

The Engel Curves of Noncooperative Households

Pierre-Andre Chiappori
Jesse Naidoo[†]

Abstract

We provide a set of necessary and sufficient conditions for a system of Engel curves to have been generated by a noncooperative model of family behavior. These conditions fully characterize the local behavior of household-level consumption in the cross-section, i.e. as a function of total income and distribution factors. In this setting, any demand system compatible with a noncooperative model is also compatible with a collective model, but the converse is not true. We describe how these nested conditions may be tested using standard instrumental-variables strategies. **JEL Classification: D7, D13.**

1 Introduction

Most households consist of more than one person, and in the last three decades economists have developed models which recognize that fact. Models that emphasize potential conflicts of interest within the household have been used to study a range of topics, including savings decisions (Anderson and Baland (2002), Schaner (2015)), labour supply (Chiappori, Fortin, and Lacroix (2002), Vermeulen (2005)), children’s health and education (Rangel (2006), LaFave and Thomas (2017)) the spread of HIV (Anderson (2018)), and even domestic violence (Angelucci (2008), Hidrobo, Peterman, and Heise (2016)).

The above list is not exhaustive. It also represents progress relative to standard practice before the 1990s, which was to adopt a “unitary” representation; unitary models are those in which a household behaves as if it is a single decision maker. This practice has been widely criticised, in particular because one of its key predictions - income pooling - has been rejected empirically, in many different contexts.¹ If a unitary model is inaccurate, the estimation of important behavioural parameters such as labour supply elasticities may be compromised. The failure to account for non-unitary features may also distort the policy conclusions that can be drawn from these models, as in the recent literature on intrahousehold inequality (Lise and Seitz (2011), Dunbar, Lewbel, and Pendakur (2013)).

A leading alternative to unitary models are “collective” models, which assume that households make Pareto efficient decisions. However, several authors have studied an alternative, noncooperative approach, which attempts to explain household behaviours as the (Nash) equilibrium of a game involving private contributions

*Corresponding author: jesse.naidoo@up.ac.za.

[†]We thank the editor and three anonymous referees for their comments and suggestions. This paper is partly based on Naidoo (2015), Naidoo’s Ph.D. thesis at the University of Chicago, and was largely written while Naidoo was a postdoc at the Harris School of Public Policy at the University of Chicago. He thanks Azeem Shaikh, Alessandra Voena, Robert Garlick, Chen Yeh, Stephane Bonhomme, Magne Mogstad, Philip Barrett and especially Derek Neal for their comments. Chiappori gratefully acknowledges financial support from the NSF (grant 1124277). **Current version:** October 03, 2019.

¹Three such rejections are Duflo (2003), Lundberg, Pollak, and Wales (1997), and Thomas (1990).

to public goods.² In this paper we spell out the testable implications of a generic public goods game, and we compare them to the implications of collective models.

1.1 Individual Endowments and Distribution Factors

Any noncooperative model of public goods provision must rely on a notion of individual endowments or income. But standard approaches to noncooperative behaviour invoke a stronger assumption - that individuals not only fund their contributions to the public good from their own income, but also that these individual incomes are directly observable.³ Hence, any - possibly implicit - transfers between the spouses are observed by the econometrician; in addition, individual property rights on any income jointly received by the household are perfectly known.

This assumption introduces unduly restrictive and largely unrealistic limitations. Straightforward as it may sound in theory, the notion of individual income *within the household* is far from clear in practice. For instance, several empirical studies have shown that such factors as laws governing divorce or the sex ratio prevailing on the relevant marriage market may affect consumption decisions. At first glance, this fact may seem to contradict a basic property of noncooperative models, namely that behavioural outcomes depend only on individual incomes. After all, these are not altered by laws governing divorce, at least for couples who remain married. A more sophisticated view, however, would recognize that the relevant notion of “individual incomes” is more complex, if only because it must include the current value of future income at various dates and in various states of the world. As such, it typically involves current beliefs about future events such as divorce; if individuals realize that they *may* divorce, these expectations will typically influence their choices even in a noncooperative context. Then one may expect any variables that influence the individual costs of divorce to affect household behaviour. For instance, the (expected) split of income from jointly owned assets may well affect household consumption decisions, although it may be difficult to say exactly who is entitled to what in practice - even if each spouse’s formal legal claims are clear. So the assumption that the income distribution is known is not very appealing.

However, we may still be able to observe predictors of this distribution. As we just discussed, laws governing divorce are likely to affect the distribution of joint income between spouses. More abstractly, we could define a *distribution factor* to be a variable which affects consumption decisions but does not affect preferences

²The goods under consideration are publicly consumed *within the household*, in the sense that they enter all members’ utilities. Various aspects of children’s welfare are the most obvious candidates for a “public good” within the household, but many others could be considered. For example, most durables, including housing itself, would naturally fall into this class.

³D’Aspremont and Dos Santos Ferreira (2014) is an exception among papers that take a “differentiable” approach to the analysis of household models; they provide a noncooperative model in which there are unobserved transfers between the spouses. Among the papers that exploit global features of demand data to identify and estimate household models, Cherchye, Demuyck, and De Rock (2011) also allows for unobserved individual endowments.

or the set of feasible allocations. Then any variable that affects an individual’s position on the “marriage market”, to use a seminal concept introduced by Becker, is a potential distribution factor, provided that its variations are able to influence intrahousehold allocation of resources.⁴ The notion of distribution factors is an important one in the collective literature; in a cross-sectional context such as the one we consider here, when at least two distribution factors are observable, the collective model generates strong testable restrictions on behavior (see Bourguignon, Browning, and Chiappori (2009)).

1.2 Main Result

In this paper, we extend those results to a noncooperative setting. Specifically, we propose a general model of a noncooperative household in which spouses take their private endowments as given, and make voluntary contributions to the set of public goods. Unlike most of the existing literature, we do *not* take these endowments to be observable by the econometrician; we only assume that there exists a reduced-form relationship between the distribution factors and the (unobserved) split of total income between the spouses, which ultimately generates individual endowments.⁵ In this sense, our model generalizes noncooperative models such as Leuthold (1968), Ashworth and Ulph (1981), Bergstrom, Blume, and Varian (1986), Chen and Woolley (2001) or Doepke and Tertilt (2017). We show that these assumptions alone are sufficient to generate strong testable restrictions upon observed behavior.

Our approach is designed to remain compatible with many interpretations. It may be the case that transfers take place between individuals, thus altering the distribution of endowments within the household relative to, say, a legal baseline; in that case, we are totally agnostic about the economic mechanisms generating these transfers, and simply allow for their existence. Alternatively, our framework may simply reflect the econometrician’s inability to fully observe the exact distribution of income (or property rights) within the household; then we simply postulate the existence of unobserved “individual incomes”, under the sole restriction that they must add up to the household’s total income.

We focus on cross-sectional data because of its abundance. Then the data will consist of household demands as a function of total household income and a vector of distribution factors.⁶ Given cross-sectional data in which distribution factors are observed to affect household choices, we provide conditions on the observable

⁴These are not the only possibilities. Empirically, several distribution factors have been used in the literature, including: individual relative incomes, the sex ratios on the relevant marriage markets (Chiappori, Fortin, and Lacroix (2002)); family background (Attanasio and Lechene (2014)), relative non-labour income (Lundberg, Pollak, and Wales (1997)), and laws governing divorce and alimony (Chiappori et al. (2017)). See also chapter 5 of Browning, Chiappori, and Weiss (2014) for a partial review of empirical studies of the variation in intrahousehold allocation due to distribution factors.

⁵We discuss how such a relationship might arise in section 2.2 below, although none of our results rely on the validity of a particular microfoundation.

⁶Our focus on the cross-section means we do not exploit variation in relative prices. But note that even if one has longitudinal data, *relative* prices may not vary much in practice. Nevertheless, in Section 5.3 we contrast our results with previous ones about the testability and identification of noncooperative and collective models in the presence of price variation.

demands that are - locally - necessary and sufficient for the data to have been generated by a noncooperative model.

We first establish that whenever distribution factors appear to matter, then the same distribution factor proportionality conditions - introduced by Browning and Chiappori (1998) and Chiappori, Fortin, and Lacroix (2002) - that are necessary and sufficient for a collective rationalization of demand data, are also implied by a noncooperative model. This means that without price variation, any demand that is compatible with a noncooperative model is also compatible with at least one collective model.

