically significantly negative). However, they are a little larger in magnitude than those reported in the baseline. For people living in the largest households, efficiency loss is close to zero, so the changes in resource shares dominate the changes in money metric utility. Thus, in these households, men gain 1.6 per cent (marginally statistically significant) and women and children lose (each loss is insignificant, but their total loss is marginally statistically significant).

We have three main bottom line empirical results. First, our instruments for the effects of abuse on consumption and for endogeneity of total expenditures are empirically relevant and pass overidentification tests for exogeneity. Second, we find that domestic abuse is a cooperation factor, i.e., it does affect the efficiency of household consumption. We find losses due to decreased sharing and cooperation on the order of 0 to 7 per cent of the household's total budget in households where domestic abuse is present. Third, we find that domestic abuse affects resource shares, with about 1.5 percent more of total household expenditures going towards men (and away from women and children, roughly equally) in households with abuse. The net effect of these shifts is that domestic abuse has little or no affect the money-metric utility from consumption for men, but reduces the money-metric utility from consumption of women and children by roughly 2.5 per cent for women and children. Note that these numbers only measure the impacts of abuse on consumption, and do not include, e.g., direct pain or sorrow from experiencing abuse.

4 Conclusions

We provide a general framework for analyzing the effects of what we call "cooperation factors" on collective household behavior. A cooperation factor is any variable that can 1. induce inefficiency in consumption by reducing cooperation and sharing, 2. affect resource shares like a distribution factor, and 3. directly affect the utility of household members (additively separable from consumption).

Examples of cooperation factors could be direct indicators cooperation (e.g., measures of time spent together on household chores), or more generally be behaviors that correlate cooperation or failures to cooperate, such as money hiding or domestic abuse. Still other examples could variables relating to private ownership of durables, or almost any variable that in the previous literature was considered to be distribution factor.

A common objection to the application of collective household models, particularly in developing countries, is that most such models assume households are Pareto efficient, while behavior like domestic abuse is evidence of inefficiency. A convenient feature of cooperation factors is that they allow for inefficiency while still maintaining the modeling advantages of efficient collectives.

We take our general cooperation factors model, simplify it to reduce data requirements, and apply it to household survey data from Bangladesh. Our empirical estimates are that abuse reduces household consumption efficiency by an amount roughly equal to a 5% reduction in total household resources. Commiting abuse also increases the man's share of the household's resources from about 35% to 36.5% (with that gain coming roughly equally out of the woman's and children's shares).

The net effect of commiting abuse on men's resources is close to zero, as their gains from increased resource shares are offset by their losses in household efficiency. So why do these men engage in abuse, since it makes the women and children worse off while having almost no net effect on the men's resources? Our model does not answer this question, but two answers are consistent with our model. One is that they may directly derive utility from commiting abuse (e.g., by feeling powerful). Another possibility is that cooperation and efficiency in consumption requires effort, and these men may get utility from avoiding this effort.

One other takeaway from these empirical results is that, while domestic abuse can cause considerable suffering and damage, the effects of this abuse on the consumption behavior of households (while statistically significant), are not very large in magnitude, and omitting these effects from the model do not greatly change estimated resource shares.

So, while estimating the impacts of violence on household behavior remains very important, these empirical results suggest that past studies that estimated resource shares (and associated individual specific poverty rates) while ignoring these inefficiencies, like DLP or Calvi (2019), would not have had their conclusions change much if they had been able to observe and control for domestic abuse. Future work should include investigating the impacts of abuse and other potential sources of inefficiency (i.e., other potential cooperation factors), to see if they have larger effects in other countries or among particular subpopulations.

5 Appendix A: Formal Assumptions and Proofs

Here we formally derive our model, and prove that it is identified. To simplify the derivations and assumptions, we first prove results without unobserved random utility parameters (as would apply if, e.g., our data consisted of many observations of a single household, or of many households with no unobserved variation in tastes). We then later add unobserved error terms to the model.

Let f, r, y, p, π, and z be as defined in the main text (recalling that r contains all of the elements of both $\widetilde{r}$ and v). Note that the first few Lemmas below will not impose the restriction that f only equal two values.

ASSUMPTION A1: Conditional on f, r, y, p, π, and z, the household chooses quantities to consume using the program given by equation (9).

Assumption A1 describes the collective household's conditionally efficient behavior. For each household member j, U_j is that member's utility function over consumption goods, u_j is that members additional utility or disutility associated with f, and ω_j is that member's Pareto weight.

As can be seen by equation (9), the way that private assignable goods q_j differ from other goods g is that each q_j only appears in the utility function of individual j (which makes it assignable to that member) and these goods are unaffected by the matrix A_f in the budget constraints, meaning that they are not shared or consumed jointly (which makes them private goods).

We next assume some regularity conditions. These assumptions ensure sensible and convenient restrictions on economic behavior like no money illusion, preferring larger consumption bundles to smaller ones, and the absence of corner solutions in the household's maximization problem.

ASSUMPTION A2: Each $\omega_j(f, z, p, \pi, y)$ function is differentiable and homogeneous of degree zero in p, π, and y. Each $U_j(q_j, g_j, z)$ function is concave, strictly increasing, and twice continously differentiable in g_j and q_j. For each f, the matrix A_f is nonsingular with all nonnegative elements and a strictly positive diagonal. The variable y and each element of p and π are all strictly positive, and the maximizing values of $g_1, q_1, ...g_J, q_J$ in Assumption A1 are all strictly positive.