The converse, however, is not true. We provide a second set of restrictions that characterise Engel curves derived from a noncooperative model. These restrictions take the form of first, second, and third-order partial differential equations. Demands arising from a collective model will typically not satisfy these restrictions. Thus, in the absence of price variations, the set of demand systems generated by a noncooperative household are a strict subset of those generated by a collective one, at least from a local, differentiable perspective.⁷ This property has an important, direct implication: it is possible to test certain noncooperative models against collective alternatives. Most tests of intrahousehold efficiency (or, equivalently, of the collective model) have not specified an alternative; Del Boca and Flinn (2012) and Del Boca and Flinn (2014) are exceptions.

Because we focus on the local properties of demand for a nonparametric class of models, our conditions are stated as pointwise restrictions on demand that must hold everywhere in the neighbourhood of a given point. Obviously, testing such stringent conditions is impossible, so we exhibit two strategies for implementing tests in practice. One approach is to impose a parametric demand system, and translate our PDEs into restrictions on its coefficients. We work through such an example in section 4. Another, potentially more appealing, strategy would be to estimate a flexible local approximation to the demand system. In section 3.3, we exploit the results of Hausman et al. (1991) to show how this can be done using standard instrumental-variables methods. Instruments are necessary because the local approximations to the \mathbf{z} -conditional demands involve regressing the consumption of each good on the consumption of a public good, and thus they raise, by construction, a serious endogeneity problem. Fortunately, natural instruments for the consumption of the public good are available: namely, the distribution factors.

Our results hold because the only equilibria in a noncooperative model in which distribution factors affect demand involve free-riding. This fact leads to a specific kind of additive separability in the household's demand for private goods. The type of free-riding involved in these equilibria also leads to some exclusion

⁷“Differentiable” approaches take as their data demand as differentiable functions of prices and incomes; restrictions are then expressed in terms of partial differential equations that these functions must satisfy. By contrast, “revealed preference” approaches take as data a finite set of observations of prices and quantities; conditions then take the form of a finite set of inequalities.

restrictions in the demand for public goods. These restrictions arise because in the relevant noncooperative equilibria, spouses contribute to disjoint sets of public goods. It is important to stress that the testable implications of our framework come only from the assumption that spouses make private contributions to public goods, which is fundamental to any noncooperative model. In particular, we do not rely on arbitrary restrictions on preferences, or on the way in which distribution factors affect the intrahousehold division of wealth.

One limitation of our results is that they do not allow the econometrician to be completely agnostic about preferences: she must know that there is at least one public good within the household. Our results are also limited in that we cannot deliver the nonparametric identification of all aspects of preferences - hence our title.⁸

Distinguishing efficient behaviour from this type of noncooperative alternative is relevant for the measurement of inequality, as in Lise and Seitz (2011), and may have implications for matching patterns on the marriage market, as Pollak (2019) argues. It is also a necessary input into the efficient design of certain types of policy: whether to target certain public transfers by gender, for example, requires some knowledge about the magnitude of possible misallocation within the household. The same is true of choices between transfers in cash or in kind.

1.3 The Empirical Content of Non-Unitary Models: Previous Results

The empirical content of collective models, by contrast, has been fully characterised for some time. When price variations are observed, Browning and Chiappori (1998), using a “differentiable” approach, show that the demand functions generated by a collective model have a Slutsky matrix that satisfies a “symmetry plus rank one” (SR1) condition; later, Chiappori and Ekeland (2006) showed that these conditions - as well as some technical smoothness restrictions - are sufficient for (local) integration. Cherchye, Rock, and Vermeulen (2007) and Cherchye, Rock, and Vermeulen (2009) provide corresponding restrictions based on a Revealed Preference approach; that is, they provide a complete characterization of the set of price-consumption pairs that can be generated by a collective model.

In a cross-sectional context, Bourguignon, Browning, and Chiappori (2009) have shown that a condition called distribution factor proportionality (DFP), introduced by Browning and Chiappori (1998) and Chiappori, Fortin, and Lacroix (2002), is necessary for given Engel curves to be compatible with a collective model; and that, conversely, demand functions satisfying DFP can always be rationalised as stemming from a collective

⁸Sufficient conditions for the full identification of a noncooperative model are not yet known, but presumably at least some price variation would be required. See section 5.3 below.

model where all commodities are publicly consumed. These properties have been tested in several articles; empirical evidence seems broadly supportive of the collective model, in the sense that most studies fail to reject it.⁹

Much less is known about the demand functions arising from a noncooperative model. When prices vary, Lechene and Preston (2011) derive necessary conditions on the structure of the Slutsky matrix - conditions which are qualitatively similar to, although less restrictive than, the SR1 conditions of Browning and Chiappori (1998) - but whether these conditions are sufficient is not yet known. Cherchye, Demuynck, and De Rock (2011) approach the same question from a revealed preferences perspective; they provide a set of linear inequalities which, if satisfied, allow the observed aggregate consumption of the household to be decomposed into two distinct sets of price-consumption pairs, each of which must satisfy GARP.¹⁰ These inequalities are necessary and sufficient for a noncooperative rationalization of the data, in the usual revealed-preference sense. However, absent price variation - i.e., when demands only depend on incomes and distribution factors - nothing is known about the restrictions implied by the noncooperative approach. We fill that gap in this paper.

2 Noncooperative Engel Curves

2.1 Notation and Aggregate Constraints

We start with a general, noncooperative model of the household consisting of a public goods game. Note that it is necessary to have at least one public good in the household: otherwise, the first welfare theorem will imply that noncooperative decision-making is efficient, removing the distinction between noncooperative and collective models.

Consider a two-person household with members named a and b . The household has total wealth y , which can be used to consume several goods. Some of these goods are public; so let Q_j be the household's consumption of the j -th public good, and say there are N of these. Let q_i^a be a 's consumption of the i -th private good, and similarly let q_i^b be b 's consumption of it. At the household level, the econometrician observes the vector $(Q_1, \dots, Q_N, q_1, \dots, q_n)$, where $q_i = q_i^a + q_i^b$ is the aggregate consumption of private good i .

For now, we will assume the identities of the public and private goods are known; relaxing this assumption

⁹Some prominent examples are Browning and Chiappori (1998) for the SR1 prediction involving price variations, Chiappori, Fortin, and Lacroix (2002) for labour supply models, and Browning et al. (1994), Bobonis (2009) and Attanasio and Lechene (2014), for the ‘‘Engel curve’’ framework.

¹⁰They also provide an example of a finite set of possible price-quantity observations which are consistent with a noncooperative model in which all goods are public, but not with any collective model. Note, however, that approaches based on Revealed Preference are intrinsically global, whereas the differentiable perspective typically generates local conditions. We return to this distinction in Section 5.4.

will be discussed in section 5.1. Since relative prices are constant, we may choose the physical units of the goods such that all prices are unity, so the household’s budget constraint is

$$\sum_{i=1}^n q_i + \sum_{j=1}^N Q_j = y, \tag{1}$$

where y denotes the household’s total income y .

We assume that the econometrician observes a vector \mathbf{z} of dimension $K \geq 2$; below, we discuss the economic and econometric assumptions \mathbf{z} should satisfy. For now it is enough that distribution of consumption varies with \mathbf{z} , conditional on y . The data consist of the functions $Q_j(y, \mathbf{z})$, for $j = 1, \dots, N$ and $q_i(y, \mathbf{z})$ for $i = 1, \dots, n$, observed over some open neighbourhood \mathcal{O} of a given point $(\bar{y}, \bar{\mathbf{z}})$. We further assume that observed demand functions are smooth.

2.2 Distribution Factors in Noncooperative and Collective Models

Next, we want to capture the notion that individuals privately (and noncooperatively) contribute to the household’s public goods. They fund their private consumption and their contributions to the public goods out of their individual endowments. The exact value of the endowments may not be observable to the econometrician. So, we assume that there is a function $\rho(y, \mathbf{z})$ linking total income y and a vector of observables \mathbf{z} to the intrahousehold distribution of endowments, so that a receives $\rho(y, \mathbf{z})$ while b gets $y - \rho(y, \mathbf{z})$. Here, the variables in \mathbf{z} are *distribution factors*, in the sense that they do not affect preferences or the budget set, but only the intrahousehold income distribution. We do *not* take a stand on how or why these variables affect intrahousehold distribution. That is, we do not impose any restrictions on the functional form of $\rho(y, \mathbf{z})$. Instead, we treat $\rho(y, \mathbf{z})$ as a reduced-form relationship that could have several possible microfoundations.

One possibility directly generalizes the standard noncooperative framework, in which individuals only control their own income. But, as we argued above, the notion of “individual incomes” within a couple is often ambiguous and hard to define (let alone measure): while individual property rights can precisely be identified for some sources of income, other types of income are jointly received by the household, and their allocation between spouses is not observable.

For instance, suppose that a and b supply labour inelastically, earning y^a and y^b respectively; and that they have some joint assets generating an income flow of y^{ab} . Then $y = y^a + y^b + y^{ab}$ is their total income.

Suppose also that the law is such that each spouse has sole - and costlessly enforced - claim to their own labour income, but that a is entitled to a share $\sigma(\mathbf{z})$ of the joint income. Here \mathbf{z} could represent legal rules, the offers of their extended families (if some component of y^{ab} is a claim on inheritance or informal insurance, say), or other social or institutional factors; and σ summarizes their effects.

If all the components of the couple's wealth are observed, y^a/y^b is a natural distribution factor, but it is evidently not the only one available. And for the purposes of isolating variation in consumption choices that are driven by factors outside the household, it may not even be the most empirically useful one - that will depend on (among other things) the joint distribution of \mathbf{z} and the various components of income.

Another possibility - not entirely distinct from the first - is that the spouses play some type of (noncooperative) bargaining game to allocate income, and that \mathbf{z} describes the conditions on the marriage market which affect the equilibrium of this game. Then, given the equilibrium of this "first stage" game, the spouses make consumption decisions.

Generally speaking, the apparent effects of distribution factors are often interpreted in terms of a collective model - i.e., \mathbf{z} is a determinant of the Pareto weights within marriage. But the fact that distribution factors have detectable effects does not imply anything about the economic mechanisms underlying those effects. Since we want to explore the implications of the alternative, noncooperative model, it seems only fair to allow the same level of richness and flexibility in both. In a noncooperative framework of our kind, the only channel for distribution factors to affect demand is through the effective intrahousehold distribution of wealth. Still, such effects may arise in both the collective or noncooperative models.

2.3 Individual Behavior and Equilibrium

Lastly, we need to describe individual behaviour in our framework. Following the existing literature, we use (noncooperative) Nash equilibrium as our solution concept; that is, individuals each choose their private consumption and individual contributions to public goods taking as given their partner's behaviour. Then, the observed consumption vector $(Q_1, \dots, Q_N, q_1, \dots, q_n)$ solves the following system:

$$\max_{Q^a, q^a} U^a(Q^a + Q^b, q^a) \quad \text{s.t.} \quad \sum_i q_i^a + \sum_j Q_j^a = \rho(y, \mathbf{z}) \quad (2)$$

$$\max_{Q^b, q^b} U^b(Q^a + Q^b, q^b) \quad \text{s.t.} \quad \sum_i q_i^b + \sum_j Q_j^b = y - \rho(y, \mathbf{z}) \quad (3)$$

Here, U^a and U^b represent the preferences of the household members, $Q^x = (Q_1^x, \dots, Q_N^x)$, for $x = a, b$ denotes agent x 's contribution to the public goods. Of course, in household-level data, we only observe the aggregate

consumption vector.

3 Results

Our analysis starts from a basic result, initially derived by Bergstrom, Blume, and Varian (1986) in the case of a single public good, and later extended by Browning, Chiappori, and Lechene (2010) to a general setting. The result states that there are two possible types of equilibria in the game defined above. One type - which, following Lundberg and Pollak (1993), could be called a “separate spheres” equilibrium - requires that each public good is contributed to by *only one* of the agents. Alternatively, there may be one (but generically not more than one) commodity to which both agents contribute. In the latter case, income pooling must hold: $\partial Q_j / \partial z_k$ and $\partial q_i / \partial z_k$ are zero for all goods i, j and all distribution factors k in some open set. In general, and for given preferences, to which of the two types the Nash equilibrium belongs depends on the parameters (prices, incomes, etc.). For instance, if one continuously changes the internal allocation of resources, the equilibrium typically switches (possibly several times) from one type to the other.¹¹

The properties of aggregate demands depend on the type of equilibrium that obtains. The second type of equilibrium - overlapping contributions to one public good - is straightforward: demand is unaffected by the distribution factors, so it is locally compatible with *any* model of household behaviour, whether unitary, collective or noncooperative. These cases require no further investigation, because the testable implications of unitary models have been fully characterized already.

Less obvious is the “separate spheres” case, where each public good is contributed by only one of the agents, which we consider next. Again, our approach is local; that is, we consider a point $(\bar{y}, \bar{z}) \in \mathcal{O}$ such that $\partial Q_j / \partial z_l(\bar{y}, \bar{z}) \neq 0$ for at least one public good j and one distribution factor l , and derive conditions that noncooperative demands must satisfy in some open neighbourhood of that point.

3.1 Distribution Factor Proportionality

Our first result is that “distribution factor proportionality” (DFP), a property of demand systems first introduced by Browning and Chiappori (1998) as an implication of efficiency, is also an implication of the noncooperative model. To see this, let α_i and β_i respectively denote a ’s and b ’s Marshallian demand for commodity i . Then:

¹¹A sufficiently large redistribution may cause the household to switch between an income pooling equilibrium and a “separate spheres” equilibrium (or vice versa); see Browning, Chiappori, and Lechene (2010). The introduction of a large cash transfer program might be interpreted in such a way, although they would also have effects on the level of the household’s budget. And, of course, the conditionalities built into many such programs may alter the relative prices facing the household.

$$q_i = \alpha_i(\rho(y, \mathbf{z})) + \beta_i(y - \rho(y, \mathbf{z})) \quad (4)$$

$$Q_j = \begin{cases} \alpha_j(\rho(y, \mathbf{z})) & \text{if } a \text{ contributes to } j \\ \beta_j(y - \rho(y, \mathbf{z})) & \text{if } b \text{ contributes to } j \end{cases} \quad (5)$$

and differentiating, we have:

$$\frac{\partial q_i(y, \mathbf{z}) / \partial z_k}{\partial q_i(y, \mathbf{z}) / \partial z_l} = \frac{\partial Q_j(y, \mathbf{z}) / \partial z_k}{\partial Q_j(y, \mathbf{z}) / \partial z_l} = \frac{\partial \rho(y, \mathbf{z}) / \partial z_k}{\partial \rho(y, \mathbf{z}) / \partial z_l}. \quad (6)$$

Summarizing, we have the following.

PROPOSITION 1. *In a noncooperative separate spheres equilibrium, household demands satisfy the Distribution Factor Proportionality (DFP) property: for all goods i, j and distribution factors k, l ,*

$$\frac{\partial q_i(y, \mathbf{z}) / \partial z_k}{\partial q_i(y, \mathbf{z}) / \partial z_l} = \frac{\partial Q_j(y, \mathbf{z}) / \partial z_k}{\partial Q_j(y, \mathbf{z}) / \partial z_l}. \quad (7)$$

An equivalent formulation uses the notion of a \mathbf{z} -conditional demand system, introduced by Bourguignon, Browning, and Chiappori (2009). We may, with no loss of generality, consider a $(\bar{y}, \bar{\mathbf{z}})$ such that $\partial Q_1 / \partial z_1(\bar{y}, \bar{\mathbf{z}}) \neq 0$. Then the function $Q_1(y, \bar{\mathbf{z}})$ can be locally inverted as $z_1 = \zeta(y, Q_1, z_2, \dots, z_K)$. For the goods other than good 1, define the z -conditional demand by:

$$\chi_i(y, Q_1, z_2, \dots, z_K) = q_i(y, \zeta(y, Q_1, z_2, \dots, z_K), z_2, \dots, z_K).$$

Then the previous result can be restated as:

COROLLARY 1. *The z -conditional demands generated by a noncooperative model satisfy $\partial \chi_i / \partial z_k = 0$ for all goods i and all distribution factors $k \geq 2$.*

For a proof, see Bourguignon, Browning, and Chiappori (2009).