LEMMA 1 Let Assumptions A1 and A2 hold. Then there exist positive resource share functions $\eta_j(p, \pi, y, f, z)$ such that $\sum_{j=1}^{J} \eta_j(p, \pi, y, f, z) = 1$, and the household's demand function for goods is given by each member j solving the program

$$\max_{g_j, q_j} U_j(q_j, g_j, z) \tag{15}$$

$$\text{such that } \; p' A_f g_j + \pi_j q_j = \sum_{j=1}^{J} \eta_j(p, \pi, y, z, f)\, y \; \text{ and } g = A_f \sum_{j=1}^{J} g_j.$$

To prove Lemma 1, first observe that the values of $g_1, q_1, ...g_J, q_J$ that maximize equation (9) are equivalent to the values that maximize

$$\max_{g_1, q_1, ...g_J, q_J} \sum_{j=1}^{J} U_j(q_j, g_j, z)\, \omega_j(p, \pi, y, f) \tag{16}$$

given the same budget constraint. because the terms in equation (9) that are not in (16) do not depend on $g_1, q_1, ...g_J, q_J$. With that replacement, the proof of Lemma 1 then follows immediately from the results derived in BCL. BCL only considered $J = 2$, but the extension of this Lemma to more than two household members, and to carrying the additional covariates, is straightforward. Note that the resource share functions η_j in Lemma 1 do not depend on r, because r, including the component v, does not appear in either equation (16) or in the budget constraint, and so cannot affect the outcome quantities.

Our empirical work will make use of cross section data, where price variation is not observed. Most of the remaining assumptions we make about resource shares and about the U_j component of utility are the same, or similar, to those made by DLP, and for the same reason: to ensure identification of the model, and without requiring price variation.

ASSUMPTION A3. The resource share functions $\eta_j(p, \pi, y, f, z)$ do not depend on y.

DLP give many arguments, both theoretical and empirical, supporting the assumption that resource shares do not vary with y. Given Assumption A3, we hereafter write the resource share function as $\eta_j(\pi, p, f, z)$.

For the next assumption, recall that an indirect utility function is defined as the function of prices and the budget that is obtained when one substitutes an individual's demand functions into their direct utility function.

ASSUMPTION A4. For each household member j, the direct utility function $U_j(g_j, q_j, z)$, when facing prices p and π and having the budget y, has the associated indirect utility function

$$V_j(\pi_j, p, y, z) = [\ln y - \ln S_j(\pi_j, p, z)] M_j(\pi_j, p, z) \tag{17}$$

For some functions S_j and M_j.

Assumption A4 says that household members each have utility functions in the class that Muellbauer (1974) called PIGLOG (price independent, generalized logarithmic) preferences. As noted in the main text, this is a class of functional forms that is widely known to fit empirical continuous consumer demand data well. Examples of popular models in this class include the Christensen, Jorgenson, and Lau (1975) Translog demand system and Deaton and Muellbauer's (1980) AIDS (Almost Ideal Demand System) model.[10]

LEMMA 2: Let Assumptions A1, A2, A3, and A4 hold. Then the value of $U_j(q_j, g_j, z)$ attained by household member j is given by

$$U_j = [\ln \eta_j(\pi, A_f p, f, z) + \ln y - \ln S_j(\pi_j, A_f p, z)] M_j(\pi_j, A_f p, z) \quad (18)$$

and the household's demand functions for the private assignable goods q_j are

$$q_j = \eta_j(\pi, A_f p, f, z) y \left(\frac{\partial \ln S_j(\pi_j, A_f p, z)}{\partial \pi_j} - \frac{\partial \ln M_j(\pi_j, A_f p, z)}{\partial \pi_j} (\ln(\eta_j(\pi, A_f p, f, z) y) - \ln S_j(\pi_j, A_f p, z)) \right) \quad (19)$$

To prove Lemma 2, observe that by Lemma 1, household member j maximizes the utility function $U_j(q_j, g_j, z)$ facing shadow prices $A_f' p$ and π_j and having the shadow budget $\eta_j(\pi, A_f p, f, z) y$. Therefore, using the definition of indirect utility, member j's attained utility level $U_j(q_j, g_j, z)$ is given by $V_j(\pi_j, A_f' p, \eta_j(\pi, A_f p, f) y)$, which by Assumption A4 equals equation (18). Next, a property of regular indirect utility functions is that the corresponding demand functions can be obtained by Roy's identity. Equation (19) is obtained by applying Roy's identity to equation (17) for the private assignable goods q_j, and then replacing p and y in the result with $A_f' p$ and $\eta_j(\pi, A_f p, f) y$.

We could similarly obtain the demand functions for other goods g, as in BCL, but these will be more complicated due to the sharing, with Roy's identity being applied to each member to obtain each g_j demand function, and substituting the results into $g = A_f \sum_{j=1}^{J} g_j$. However, our empirical analyses will only make use of the private assignable goods q_j with demands given by equation (19).

ASSUMPTION A5. Let $\ln M_j(\pi_j, A_f p, z) = m_j(A_f p, z) - \beta(z) \ln \pi_j$ for some functions m_j and β.

There are two restrictions embodied in Assumption A5. One is that the functional form

[10] Most more recent alternatives, like so-called " rank three" demand systems, are used for data from countries where the distribution of y is large, and more complicated budget responses are needed to capture behavior at both low and high income levels. Other popular demand models, like the multinomial logit based models widely used in the industrial organization literature, are designed for use with discrete demand data and are unsuitable for the type of continuous consumer demand data we analyze here.

of $\ln M_j$ in terms of prices is linear and additive in $\ln \pi_j$, and the other is that the function $\beta(z)$ does not vary by j. The functional form restriction of log linearity in log prices is a common one in consumer demand models, e.g., the function M_j in Deaton and Muellbauer's (1980) AIDS (Almost Ideal Demand System) satisfies this restriction. Assumption A5 could be further relaxed by letting β depend on p (though not on A_f) without affecting later results.

To identify their model, DLP define and use a property of preferences called similarity across people (SAP), and provide empirical evidence in support of SAP. The restriction that β not vary by j suffices to make SAP hold for the private assignable goods (but not necessarily for other goods).

ASSUMPTION A6. Let $\ln S_j(\pi_j, A_f p, z) = \ln s_j(\pi_j, p, z) - \ln \delta(A_f p, z)$ for some functions s_j and δ. Without loss of generality, let $\ln \delta(A_0 p, z) = 0$.

Assumption A6 assumes separability of the effects of π_j and f on the function S_j. Assuming $\ln \delta(A_0 p, z) = 0$ in Assumption A6 is without loss of generality, because if it does not hold then one can make it hold if one redefines δ and s_j by subtracting $\ln \delta(A_0 p, z)$ from both $\ln \delta(f, p, z)$ and $\ln s_j(\pi_j, p, z)$. DLP discuss various ways in which the matrix A_f can drop out of a function of prices, as required in the function s_j.[11] This assumption is not vital, but will be helpful for making the cost of an inefficient choice of f identifiable

It will be convenient to express our demand functions in budget share form. Define $w_j = q_j \pi_j / y$. This budget share is the fraction of the household's budget y that is spent on buying person j's assignable good q_j.

LEMMA 3: Given Assumptions A1 to A6, the value of $U_j(q_j, g_j, z)$ attained by household

[11] For example, one way A_f drops out is if A_f is block diagonal, with one block that does not vary by f, and with s_j only depending on π_j and the prices in that block. Alternatively, linear constraints could be imposed on the elements of A_f, with s_j depending only on the corresponding functions of prices, that, by these constraints, do not vary with A_f. Analogous restrictions are oftern imposed on demand systems. For example as shown in Lewbel (1991), the Translog demand system as implemented by Jorgenson, and Slesnick (1987) imposes a linear constraint on its Barten (1964) scales, that results in a restriction like this on its equivalence scales. Note that BCL refer to the diagonal elements of A_f as Barten technology parameters, due to their equivalent to Barten scales.

member j is given by

$$\left[\ln \eta_j \left(\pi, A_f p, f, z\right) + \ln y - \ln s_j \left(\pi_j, p, z\right) + \ln \delta \left(A_f p, z\right)\right] \left[m_j \left(A_f p, z\right) - \beta \left(z\right) \ln \pi_j\right] \tag{20}$$

and the budget share demand functions for each private assignable good are given by

$$w_j = \eta_j \left(\pi, A_f p, f, z\right) \left[\gamma_j \left(\pi_j, p, z\right) + \beta \left(z\right) \left(\ln y + \ln \eta_j \left(\pi, A_f p, f, z\right) + \ln \delta \left(A_f p, z\right)\right)\right]. \tag{21}$$

where the function γ_j is defined by

$$\gamma_j \left(\pi_j, p, z\right) = \frac{\partial \ln s_j \left(\pi_j, p, z\right)}{\partial \ln \pi_j} - \beta \left(z\right) \ln s_j \left(\pi_j, p, z\right)$$

The proof of Lemma 3 consists of substituting the expressions for M_j and S_j given by Assumptions A5 and A6 into the equations given by Lemma 2, and converting the quantity q_j into the budget share w_j.

ASSUMPTION A7. Market prices p and π are the same for all households.

Our data come from a single time period, which (assuming the law of one price) justifies assuming p and π are the same across all households. This assumption makes our demand functions reduce to Engel curves. For simplicity, we abuse notation here and redefine objects that were functions of $A_f p$ as just functions of f, since with fixed prices the only source of variation of $A_f p$ is just variation in f).

LEMMA 4: Given Assumptions A1 to A7, the value of $U_j \left(q_j, g_j, z\right)$ attained by household member j is given by

$$\left[\ln \eta_j \left(f, z\right) + \ln y - \ln s_j \left(z\right) + \ln \delta \left(f, z\right)\right] M_j \left(f, z\right) \tag{22}$$

and the budget share Engel curve functions $w_j = W_j \left(f, z, y\right)$ for each private assignable good are given by

$$W_j(f,z,y) = \eta_j(f,z)\left[\gamma_j(z) + \beta(z)\left(\ln y + \ln \eta_j(f,z) + \ln \delta(f,z)\right)\right]. \qquad (23)$$

Lemma 4 entails a small abuse of notation, where we haved absorbed the values of p and π into the definitions of all of our functions, noting that any function of $A_f p$ remains a function of f even if p is a constant. Lemma 4 is just rewriting Lemma 3 after dropping the prices.

LEMMA 5: Let Assumptions A1 to A7 hold. Let $W_j(f,z,y)$ be defined by equation (23) for $j = 1, ..., J$. Given functions $W_j(f,z,y)$, the functions $\eta_j(f,z)$, $\delta(f,z)$,$\gamma_j(z)$, and $\beta(z)$ are identified.

To prove Lemma 5, observe first by equation (23) that $\eta_j(f,z)\beta(z) = \partial W_j(f,z,y)/\partial \ln y$. Next, since resource shares sum to one, we can identify $\beta(z)$ and $\eta_j(f,z)$ by

$$\beta(z) = \sum\nolimits_{j=1}^{J} \frac{\partial W_j(f,z,y)}{\partial \ln y} \quad \text{and} \quad \eta_j(f,z) = \frac{1}{\beta(z)} \frac{\partial W_j(f,z,y)}{\partial \ln y}$$

Next, define $\rho_j(f,z,y)$ by

$$\rho_j(f,z,y) = \frac{W_j(f,z,y)}{\eta_j(f,z)} - \beta(z)\left(\ln y + \ln \eta_j(f,z)\right)$$

The function $\rho_j(f,z,y)$ is identified because it is defined entirely in terms of identified functions. By equation (23), $\rho_j(f,z,y) = \gamma_j(z) - \beta(z)\ln \delta(f,z)$. It follows from Assumption A6 that $\ln \delta(0,z) = 0$, so $\gamma_j(z)$ and $\delta(f,z)$ are identified by

$$\gamma_j(z) = \rho_j(0,z,y) \quad \text{and} \quad \ln \delta(f,z) = \frac{\rho_j(f,z,y) - \rho_j(0,z,y)}{\beta(z)}$$

evaluated at any value of y (or, e.g., averaged over y).