The DFP property means that all distribution factors affect demand through a scalar index, and when many goods are demanded, this fact creates cross-equation restrictions. In collective models, DFP follows from

the fact that all distribution factors operate through the Pareto weight. But in noncooperative models, too, distribution factors act only through a scalar index: namely ρ , which represents the household's effective division of wealth.

Another corollary is that the demand generated by a noncooperative model can also be rationalised by a collective model - of course with different preferences:

COROLLARY 2. *Take a solution curve $(Q(y, \mathbf{z}), q(y, \mathbf{z}))$ lying in \mathbb{R}^{N+n} to the programs (2) and (3). Then $(Q(y, \mathbf{z}), q(y, \mathbf{z}))$ is locally collectively rationalisable in the sense that there exists two utilities $V^1(Q, q), V^2(Q, q)$ and a nonnegative function $\mu(y, \mathbf{z})$ such that (Q, q) solves the social planner's problem*

$$\max V^1(Q, q) + \mu(y, \mathbf{z})V^2(Q, q)$$

subject to the budget constraint (1) in some open neighbourhood of $(\bar{y}, \bar{\mathbf{z}})$.

This follows by combining Proposition 2 of Bourguignon, Browning, and Chiappori (2009) and our Proposition 1. So, if we consider a common version of noncooperative household models, where z is a scalar representing a 's share, and the distribution of private endowments is observed - i.e., $\rho(y, z) = zy = y^a$ for some observable y^a - any demand function derived from that setting can also be rationalised by a collective model. However, the economic interpretation of such a collective model may differ substantially from that of the noncooperative model that generated the data.¹²

3.2 Additional Restrictions

The noncooperative model implies further restrictions on demand, at least in the presence of public goods. Suppose DFP holds, so that without loss of generality, we can assume there is only one distribution factor. We also assume that the income elasticities for all of the public goods are nonzero for both a and b , so that variations in relative income induce variation in the consumption of the public goods.

Notice that we can, on the basis of the z - conditional demands, partition the set of public goods $1, \dots, N$ such that goods from distinct subsets are contributed by distinct members. Take an arbitrary public good - say Q_1 - and assume it is contributed by person a , which can be ensured by a suitable definition of ρ . If $N = 1$, so there are no other public goods, we are done. And if $N \geq 2$, notice that another public good j is contributed by a if and only if

¹²For instance, a careful reading of the proof of Proposition 2 of Bourguignon, Browning, and Chiappori (2009) reveals that the rationalizing model constructed there is such that all goods are *public*. If, however, one imposes that some commodities must be private, then the collective model imposes restrictions over and above DFP for these commodities, although these conditions are quite intricate. Those conditions were not derived in Bourguignon, Browning, and Chiappori (2009).

$$\frac{\partial \Xi_j}{\partial Q_1}(y, Q_1) > 0$$

where $\Xi_j(y, Q_1)$ is the z -conditional demand for public good j . This is because, conditional on y , increases in Q_1 only happen when ρ increases.¹³ So, we may assume that we know which public goods are contributed by each member, at least up to a permutation of names.

PROPOSITION 2. *Order the public goods such that a contributes to goods $1, \dots, L$ and b contributes to $L+1, \dots, N$. Let χ_j represent the z -conditional demand for the j -th private good, and let Ξ_k represent the demand for the k -th public good. In a neighbourhood of any point (\bar{y}, \bar{z}) such that $\partial Q_1 / \partial z_k(\bar{y}, \bar{z}) \neq 0$ for at least one distribution factor k , the z -conditional demand functions generated by the noncooperative model must satisfy the following partial differential equations: for all private goods i, j ,*

$$\frac{\partial^2 \chi_i / \partial y \partial Q_1}{\partial^2 \chi_i / \partial y^2} - \frac{\partial^2 \chi_j / \partial y \partial Q_1}{\partial^2 \chi_j / \partial y^2} = 0 \quad (8)$$

$$\text{and } \frac{\partial}{\partial y} \left(\frac{\partial^2 \chi_i / \partial y \partial Q_1}{\partial^2 \chi_i / \partial y^2} \right) = 0; \quad (9)$$

for all public goods contributed by a , $2 \leq j \leq L$,

$$\frac{\partial \Xi_j(y, Q_1)}{\partial y} = 0; \quad (10)$$

and for all public goods $L+1 \leq k, l \leq N$ contributed by b ,

$$\frac{\partial \Xi_k(y, Q_1) / \partial Q_1}{\partial \Xi_k(y, Q_1) / \partial y} = \frac{\partial \Xi_l(y, Q_1) / \partial Q_1}{\partial \Xi_l(y, Q_1) / \partial y} \quad (11)$$

$$\text{and } \frac{\partial}{\partial y} \left(\frac{\partial \Xi_k(y, Q_1) / \partial Q_1}{\partial \Xi_k(y, Q_1) / \partial y} \right) = 0. \quad (12)$$

Lastly, the effective division of wealth ρ is identified up to an additive constant; for any value of that constant, individual Engel curves are identified up to one common constant each.

Proof. See Appendix. □

¹³Formally, let ξ be the inverse of a 's Engel curve for Q_1 ; then $\rho = \xi(Q_1)$, and the z -conditional demand for goods in a 's sphere satisfy $Q_j = A_j(\xi(Q_1))$, with both A_j and ξ increasing. On the other hand, for the goods contributed by b , their z -conditional demands satisfy $Q_j = B_j(y - \xi(Q_1))$, with B_j increasing. Conditional on y , the demands for the public goods have a one-factor structure, where all variation is driven by the unobserved scalar ρ .

Some intuition for these PDEs - which follows from the steps in the proof - is as follows: suppose a 's Engel curve for the first public good was known. In a separate spheres equilibrium where only a contributes to the first public good, we could locally invert a 's demand for Q_1 to obtain ρ ; then, the joint distribution of $(\rho, y - \rho)$ and the consumption of each of the private goods identifies the slope of both a and b 's Engel curves for the private goods. Of course, we do not know a 's Engel curves, but we will be able to recover ρ up to a constant by exploiting the fact that aggregate demand for each private good is the sum of a 's demand (which is a function of ρ alone) and b 's demand (which is a function of $y - \rho$).

This additive separability, and the observability of aggregate wealth y , forces each of the cross-partial derivatives $\partial^2 \chi_i / \partial y \partial Q_1$ to be proportional to $-\partial^2 \chi_i / \partial y^2$ (since, conditional on y , increases in ρ are decreases in $y - \rho$). The common constant of proportionality turns out to be the derivative of the inverse of a 's Engel curve for Q_1 . This identifies ρ up to a constant, and gives the exclusion restriction (9). Finally, the slopes of both partners' Engel curves for the private goods can then be recovered from the first derivatives of the z -conditional demands.

In sum, the conditions on private demands express the fact that under the noncooperative model, the inverse Engel curve $\rho(Q_1)$ is well-defined (as reflected in (8)) and is a function of Q_1 only (as reflected in (9)). The conditions on the other public goods reflect the fact that spouses contribute to disjoint sets of public goods. Obviously, the identification of preferences and of the intrahousehold distribution of wealth is only up to a permutation of the spouses' identities, unless additional information on preferences - for example, in the form of an assignable good - is available.¹⁴

3.3 Implementation via Instrumental Variables

Propositions 1 and 2 suggest a natural testing strategy: first, test whether DFP holds in the data. If it does, proceed to test whether the cross-equation restrictions (8) and the exclusion restrictions (9) hold.¹⁵ The first step is by now standard, and has been implemented in a collective framework by various authors, e.g., Attanasio and Lechene (2014). We now discuss how to implement tests of the second set of restrictions - and thus, of the noncooperative model - using standard instrumental-variables methods. We make no claim that the specification we suggest here will be statistically optimal, but at least the identification conditions are easy to interpret.

¹⁴This is the approach followed by Boone et al. (2014), who use leisure as an assignable good and treat expenditures on children, and non-child expenditures, as two distinct public goods.