Lemma 5 shows that, given the household demand functions, the resource share functions $\eta_j(f,z)$ are identified, so our model, like DLP, overcomes the problem in the earlier collective household literature of (the levels of) resource shares not being identified. Lemma 5 also shows identification of the preference related functions $\gamma_j(z)$ and $\beta(z)$, and identification of our new cost of inefficiency function $\delta(f,z)$.

Define the function $R(f,z,v,y)$ by

$$R(f,y,v,z) = (\ln \eta_1(f,z) + \ln y - \ln s_1(z) + \ln \delta(f,z)) M_1(f,z) + u_1(f,v,z)$$

LEMMA 6: Let Assumptions A1 to A7 hold. Assume f is chosen by member 1 to maximize his own attained utility level. Then f is given by $f = \arg\max R(f,y,v,z)$.

Member 1's utility level is $U_1 + u_1$. The proof of Lemma 6 is then that, by equation (22) and the definition of u_1, for any f the level of $U_1 + u_1$ attained by member 1 is given by the function $R(f,y,v,z)$.

The above analyses apply to a single household. Our data will actually consist of a cross section of households, each only observed once. To allow for unobserved variation in tastes across households in a conveniently tractible form, replace the function $\ln S_j(\pi_j, A_f p, z)$ with $\ln S_j(\pi_j, A_f p, z) - \widetilde{\varepsilon}_j$ where $\widetilde{\varepsilon}_j$ is a random utility parameter representing unobserved variation in preferences for goods. This means that that $\widetilde{\varepsilon}_j$ appears in member j's utility function U_j. We assume these taste parameters vary randomly across households, so $E(\widetilde{\varepsilon}_j \mid r, z) = 0$. Similarly, replace $u_j(f,r,z)$ with $u_j(f,r,z) + \widetilde{e}_{jf}$ where $\widetilde{e}_{jf}$ represents variation in the utility or disutility associated with the choice of f. The errors $\widetilde{e}_{jf}$ and $\widetilde{\varepsilon}_j$ can be correlated with each other and across household members.

Substituting these definitions into the above equations, we get

$$w_j = \eta_j(f,z)\left[\gamma_j(z) + \beta(z)(\ln y + \ln \eta_j(f,z) + \ln \delta(f,z)) + \varepsilon_j\right] \tag{24}$$

where $\varepsilon_j = \beta(z)\widetilde{\varepsilon}_j$ so $E(\varepsilon_j \mid r, z) = 0$, and

$$f = \arg\max\left[R(f,y,r,z) + (M_1(f,z)/\beta(z))\varepsilon_1 + \widetilde{e}_{1f}\right] \tag{25}$$

We will want to estimate the Engel curve equations (24) for $j = 1, ..., J$. Equation (25) shows that f is an endogenous regressor in these equations, because f depends on both ε_1 and $\widetilde{e}_1$.

As discussed in the main text, we do not try to empirically identify or estimate equation (25), because both the function R and errors $\widetilde{e}_{1f}$ depend on u_1, and there are likely to be many important determinents of u_1 that we cannot observe.

Another source of error in our model is that, in our data, y is a constructed variable (including imputations from home production), and so may suffer from measurement error. We will therefore require instruments for y. Our current collective household model is static. This is justified by a standard two stage budgeting (time separability) assumption, in which households first decide how much of their income and assets to save versus how much to spend in each time period, and then allocate their expenditures to the various goods they purchase. The total they spend in the time period is y, and the household's allocation of y to the goods they purchase is given by equation (9). These means that variables associated with household income and wealth will correlate with y and so are potential instruments for y.

This time separability applies to the utility functions over goods, $U_j(q_j, g_j, z)$ for each member j, but need not apply to the utility or disutility associated with f, that is, $u_j(f, v, z)$. So at least some of these income and wealth variables could be components of v (e.g., the wall thickness measure discussed earlier as a component of v is likely to be correlated with income or wealth). Let $\widetilde{r}$ denote the vector of additional potential instruments for y (that is, measures related to income or wealth that are not already included in v).

Assume there exists values v_0 and v_1 such that $u_1(f, v_0, z) \neq u_1(f, v_1, z)$. Then it follows from equation (25) that f varies with v, so v can serve as an instrument for f. Similarly, assume that $\ln y$ correlates with $\widetilde{r}$, which can serve as instruments for $\ln y$ (elements of v could also be instruments for y). Based on equation (24), we have conditional moments

$$E\left[\left(\frac{w_j}{\eta_j(f,z)} - \gamma_j(z) - \beta(z)(\ln y + \ln \eta_j(f,z) + \ln \delta(f,z))\right) \mid \widetilde{r}, v, z\right] = 0 \qquad (26)$$

Later in this Appendix we consider nonparametric identification based on these moments, but for now consider using them parametrically. If we parameterize each of the unknown functions using a parameter vector θ, then equation (26) implies unconditional moments

$$E\left[\left(\frac{w_j}{\eta_j(f,z,\theta)} - \gamma_j(z,\theta) - \beta(z,\theta)(\ln y + \ln \eta_j(f,z,\theta) + \ln \delta(f,z,\theta))\right)\phi(\widetilde{r},v,z)\right] = 0 \tag{27}$$

for any suitably bounded functions $\phi(\widetilde{r},v,z)$. Our actual estimator will consist of parameterizing the unknown functions in this expression, choosing a set of functions $\phi(\widetilde{r},v,z)$, and estimating the parameters by GMM (the generalized method of moments) based on these moments. At the end of this Appendix we discuss choice of the ϕ functions.