¹⁵On the other hand, if the data reject the proportionality restrictions (7), or their equivalent z -conditional forms, there are three possible explanations: either the distribution factors are invalid; the econometric model is misspecified (due to say selection biases or omitted variables); or, neither a collective model nor a noncooperative one can rationalize the data. Of course, these explanations are not mutually exclusive!

Consider the cross-equation restrictions (8). Take a second-order approximation to the z -conditional demand for the i -th private good:

$$q_i \approx \mu_{i0} + \mu'_{i1} \mathbf{x}^* + \mathbf{x}^{*'} \mathbf{M}_{i2} \mathbf{x}^* + \varepsilon_i \quad (13)$$

where $\mathbf{x}^* = (y^*, Q_1^*)$ is the vector consisting of the household's true income y^* and true demand for the public good Q_1^* . The restrictions (8) directly translate into the statements that particular functions of the matrix \mathbf{M}_{i2} , and the coefficients in the vector μ_{i1} , must be equal across equations indexed by i . The problem, therefore, is to derive consistent estimates of these coefficients.

If we could observe \mathbf{x}^* directly, we could estimate the coefficients in (13) by least squares for each private good i . Then testing the cross-equation restrictions (8) would be straightforward, although the power of the test might be compromised by correlations in the disturbances across goods (i.e. that ε_i might be correlated with ε_j). However, unless the joint distribution of the data is degenerate, we will have measurement error (or unobserved heterogeneity) in the observed demands, including that of the public good Q_1^* . So, suppose that what the econometrician observes is actually

$$\mathbf{x} = \mathbf{x}^* + \eta$$

where we also allow for possible measurement error in household income y .

In practice, it is standard to proxy y^* with total expenditures, instrumented by total income. The case of Q_1 is however more difficult; we need, as instruments, variables which predict the consumption of the public good but do not directly affect the consumption of private goods. The distribution factors \mathbf{z} are exactly such variables: by assumption, they do not affect preferences or the set of feasible allocations.

The reason for this is that the χ_i and Ξ_j functions are z -conditional demands; as such, *they do not directly depend on the distribution factors* - the latter matter only through the demand for the public good. That is, the distribution factors \mathbf{z} do influence the choice of Q_1 ; however, *conditional on Q_1* , theory imposes that other demands are not affected by \mathbf{z} . In other words, a theoretical prediction is that *the distribution factors \mathbf{z} are valid instruments for the demand for public good 1*.

Explicitly, suppose that

$$\mathbf{x}^* = \gamma' \mathbf{w} + \nu \quad (14)$$

where \mathbf{w} is a vector containing the distribution factors, and possibly instruments for total household income (such as total consumption), and ν is prediction error. γ is a vector of “first-stage” coefficients. We may regard (14) as a linear approximation to the composition of a ’s Engel curve for the public good with $\rho(y, \mathbf{z})$, where a ’s individual income is predicted by the distribution factors.

Econometrically, this is precisely the situation considered by Hausman et al. (1991). The nonlinearity of the conditional mean $E[q_i|\mathbf{x}^*]$ in \mathbf{x}^* implies that the instruments \mathbf{w} need to satisfy stronger independence restrictions than in the linear case. For example, the instrument validity conditions will be satisfied if ν is statistically independent of ε_i , the measurement error in the dependent variable q_i . The instrument relevance conditions are slightly more involved than the textbook case, too: for example, it is necessary that the matrix of second moments of the joint distribution of \mathbf{w} and the products of its elements be nonsingular. This rules out binary distribution factors, for instance.

The exclusion restriction (9) can be tested in a similar way, and this can even be done using single-equation methods. However, doing so would require estimating a third-order approximation to the \mathbf{z} -conditional demands, with the accompanying costs of possibly lower precision.

Notice that the empirical role of $\rho(y, \mathbf{z})$ is to determine the strength of the “first stage”, where the distribution factors are used to predict the consumption of the public good. The properties of the first stage will also depend on preferences; in the noncooperative model, for example, the slope of the contributing spouse’s Engel curve for the public good will partially determine the strength of the distribution factors as instruments.

In contrast, preferences alone determine the local properties of \mathbf{z} -conditional demands, and thus the values of the “second-stage” coefficients μ_{i1} and \mathbf{M}_{i2} , in both the noncooperative and the collective models. This is because, conditional on y , the set of equilibria in each model is one-dimensional. In collective models, the set of equilibria is parameterised by a Pareto weight - θ , say. In the noncooperative model, it is parameterised by a ’s endowment ρ . The z -conditional form of demands amounts to a change of variables, using the consumption of the public good (or one of them, if there are several), rather than the unobservable θ or ρ , to parameterise the set of equilibria.

4 An Example

The restrictions on a household’s \mathbf{z} -conditional demands described in Proposition 2 are not in general true of collective models. Note, however, that it is necessary that the aggregate Engel curves of the household are nonlinear, because otherwise the ratios in (8) are undefined. Since the literature on Engel curves systematically uses nonlinear specifications, this would not seem to be much of a problem.

We now show, on a simple but generally applicable parametric example, how our tests can be performed. In particular, we show that the conditions described in Proposition 2 are not in general true of collective models; in fact, in our example, they are generically (i.e., except for a subset of coefficients of measure zero) violated.

We use the well-known Working-Leser form for Engel curves to illustrate how Proposition 2 may be implemented. This functional form is commonly used in empirical studies; in particular, demand systems based on Deaton and Muellbauer's "Almost Ideal" Demand System Deaton and Muellbauer (1980) will have Working-Leser Engel curves. In the collective literature, it is used by Bourguignon, Browning, and Chiappori (2009).

We now show that the restrictions of Proposition 2 imply a simple set of restrictions on the parameters of the demand system. Consider Engel curves of the form

$$q_i = a_i + b_i y + c_i y \ln y + \sum_k d_i^k z^k \quad (15)$$

$$Q_j = A_j + B_j y + C_j y \ln y + \sum_k D_j^k z^k \quad (16)$$

where y denotes income or, in practice, total expenditure, and $\mathbf{z} = (z^1, \dots, z^K)$ is a vector of distribution factors. We consider the neighbourhood of some point $(\bar{y}, \bar{\mathbf{z}})$ where $\partial Q_j / \partial z^k$ is nonzero for all j and all k . Then we cannot have income pooling, implying that one agent only contributes to each public good; say it is a who contributes to good 1.

4.1 Testing DFP

The first test relates to the Distribution Factor Proportionality condition. Here, (7) implies:

$$\frac{d_i^k}{d_i^l} = \frac{D_j^k}{D_j^l} = \frac{\Delta^k}{\Delta^l}$$

for all goods i, j and all distribution factors k, l . In other words, the coefficients of the various distribution factors must be proportional across all equations. It follows that for all i, j :

$$\sum_k d_i^k z^k = d_i \sum_k \Delta^k z^k \quad \text{and} \quad \sum_k D_j^k z^k = D_j \sum_k \Delta^k z^k$$

for some d_i, D_j . In effect, there is only one distribution factor, namely $z = \sum_k \Delta^k z^k$, and the functional form of the demand system simplifies to:

$$q_i = a_i + b_i y + c_i y \ln y + d_i z \quad (17)$$

$$Q_j = A_j + B_j y + C_j y \ln y + D_j z. \quad (18)$$

4.2 Local Restrictions on the Coefficients of Engel Curves

The conditions in Proposition 2 can directly be translated into cross-equation restrictions on the coefficients of the Engel curves (17); since these curves can be directly estimated (using total expenditures for y , instrumented by household income), statistical tests directly obtain.

Starting with the private goods, we see that:

$$\chi_i(y, Q) = \left(a_i - A_1 \frac{d_i}{D_1} \right) + \left(b_i - B_1 \frac{d_i}{D_1} \right) y + \left(c_i - C_1 \frac{d_i}{D_1} \right) y \ln y + \frac{d_i}{D_1} Q_1$$

which implies that

$$\frac{\partial \chi_i(y, Q_1)}{\partial Q_1} = \frac{d_i}{D_1} \implies \frac{\partial^2 \chi_i(y, Q_1)}{\partial y \partial Q_1} = 0 \text{ for all } i. \quad (19)$$

The conditions (8) and (9) are therefore always satisfied: with Working-Leser Engel curves, the noncooperative setting implies no restrictions on the demand for private goods.¹⁶

Things are different with other public goods, though. Again as before, let public goods 1 to L be paid for by a , and those from $L + 1$ to N be contributed by b . The z -conditional demand for public good j is:

$$\Xi_j(y, Q_1) = A_j - A_1 \frac{D_j}{D_1} + \left(B_j - B_1 \frac{D_j}{D_1} \right) y + \left(C_j - C_1 \frac{D_j}{D_1} \right) y \ln y + \frac{D_j}{D_1} Q_1. \quad (20)$$

Therefore, (10) requires

¹⁶This feature, however, is directly linked to a specific feature of the functional form we use: namely, that the impact of distribution factors on the demand for public good(s) is additively separable from that of household income y . If one introduces an interaction between y and \mathbf{z} , then the conditions on z -conditional demand for private goods would be testable even with one public good.

$$\frac{\partial \Xi_j(y, Q_1)}{\partial y} = 0 \Rightarrow B_j D_1 = B_1 D_j \text{ and } C_j D_1 = C_1 D_j. \quad (21)$$

for all goods j in a 's sphere. Further, we can tell which goods are in a 's sphere by testing whether $D_j/D_1 > 0$. Now consider goods in b 's sphere. Since

$$\frac{\partial \Xi_k(y, Q_1)/\partial Q_1}{\partial \Xi_k(y, Q_1)/\partial y} = \frac{D_k}{(B_k D_1 - B_1 D_k + C_k D_1 - C_1 D_k) + (C_k D_1 - C_1 D_k) \ln y},$$

equations (11) and (12) give, for all $L + 1 \leq k, l \leq N$:

$$C_k D_1 = C_1 D_k, \quad \text{and} \quad (22)$$

$$\frac{D_k}{B_k} = \frac{D_l}{B_l}. \quad (23)$$

Note that (21) is a particular case of (22) and (23). Hence, the demand system is compatible with the noncooperative setting if and only if (22) and (23) are satisfied for all public goods. In addition, any public good k such that $B_k D_1 - B_1 D_k \neq 0$ must be contributed to by b (i.e., by the agent who does not contribute to good 1).

In practice, therefore, with a Working-Leser structure the noncooperative model requires additional proportionality conditions between the coefficients of income (the B s and C s) and those of the distribution factors (the D s). Moreover, these conditions are also sufficient. Indeed, we have that:

$$\rho(y, z) = \lambda(B_1 y + C_1 y \ln y + D_1 z) + \mu$$

for some $\lambda > 0$ and an undetermined constant μ . Individual Engel curves are affine.

Recall that Bourguignon, Browning, and Chiappori (2009) showed any set of demand functions satisfying DFP can be rationalized by at least one collective model (in particular, one in which all commodities are public). Two conclusions follow:

- * in the absence of price variations, the set of demand functions for public goods generated by a noncooperative model with public goods is a strict subset of the set generated by a collective model.

* however, the empirical distinction requires that the econometrician knows which goods are publicly consumed. As expected, if all consumptions are private, the two sets coincide.

4.3 Local Restrictions on the Coefficients of z - Conditional Demands

Alternatively, one can directly estimate the z - conditional demands by regressing the demand for each public good $j \geq 2$ on $(y, y \ln y, Q_1)$, where y denotes household total expenditures, instrumented by household total income; remember that Q_1 must be instrumented by total income and distribution factors (including possibly individual incomes). The outcome of the regression gives:

$$\Xi_j(y, Q) = \alpha_j + \beta_j y + \gamma_j y \ln y + \delta_j Q_1$$

Again, if we define a to be person who contributes to Q_1 , we can test which goods are contributed by a by testing whether $\delta_j > 0$. Now:

* for all public goods j contributed by a , we have that

$$\frac{\partial \Xi_j(y, Q_1)}{\partial y} = 0, \text{ which implies } \beta_j = \gamma_j = 0.$$

* for all other public goods k, l , we have that

$$\frac{\partial \Xi_k(y, Q_1) / \partial Q_1}{\partial \Xi_k(y, Q_1) / \partial y} = \frac{\delta_k}{(\beta_k + \gamma_k) + \gamma_k \ln y}$$

and therefore

$$\frac{\partial}{\partial y} \left(\frac{\partial \Xi_k(y, Q_1) / \partial Q_1}{\partial \Xi_k(y, Q_1) / \partial y} \right) = 0 \Rightarrow \gamma_k = 0; \text{ and further,}$$

$$\frac{\partial \Xi_k(y, Q_1) / \partial Q_1}{\partial \Xi_k(y, Q_1) / \partial y} = \frac{\partial \Xi_l(y, Q_1) / \partial Q_1}{\partial \Xi_l(y, Q_1) / \partial y}, \text{ which implies } \frac{\delta_k}{\beta_k} = \frac{\delta_l}{\beta_l}.$$

In summary, in the z - conditional demand approach:

- * the coefficients of the $y \ln y$ terms must all be zero; and
- * the coefficients of the y terms must either be zero (in which case the commodity is contributed by a) or proportional to the coefficients of the Q_1 terms across equations (for those contributed by b).

4.4 Interpretation

Let us briefly summarize the interpretation for the various possible outcomes of the tests just described:

- * *If DFs are not significant:* then observed demands are compatible with *any* model - unitary, collective or non-cooperative.
- * *If DFs are significant, but do not satisfy DFP for commodities that are known to be publicly consumed* (i.e., in our example, (4.1) is violated): then observed demands are *not* compatible with *any* model, whether unitary, collective or non-cooperative.
- * *If DFs are significant, satisfy DFP, but do not satisfy the additional restrictions in Proposition 2:* then observed demands are
 - *not* compatible with the unitary model
 - compatible with the collective model
 - *not* compatible with the non-cooperative model, which would imply either income pooling or the conditions of Proposition 2.
- * *Lastly, if both DFP and the additional restrictions in Proposition 2 are satisfied,* then the observed demands are *not* compatible with the unitary model, but compatible with both the collective and the non-cooperative model. In practice, however, one should probably favour the non-cooperative setting, at least if the failure to reject is “strong enough” (i.e., does not simply reflect the lack of power of the tests). As clearly shown by the Working-Leser example, the additional conditions are non-generic (they are satisfied for a set of parameters of measure zero), so it would be hard to believe that the true model is a collective one and these conditions just happen to be satisfied by chance.

4.5 Global Tests

The tests just described are local. Regarding the *global* structure of the demand system, two points can be made.

- * At any given point, the equilibrium is either of an “income pooling” or of a “separate spheres” type. In the first case, *none* of the various components of household demand depend on any distribution factor. In the second case, they all do, and the impact of DFs is fully described by the conditions provided.
- * In principle, an empirical estimation should adopt a “switching regime” framework; indeed, theory suggests that the two types of equilibria will occur for different values of incomes and distribution

factors. In practice, such estimations are not too common; one can however mention a paper by Boone et al. (2014), who study a model in which demand for public goods is characterized by three possible regimes (“Husband Dictatorship”, “Wife Dictatorship”, and “Split Might”), and a paper by Browning and Lechene (2001) in which the presence of caring implies that one spouse acts as a Beckerian dictator for very unequal intrahousehold distributions.

Such estimations are left for future research.

5 Extensions

5.1 Unknown Identity or Number of Public Goods

The previous results rely on the assumption that we know which commodities are publicly consumed. What if that assumption were relaxed? Our results still hold, although more care needs to be taken in applying them. DFP still holds in the noncooperative model, and nothing about the identity of the public goods needs to be known to verify it.

As for Proposition 2, suppose first that N is known, and is at least one. Then there is at least one subset of size N of the observed commodities such that the PDEs above apply when the goods in the subset are treated as public. If N is itself unknown but larger than or equal to one, there is at least one value of N such that they hold. Thus, not knowing which goods are public creates some indeterminacy, but for the purposes of falsifying the noncooperative model, this indeterminacy is irrelevant. Alternatively, Proposition 2 provides an empirical way to test assumptions about the public nature of the goods: in a noncooperative household, a commodity cannot be public unless its consumption satisfies the PDEs of Proposition 2.