Equation (27) can suffice for parametric identification and estimation, but is it still possible to nonparametrically identify the functions in this model in the presence of unobserved heterogeneity? The following Theorem shows that the answer is yes, if we make some additional assumptions. Theorem 1 shows these additional assumptions are sufficient for nonparametric identification of these functions, These additional assumptions, which are not required for parametric identification, are listed in Assumption A8.

ASSUMPTION A8. Replace the function $\ln S_j(\pi_j, A_f p, z)$ with $\ln S_j(\pi_j, A_f p, z) - \widetilde{\varepsilon}_j$, and replace $u_j(f,v,z)$ with $u_j(f,v,z) + \widetilde{e}_{jf}$. Let $\widetilde{e}_1 = \widetilde{e}_{11} - \widetilde{e}_{10}$. Define $\widetilde{y}(\widetilde{r},v,z)$ by $\ln \widetilde{y}(\widetilde{r},v,z) = E(\ln y \mid \widetilde{r},v,z)$. Assume the following: The function $\widetilde{y}(\widetilde{r},v,z)$ is differentiable in a scalar $\widetilde{r}$ with a nonzero derivative. The error $\widetilde{e}_1$ is independent of $y, \widetilde{r}, v, z$ and $(\varepsilon_j, \widetilde{e}_1)$ is independent of $\widetilde{r}$ conditional on (v,z). $E(\varepsilon_j \mid \widetilde{r},v,z) = 0$. The function $M_1(f,z)$ does not depend on f, where f is a binary variable chosen by member 1 to maximize his own attained utility level. There exist values v_1 and v_0 of v such that $u_1(f,v_1,z) \neq u_j(f,v_0,z)$.

THEOREM 1: Let Assumptions A1 to A8 hold. Then the functions $\eta_j(f,z)$, $\delta(f,z)$,$\gamma_j(z)$, and $\beta(z)$ are identified.

To prove Theorem 1, first observe that, with f binary, it follows from equation (25) that $f = 1$ if $R(1,y,v,z) + M_1(1,z)\widetilde{\varepsilon}_1 + \mu_j(1,v,z) + \widetilde{e}_{11}$ is greater than $R(0,y,v,z) + M_1(0,z)\widetilde{\varepsilon}_1 + \mu_j(0,v,z) + \widetilde{e}_{10}$. Taking the difference in these expressions, and using the assumption that

$M_1(f,z)$ doesn't depend on f, we get that $f=1$ if and only if

$$(\ln\eta_1(1,z)+\ln\delta(1,z))\,M_1(z)+\mu_j(1,v,z)-(\ln\eta_1(0,z)+\ln\delta(0,z))\,M_1(z)-\mu_j(0,v,z)+\widetilde{e}_1$$

is positive. This means that $f=\widetilde{f}(v,z,\widetilde{e}_1)$ for some function $\widetilde{f}$.

Now, again exploiting that f is binary,

$$E(w_j\mid\widetilde{r},v,z,y)=E\left[W_j(f,z,y)+\beta(z)\ln\delta(f,z)\,\widetilde{\varepsilon}_j\mid\widetilde{r},v,z,y\right]$$

$$=E[W_j(1,z,y)\,f+\beta(z)\ln\delta(1,z)\,f\widetilde{\varepsilon}_j+W_j(0,z,y)\,(1-f)+\beta(z)\ln\delta(0,z)\,(1-f)\,\widetilde{\varepsilon}_j\mid\widetilde{r},v,z,y]$$

$$\begin{aligned}&=W_j(0,z,y)+\left[W_j(1,z,y)-W_j(0,z,y)\right]E(f\mid\widetilde{r},v,z,y)\\&+\beta(z)\left[\ln\delta(1,z)-\ln\delta(0,z)\right]E(f\widetilde{\varepsilon}_j\mid\widetilde{r},v,z,y).\end{aligned}$$

Next, observe that, since $W_j(f,z,y)$ is linear in $\ln y$, $E\left[W_j(0,z,y)\mid\widetilde{r},v,z\right]=W_j(0,z,\widetilde{y})$ and $E\left[W_j(1,z,y)\mid\widetilde{r},v,z\right]=W_j(1,z,\widetilde{y})$ where $\widetilde{y}=\widetilde{y}(\widetilde{r},v,z)$. Averaging the above expression over y, and noting that $f=\widetilde{f}(v,z,\widetilde{e}_1)$, we get

$$\begin{aligned}E(w_j\mid\widetilde{r},v,z)&=W_j(0,z,\widetilde{y})+\left[W_j(1,z,\widetilde{y})-W_j(0,z,\widetilde{y})\right]E(f\mid\widetilde{r},v,z)\\&+\beta(z)\left[\ln\delta(1,z)-\ln\delta(0,z)\right]E(f\widetilde{\varepsilon}_j\mid\widetilde{r},v,z).\end{aligned}$$

and by the conditional independence assumptions regarding $\widetilde{\varepsilon}_j$ and $\widetilde{e}_1$,

$$\begin{aligned}E(w_j\mid\widetilde{r},v,z)&=W_j(0,z,\widetilde{y})+\left[W_j(1,z,\widetilde{y})-W_j(0,z,\widetilde{y})\right]E(f\mid v,z)\\&+\beta(z)\left[\ln\delta(1,z)-\ln\delta(0,z)\right]E(f\widetilde{\varepsilon}_j\mid v,z).\end{aligned}$$

Now the functions $E(w_j\mid\widetilde{r},v,z)$ and $\widetilde{y}(\widetilde{r},v,z)$ (the latter defined by $\ln\widetilde{y}(\widetilde{r},v,z)=E(\ln y\mid\widetilde{r},v,z)$) are both identified from data (and could, e.g., be consistently estimated by nonparametric regressions. So the derivatives of these expressions with respect to $\widetilde{r}$ are identified. This means that the following expression is identified.