5.2 Exclusive or Assignable Consumption

In some contexts it may be safe to assume that a subset of the observed consumption goods benefits only one of the members of the household. For example, if it is known that only the husband smokes, or that only the husband benefits from his own clothing, then the aggregate consumption of either tobacco or men’s clothing can be treated as a component of his private consumption.

Such information implies further restrictions on the demand of a noncooperative household: for example, if good n is exclusively consumed by a , then the z -conditional demand for good n must satisfy $\partial\chi_n/\partial y = 0$. This follows from the fact that in a separate spheres equilibrium where the public good is financed by a ’s contributions alone, the aggregate consumption of the public good is a “sufficient statistic” for a ’s private

income. Thus, conditional on the level of the public goods contributed by a , variations in household wealth will have no effect on the consumption of good n .

5.3 Including Price Variation

In this paper, we have considered the set of functions of (y, \mathbf{z}) - that is, functions whose domain is a neighbourhood of a given point $(\bar{y}, \bar{\mathbf{z}}) \in \mathbb{R}^{1+K}$ - which can be generated as the aggregate demand of a noncooperative household. Our approach implicitly conditions on a fixed vector of relative prices, say $\mathbf{p} \in \mathbb{R}_{++}^{K-1}$. But as mentioned in the introduction, there are already results about the identification of, and testable restrictions implied by, collective and noncooperative models when prices are allowed to vary.

Specifically, Theorem 6 of Chiappori and Ekeland (2006) establishes necessary and sufficient conditions on the (pseudo-) Slutsky matrix of household demands generated by collective models, considered as a function of prices and aggregate wealth alone.¹⁷ That is, there are known conditions on the set of functions of (\mathbf{p}, y) that fully characterise the implications of the collective model. And Proposition 2 of Bourguignon, Browning, and Chiappori (2009) fully characterises the implications of the collective model, considered as a function of aggregate wealth and distribution factors alone - i.e. as a function of (y, \mathbf{z}) . These results have not yet been combined to give a set of conditions on the set of functions of prices, wealth, and distribution factors - $(\mathbf{p}, y, \mathbf{z})$ - that are sufficient for a collective rationalization.

Our results for the noncooperative model obviously hold at each price vector in an appropriate open domain, and so new necessary conditions can be easily derived (by differentiating our PDEs with respect to prices, for example). Another strategy would be to adapt the arguments of Lechene and Preston (2011): they derive necessary conditions on the pseudo-Slutsky matrix of household demand when a 's endowment ρ is observable.

A major challenge is to find sufficient conditions for a noncooperative rationalization of a given demand system. Without such conditions, it will be difficult to design powerful tests of the collective model against a noncooperative alternative that exploit the joint variation of prices, wealth and distribution factors without resorting to parametric assumptions. However, this is likely to be the less relevant case for applied researchers: it is hard enough to find exogenous variation in two or more distribution factors.

5.4 Local and Global Restrictions

The conditions derived above take a “differentiable” approach and are therefore intrinsically local. But the noncooperative setting generates additional, global conditions. For instance, it is typically the case that

¹⁷Chiappori and Ekeland (2009) goes further, by providing conditions under which the underlying collective model can be uniquely inferred from the observation of household demand.

the equilibrium switches between the two types of equilibria described above - income pooling and separate spheres - as relative incomes vary. It follows that, for some open subset of total wealth and distribution factors, household demand should not depend on distribution factors, a property that does not hold in general in collective models. Conversely, a collective model, besides DFP, must also satisfy global constraints, reflecting for instance the fact that the consumption of each good should be strictly positive for some values of the distribution factors. A general characterization of these constraints, however, raises complex difficulties for which the differentiable perspective may not be the natural tool. Global conditions are in general more tractable through a revealed preferences approach; however, revealed preference methods typically do not work in the absence of price variation.¹⁸

6 Concluding Remarks

In some applications, more information may be available than we have allowed our imagined econometrician: for example, one may have data on individual-level consumption, on price variation, or on observable heterogeneity in preferences. There is much work still to be done in these cases. Still, household-level surveys are probably the most commonly available type of microeconomic data. For such data, our results provide a full characterization of the demand patterns that can arise from noncooperative behaviour within the household.

These restrictions are consequences of the basic mechanism at stake in a noncooperative model, namely the private provision of public goods; in particular, they require no restriction whatsoever on the nature of the division of wealth within the household.¹⁹ In the absence of price variation, any behaviour that is compatible with a noncooperative framework, *even in our generalised sense*, can also be rationalised by a collective model. The converse, however, is not true. Noncooperative models generate restrictions on household demand beyond those implied by collective ones.

This finding stands in contrast to the findings of Cherchye, Demuynck, and De Rock (2011), who show that the restrictions implied by the collective and noncooperative models are disjoint. A possible explanation of the difference is that their revealed preference approach requires price variations, while our analysis relies on variation in income and distribution factors; again, we do not currently know of a necessary and sufficient characterization of demand functions deriving from a noncooperative framework in the “differentiable” approach with price variation.²⁰ Another possible reason for the difference between our conclusions and

¹⁸See Cherchye, Demuynck, and De Rock (2011).

¹⁹However, the fact that this relationship can be taken as given reflects a “partial equilibrium” assumption about the marriage market. If one wanted to jointly model intrahousehold allocation and equilibrium on the marriage market, the function $\rho(y, \mathbf{z})$ would have to be treated as endogenous. We thank an anonymous referee for making this point.

²⁰Note that, in many cases, providing such a characterization requires imposing an equilibrium selection device when there are

theirs is the local nature of our results, while a revealed preference approach is global by nature. Obviously, more work is needed on that topic.

Columbia University

University of Pretoria and University of Chicago

Appendix: Proof of Proposition 2

We start with the case $N = 1$. Say it is a who contributes to Q . Then:

$$Q(y, \mathbf{z}) = A(\rho(y, \mathbf{z})) \tag{A-1}$$

$$\text{and } q_i(y, \mathbf{z}) = \alpha_i(\rho(y, \mathbf{z})) + \beta_i(y - \rho(y, \mathbf{z})), \quad i = 1, \dots, n \tag{A-2}$$

where A is a 's Marshallian demand for the public good. Since $\partial Q / \partial z \neq 0$, we know that $A' \neq 0$; therefore (A-1) can be locally inverted:

$$Q(y, \mathbf{z}) = A(\rho(y, \mathbf{z}))$$

so we may write

$$\rho(y, \mathbf{z}) = \xi(Q(y, \mathbf{z}))$$

for $\xi = A^{-1}$; note that this function takes a scalar argument. Then the z -conditional demand for commodity i , corresponding to Q , is thus:

$$\chi_i(y, Q) = \alpha_i(\xi(Q)) + \beta_i(y - \xi(Q)), \tag{A-3}$$

implying that:

$$\frac{\partial \chi_i}{\partial y} = \beta'_i \text{ and } \frac{\partial \chi_i}{\partial Q} = (\alpha'_i - \beta'_i) \xi'(Q). \tag{A-4}$$

multiple equilibria.

Therefore,

$$\frac{\partial^2 \chi_i}{\partial y^2} = \beta_i'' \text{ and } \frac{\partial^2 \chi_i}{\partial y \partial Q} = -\beta_i'' \xi'(Q)$$

and:

$$\xi'(Q) = -\frac{\partial^2 \chi_i / \partial y \partial Q}{\partial^2 \chi_i / \partial y^2}. \quad (\text{A-5})$$

This implies two sets of PDEs:

$$\begin{aligned} \frac{\partial^2 \chi_i / \partial y \partial Q}{\partial^2 \chi_i / \partial y^2} - \frac{\partial^2 \chi_j / \partial y \partial Q}{\partial^2 \chi_j / \partial y^2} &= 0 \quad \forall i, j \\ \text{and } \frac{\partial}{\partial y} \left(\frac{\partial^2 \chi_i / \partial y \partial Q}{\partial^2 \chi_i / \partial y^2} \right) &= 0 \quad \forall i \end{aligned}$$

Conversely, if these conditions are satisfied, then (A-5) identifies ξ , therefore ρ , up to an additive constant. This constant is obviously not identifiable, and we may normalise it to be 0. Then (A-4) identifies β'_i and α'_i for all i , therefore β_i and α_i are identified up to one common additive constant from (A-3) which shows that the PDEs are locally sufficient.