$$\frac{\partial E\left(w_j \mid \widetilde{r}, v, z\right)}{\partial \ln \widetilde{r}} / \frac{\partial \ln \widetilde{y}\left(\widetilde{r}, v, z\right)}{\partial \ln \widetilde{r}} = \frac{\partial W_j\left(0, z, \widetilde{y}\right)}{\partial \ln \widetilde{y}} + \frac{\partial\left[W_j\left(1, z, \widetilde{y}\right) - W_j\left(0, z, \widetilde{y}\right)\right]}{\partial \ln \widetilde{y}} E\left(f \mid v, z\right) \tag{28}$$

Taking the difference between the above expression evaluated at $v = v_1$ and at $v = v_0$ then gives (and so identifies)

$$\frac{\partial\left[W_j\left(1, z, \widetilde{y}\right) - W_j\left(0, z, \widetilde{y}\right)\right]}{\partial \ln \widetilde{y}} \left[E\left(f \mid v_1, z\right) - E\left(f \mid v_0, z\right)\right]$$

and, since $E\left(f \mid v, z\right)$ is also identified, this identifies $\partial\left[W_j\left(1, z, \widetilde{y}\right) - W_j\left(0, z, \widetilde{y}\right)\right] / \partial \ln \widetilde{y}$. We can then solve equation (28) for $\partial W_j\left(0, z, \widetilde{y}\right) / \partial \ln \widetilde{y}$ where all the terms defining this derivative are identified. Taken together, the last two steps identify $\partial W_j\left(f, z, \widetilde{y}\right) / \partial \ln \widetilde{y}$ for $f = 0$ and for $f = 1$.

Given these identified functions and derivatives, we may then duplicate the proof of Lemma 5, (replacing y with $\widetilde{y}$, to show that the functions $\beta\left(z\right)$, $\eta_j\left(f, z\right)$, $\gamma_j\left(z\right)$, and $\delta\left(f, z\right)$ are identified.

6 Appendix B: Complete Estimates

Appendix Table A1: Full Estimates, Baseline Model
Number of obs = 2,866
Number of parameters = 111
Number of moments = 1803
Initial weight matrix: Unadjusted
GMM weight matrix: Robust
(Std. Err. adjusted for 281 clusters in uzcode)

	Coef.	Robust Std. Err.	z	P>|z|
eta_m				
one	0.348	0.005	65.820	0.000
avg_age_men	-0.001	0.001	-0.890	0.371
avg_age_women	0.003	0.001	1.750	0.080
avg_edu_men	0.001	0.001	1.450	0.148
avg_edu_women	0.001	0.001	1.560	0.119
avg_age_children	-0.025	0.004	-5.480	0.000

frac_girl	-0.007	0.003	-2.270	0.023
ln_dowry	0.002	0.000	6.620	0.000
ln_real_wealth	-0.001	0.000	-2.150	0.031
m1_f1_c1	0.025	0.007	3.570	0.000
m1_f1_c3	-0.072	0.010	-7.320	0.000
m1_f1_c4	-0.074	0.007	-10.260	0.000
m1_f2_c1	-0.004	0.006	-0.640	0.519
m1_f2_c2	-0.107	0.007	-15.770	0.000
m2_f1_c1	0.132	0.011	12.280	0.000
m2_f1_c2	0.055	0.009	6.320	0.000
m2_f2_c1	0.097	0.009	10.820	0.000
m2_f2_c2	-0.017	0.008	-2.260	0.024
f	0.014	0.002	7.110	0.000
gamma_m				
one	0.267	0.008	32.490	0.000
avg_age_men	0.006	0.001	4.440	0.000
avg_age_women	-0.010	0.003	-4.010	0.000
avg_edu_men	-0.002	0.001	-2.240	0.025
avg_edu_women	-0.001	0.001	-1.650	0.098
avg_age_children	-0.093	0.007	-13.220	0.000
frac_girl	0.025	0.006	4.190	0.000
ln_dowry	0.000	0.001	0.100	0.923
ln_real_wealth	0.002	0.001	2.450	0.014
m1_f1_c1	0.032	0.010	3.240	0.001
m1_f1_c3	0.036	0.017	2.130	0.034
m1_f1_c4	0.021	0.011	1.980	0.048
m1_f2_c1	-0.014	0.011	-1.350	0.177
m1_f2_c2	0.017	0.012	1.490	0.137
m2_f1_c1	0.038	0.013	2.840	0.005
m2_f1_c2	0.141	0.015	9.310	0.000
m2_f2_c1	-0.015	0.011	-1.420	0.156
m2_f2_c2	0.122	0.017	7.000	0.000
beta				
one	-0.151	0.007	-22.950	0.000
avg_age_men	0.005	0.001	5.330	0.000
avg_age_women	-0.001	0.002	-0.310	0.760
avg_edu_men	0.001	0.001	2.760	0.006
avg_edu_women	0.001	0.001	1.490	0.137
avg_age_children	-0.018	0.005	-3.550	0.000
frac_girl	0.009	0.004	2.100	0.036
ln_dowry	0.002	0.000	4.120	0.000
ln_real_wealth	0.001	0.000	2.400	0.016
m1_f1_c1	0.010	0.008	1.310	0.189
m1_f1_c3	0.000	0.011	-0.030	0.974

m1_f1_c4	-0.008	0.008	-0.960	0.339
m1_f2_c1	0.007	0.009	0.770	0.441
m1_f2_c2	-0.023	0.007	-3.130	0.002
m2_f1_c1	0.026	0.009	2.860	0.004
m2_f1_c2	0.060	0.010	6.120	0.000
m2_f2_c1	-0.001	0.008	-0.090	0.931
m2_f2_c2	0.037	0.010	3.770	0.000
lndelta				
one	-0.053	0.018	-2.980	0.003
eta_f				
one	0.305	0.005	63.190	0.000
avg_age_men	-0.002	0.001	-2.360	0.018
avg_age_women	-0.010	0.002	-6.600	0.000
avg_edu_men	-0.001	0.001	-1.220	0.224
avg_edu_women	-0.001	0.001	-1.180	0.238
avg_age_children	0.018	0.004	4.070	0.000
frac_girl	0.011	0.004	2.940	0.003
ln_dowry	-0.002	0.000	-5.630	0.000
ln_real_wealth	-0.001	0.000	-1.330	0.182
m1_f1_c1	0.022	0.005	4.250	0.000
m1_f1_c3	0.004	0.010	0.410	0.680
m1_f1_c4	-0.045	0.008	-5.400	0.000
m1_f2_c1	0.124	0.008	16.360	0.000
m1_f2_c2	0.095	0.007	14.180	0.000
m2_f1_c1	-0.041	0.008	-4.990	0.000
m2_f1_c2	-0.025	0.007	-3.610	0.000
m2_f2_c1	0.052	0.006	8.180	0.000
m2_f2_c2	0.068	0.007	9.930	0.000
f	-0.009	0.002	-4.710	0.000
gamma_f				
one	0.259	0.009	27.750	0.000
avg_age_men	0.012	0.002	7.320	0.000
avg_age_women	0.003	0.003	0.870	0.383
avg_edu_men	0.001	0.001	0.660	0.511
avg_edu_women	0.001	0.001	1.110	0.268
avg_age_children	-0.155	0.008	-18.260	0.000
frac_girl	0.009	0.007	1.360	0.173
ln_dowry	0.005	0.001	9.370	0.000
ln_real_wealth	-0.001	0.001	-0.890	0.372
m1_f1_c1	0.032	0.011	3.010	0.003
m1_f1_c3	-0.037	0.014	-2.730	0.006
m1_f1_c4	-0.027	0.015	-1.870	0.062
m1_f2_c1	0.025	0.014	1.750	0.081
m1_f2_c2	-0.035	0.011	-3.090	0.002

m2_f1_c1	0.022	0.014	1.530	0.126
m2_f1_c2	0.041	0.013	3.200	0.001
m2_f2_c1	0.019	0.014	1.360	0.175
m2_f2_c2	0.015	0.014	1.020	0.308
J test stat is 1741, 1692 df, p-val=0.1993				

7 References

Abraham, A., and S. Laczo, (2017), "Efficient risk sharing with limited commitment and storage," *The Review of Economic Studies*, 85(3), 1389-1424.

Anderberg, D. and H. Rainer, (2013), "Economic abuse: A theory of intrahousehold sabotage," *Journal of Public Economics*, 97(C), 282-295.

Angelucci, M., (2008), "Love on the Rocks: Domestic Violence and Alcohol Abuse in Rural Mexico," *The B.E. Journal of Economic Analysis & Policy*, 8(1), 1-43.

Ashraf, N. (2009) "Spousal Control and Intra-household Decision Making: An Experimental Study in the Philippines," *American Economic Review*, 99(4), 1245-1277.

Babu, B.V. and S.K. Kar, (2010), "Domestic violence in Eastern India: Factors associated with victimization and perpetration," Public Health, 124(3), 136-148

Banks, J., R. Blundell and A. Lewbel, (1997), "Quadratic Engel Curves and Consumer Demand," Review of Economics and Statistics, 79(4), 527-539.

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

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

Barten, A. P., (1964) "Family Composition, Prices, and Expenditure Patterns, " *Econometric Analysis for National Economic Planning: 16th Symposium of the Colston Society*

Edited by: Hart, P., Mills, G. and Whitaker, J. D. 277-292. London: Butterworth.

Becker, G. S., (1965) "A Theory of the Allocation of Time," *The Economic Journal*, 75(299), 493-517.

Becker, G. S., (1981) *A Treatise on the Family.* Cambridge, MA: Harvard University Press.

Bloch, F., and V. Rao, (2002) "Terror as a bargaining instrument: A case study of dowry violence in rural India," *American Economic Review*, 92(4), 1029-104

Blundell, R., P.-A. Chiappori and C. Meghir, (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.

Bourguignon, F. and P.-A. Chiappori, (1994) "The Collective Approach toHousehold Behavior" in The Measurement of Household Behavior, R. Blundell, I. Preston, and I. Walker, Eds., 70-85, Cambridge: Cambridge University Press.

Boyle, M.H., K. Georgiades, J. Cullen, and Y. Racine, (2009) "Community influences on intimate partner violence in India: Women's education, attitudes towards mistreatment and standards of living," *Social Science and Medicine*, 69(5), 691-697.

Brown, C., Calvi, R. and Penglase, J., 2018. Sharing the Pie: Undernutrition, Intra-household Allocation, and Poverty. Intra-Household Allocation, and Poverty (June 1, 2018).

Browning, M., F. Bourguignon and P.-A. Chiappori and V. Lechene, (1994) "Income and Outcomes: A Structural Model of Intrahousehold Allocation," *Journal of Political Economy*, 102(6), pages 1067-96.

Browning, M., and P.-A. Chiappori, (1998) "Efficient Intra-Household Allocations: A General Characterization and Empirical Tests," *Econometrica*, 66(6), 1241-1278.

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

Calvi, R., (2019) "Why Are Older Women Missing in India? The Age Profile of Bargaining Power and Poverty," mimeo.

Castilla, C. and T. Walker (2013), "Is Ignorance Bliss? The Effect of Asymmetric Information between Spouses on Intra-Household Allocations," *American Economic Review: Papers & Proceedings* 103(3): 263-268.

Cherchye, L., B. De Rock and F. Vermeulen, (2011)"The revealed preference approach to collective consumption behavior: testing and sharing rule recovery," *Review of Economic Studies*, 78(1), 176-198.

Cherchye, L., B. De Rock and F. Vermeulen, (2012) "Married with Children: A Collective Labor Supply Model with Detailed Time Use and Intrahousehold Expenditure Information," *American Economic Review*, 102(7), 3377-3405.

Cherchye, L.,B. De Rock, T. Demuynck, and F. Vermeulen, (2014) "Household consumption when the marriage is stable, " Unpublished manuscript.

Cherchye, L., B. De Rock, A. Lewbel and F. Vermeulen, (2015) "Sharing Rule Identification for General Collective Consumption Models," *Econometrica*, 83(5), 2001-2041.

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

Chiappori, P.-A., (1992) "Collective Labor Supply and Welfare," *Journal of Political Economy*, 100(3), pages 437-67.

Chiappori, P.-A., B. Fortin and G. Lacroix, (2002) "Marriage Market, Divorce Legislation, and Household Labor Supply," *Journal of Political Economy*, 110(1), 37-72.

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

Chiappori, P.-A., and M.Mazzocco, (2017) "Static and intertemporal household decisions," *Journal of Economic Literature*, 55(3), 985-1045.

Christensen, L. R., D. W. Jorgenson, and L. J. Lau, (1975) "Transcendental logarithmic utility functions," *The American Economic Review*, 65(3), 367-383.

Couprie, H., E. Peluso and A. Trannoy, (2010) "Is power more evenly balanced in poor households?," *Journal of Public Economics*, 94(7-8), pages 493-507.

Deaton, A. and J. Muellbauer, (1980) "An Almost Ideal Demand System," *American Economic Review*, 70(3), 312-326.

Dercon, S. and P. Krishnan (2003) "In sickness and in health: Risk sharing within households in rural Ethiopia," *The Journal of Political Economy*, 108(4) 688-727.

Doepke, M. and F. Kindermann, (2019) "Bargaining over Babies: Theory, Evidence, and Policy Implications," *American Economic Review*, 109(9), 3264-3306.

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

Dunbar, G., A. Lewbel and K. Pendakur, (2019) "Identification of Random Resource Shares in Collective Households Without Preference Similarity Restrictions," forthcoming, *Journal of Business and Economic Statistics.*

Gobbi, P. E., (2018), "Childcare and commitment within households," *Journal of Economic Theory*, 176(C), 503-551.

Gutierrez, F. H., (2018). "The Intra-household Allocation of Resources in the Presence of Common-pool Goods: Theory and Application to the Measurement of Indifference Scales," Vanderbilt University unpublished manuscript.

Hansen, L. P., (1982) "Large sample properties of generalized method of moments estimators," *Econometrica: Journal of the Econometric Society*, 1029-1054.

Hidrobo, M., A. Peterman, and L. Heise, (2016) "The effect of cash, vouchers, and food transfers on intimate partner violence: evidence from a randomized experiment in Northern Ecuador," *American Economic Journal: Applied Economics*, 8(3), 284-303.

Hughes, C., M. Bolis, R. Fries, and S. Finigan, (2015) "Women's economic inequality and domestic violence: exploring the links and empowering women," *Gender and Development*, 23(2), 279-297.

Jorgenson, D. W., and D. T. Slesnick, (1987) "Aggregate consumer behavior and household equivalence scales," *Journal of Business & Economic Statistics*, 5(2), 219-232.

Koç, A., and C. Erkin, (2011). "Development, Women's Resources and Domestic Violence," *Globalization, Religion and Development*, 129-148.

Krishnan, S., C. H. Rocca, A. E. Hubbard, K. Subbiah, J. Edmeades, and N. S. Padian, (2010) "Do changes in spousal employment status lead to domestic violence? Insights from a prospective study in Bangalore," *India Social Science & Medicine*, Volume 70, Issue 1, pp. 136-143

Lewbel, A., (1991) "Cost of Characteristics Indices and Household Equivalence Scales," *European Economic Review*, 1991, 35, 1277-1293

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

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

Lise, J., and S. Seitz, (2011) "Consumption Inequality and Intra-Household Allocations," *Review of Economic Studies*, 2011(1), 328-355.

Lise, J., and K.Yamada, (2014) "Household sharing and commitment: Evidence from panel data on individual expenditures and time use (No. W14/05)," IFS Working Papers.

Lundberg, S. and R. Pollak, (1993), "Separate Spheres Bargaining and the Marriage Market," *Journal of Political Economy*, 101(6), 988-1010.

Mazzocco, M., (2007). "Household intertemporal behaviour: A collective characterization and a test of commitment," *The Review of Economic Studies*, 74(3), 857-895.

Menon, M., K. Pendakur and F. Perali, (2012) "On the Expenditure-Dependence of Resource Shares," *Economics Letters*, 117, 739-742.

Muellbauer, J. N. J., (1976) "Community Preference Fields and the Representative Consumer," *Econometrica*, 44(5), 979-999.

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

Ramos, A., (2016) "Household Decision Making with Violence: Implications for Transfer Programs," mimeo. Barcelona: Universitat Autonoma de Barcelona.

Rao, V., (1997) "Wife-beating in rural South India: A qualitative and econometric analysis," *Social Science & Medicine* 44(8), 1169-1180.

Schaner, S., (2015) "Do opposites detract? Intrahousehold preference heterogeneity and inefficient strategic savings," *American Economic Journal: Applied Economics*, 7(2), 135-74.

Schaner, S., (2017) "The cost of convenience? Transaction costs, bargaining power, and savings account use in Kenya," *Journal of Human Resources*, 52(4), 919-945.

Udry, C., (1996) "Gender, agricultural production, and the theory of the household," *The Journal of Political Economy*, 104(5), 1010-1046.

Vermeulen, F., (2002) "Collective Household Models: Principles and Main Results," *Journal of Economic Surveys*, 16(4), 533-64.

Vogler, C., and J. Pahl (1994) "Money, power and inequality within Marriage, " *The Sociological Review*, 42(2), 263-88.