Lastly, this argument can readily be extended to the case $N \geq 2$. Since $\partial Q_1 / \partial z(\bar{y}, \bar{\mathbf{z}}) \neq 0$, we must be in the “separate sphere” context where each public good is contributed to by one agent only. Now, one can always choose a point $(\bar{y}, \bar{\mathbf{z}})$ such that $\partial Q_j / \partial z(\bar{y}, \bar{\mathbf{z}}) \neq 0$ for all j . Order the public goods such that a contributes to goods $1, \dots, L$ and b contributes to $L + 1, \dots, N$. Then

$$Q_j(y, \mathbf{z}) = A_j(\rho(y, \mathbf{z})), \quad j = 1, \dots, L \quad (\text{A-6})$$

$$Q_k(y, \mathbf{z}) = B_k(y - \rho(y, \mathbf{z})), \quad k = L + 1, \dots, N \quad (\text{A-7})$$

$$\text{and } q_i(y, \mathbf{z}) = \alpha_i(\rho(y, \mathbf{z})) + \beta_i(y - \rho(y, \mathbf{z})), \quad i = 1, n \quad (\text{A-8})$$

If we assume, without loss of generality, that $L \geq 1$, we can compute the z -conditional demands corresponding to Q_1 as before. We get:

$$\Xi_j(y, Q_1) = A_j(\xi(Q_1)), \quad j = 2, \dots, L \quad (\text{A-9})$$

$$\Xi_k(y, Q_1) = B_k(y - \xi(Q_1)), \quad k = L + 1, \dots, N \quad (\text{A-10})$$

$$\chi_i(y, Q_1) = \alpha_i(\xi(Q_1)) + \beta_i(y - \xi(Q_1)), \quad (\text{A-11})$$

Now (A-9) requires

$$\frac{\partial \Xi_j(y, Q_1)}{\partial y} = 0, \quad j = 2, \dots, L$$

while from (A-10)

$$\frac{\partial \Xi_k(y, Q_1)}{\partial y} = B'_k \quad \text{and} \quad \frac{\partial \Xi_k(y, Q_1)}{\partial Q_1} = -B'_k \xi'(Q_1), \quad k = L + 1, \dots, N \quad (\text{A-12})$$

which gives:

$$\xi'(Q_1) = -\frac{\partial \Xi_k(y, Q_1) / \partial Q_1}{\partial \Xi_k(y, Q_1) / \partial y} \quad (\text{A-13})$$

for $k = L + 1, \dots, N$.

Conversely, if these conditions hold, then ξ is identified up to an additive constant; for any value of this constant, A_j and B_k are identified for all (j, k) . Finally, the conditions on $q_i(y, \mathbf{z})$ and the identification of individual Engel curves are derived as when $N = 1$ above.

- Anderson, Siwan. 2018. "Legal Origins and Female HIV." *American Economic Review* 108 (June): 1407–39. <https://doi.org/10.1257/aer.20151047>.
- Anderson, Siwan, and Jean-Marie Baland. 2002. "The Economics of ROSCAs and Intrahousehold Resource Allocation." *Quarterly Journal of Economics* 117 (3). Oxford University Press: 963–95. <http://qje.oxfordjournals.org/content/117/3/963.short>.
- Angelucci, Manuela. 2008. "Love on the Rocks: Domestic Violence and Alcohol Abuse in Rural Mexico." *The B.E. Journal of Economic Analysis & Policy* 8 (1). <https://doi.org/10.2202/1935-1682.1766>.
- Ashworth, J A, and D T Ulph. 1981. "Household Models." In *Taxation and Labour Supply*, edited by Charles Brown. London: Allen & Unwin.
- Attanasio, Orazio P, and Valérie Lechene. 2014. "Efficient Responses to Targeted Cash Transfers." *Journal of Political Economy* 122 (1). The University of Chicago Press: 178–222. <http://www.jstor.org/stable/info/10.1086/674968>.
- Bergstrom, Theodore C, Lawrence E Blume, and Hal R Varian. 1986. "On the Private Provision of Public Goods." *Journal of Public Economics* 29 (1). Elsevier: 25–49. <http://www.sciencedirect.com/science/article/pii/0047272786900241>.
- Bobonis, Gustavo J. 2009. "Is the Allocation of Resources within the Household Efficient? New Evidence from a Randomized Experiment." *Journal of Political Economy* 117 (3). The University of Chicago Press: 453–503. <http://www.jstor.org/stable/10.1086/600076>.
- Boone, Jan, Karen van der Wiel, Arthur van Soest, and Frederic Vermeulen. 2014. "Kinky choices, dictators and split might: a non-cooperative model for household consumption and labor supply." *IZA Journal of Labor Economics* 3 (1): 1–19. <https://doi.org/10.1186/2193-8997-3-11>.
- Bourguignon, Francois, Martin Browning, and Pierre-Andre Chiappori. 2009. "Efficient Intra-Household Allocations and Distribution Factors: Implications and Identification." *Review of Economic Studies* 76 (2). Oxford University Press: 503–28. <http://restud.oxfordjournals.org/content/76/2/503.short>.
- Browning, Martin, Francois Bourguignon, Pierre-Andre Chiappori, and Valérie Lechene. 1994. "Income and Outcomes: A Structural Model of Intrahousehold Allocation." *Journal of Political Economy* 102 (6). The University of Chicago Press: 1067–96. <http://www.jstor.org/stable/10.2307/2138780>.
- Browning, Martin, and Pierre-Andre Chiappori. 1998. "Efficient Intra-Household Allocations: A General Characterization and Empirical Tests." *Econometrica* 66 (6). JSTOR: 1241–78. <http://www.jstor.org/stable/10.2307/2999616>.

- Browning, Martin, Pierre-Andre Chiappori, and Yoram Weiss. 2014. *The Economics of the Family*. New York: Cambridge University Press.
- Browning, Martin, Pierre-André Chiappori, and Valérie Lechene. 2010. “Distributional Effects in Household Models: Separate Spheres and Income Pooling.” *Economic Journal* 120 (545): 786–99. <https://doi.org/10.1111/j.1468-0297.2009.02311.x>.
- Browning, Martin, and Valérie Lechene. 2001. “Caring and Sharing: Tests Between Alternative Models of Intra-household Allocation.” University of Copenhagen, Department of Economics.
- Chen, Zhiqi, and Frances Woolley. 2001. “A Cournot-Nash Model of Family Decision Making.” *Economic Journal* 111 (474): 722–48. <https://doi.org/10.1111/1468-0297.00657>.
- Cherchye, Laurens, Thomas Demuynck, and Bram De Rock. 2011. “Revealed Preference Analysis of Non-Cooperative Household Consumption.” *Economic Journal* 121 (555): 1073–96. <https://doi.org/10.1111/j.1468-0297.2011.02432.x>.
- Cherchye, Laurens, Bram de Rock, and Frederic Vermeulen. 2007. “The Collective Model of Household Consumption: A Nonparametric Characterization.” *Econometrica* 75 (2): 553–74. <http://doi.wiley.com/10.1111/j.1468-0262.2006.00757.x>.
- . 2009. “Opening the Black Box of Intrahousehold Decision Making: Theory and Nonparametric Empirical Tests of General Collective Consumption Models.” *Journal of Political Economy* 117 (6). The University of Chicago Press: 1074–1104. <http://www.jstor.org/stable/10.1086/649563>.
- Chiappori, Pierre-Andre, and Ivar Ekeland. 2006. “The Micro Economics of Group Behavior: General Characterization.” *Journal of Economic Theory* 130 (1): 1–26. <http://www.sciencedirect.com/science/article/pii/S0022053105000566>.
- . 2009. “The Microeconomics of Efficient Group Behavior: Identification.” *Econometrica* 77 (3). Wiley Online Library: 763–99. <http://onlinelibrary.wiley.com/doi/10.3982/ECTA5929/abstract>.
- Chiappori, Pierre-Andre, Murat Iyigun, Jeanne Lafortune, and Yoram Weiss. 2017. “Changing the Rules Midway: The Impact of Granting Alimony Rights on Existing and Newly Formed Partnerships.” *Economic Journal* 127 (604). New York: 1874–1905. <https://doi.org/10.1111/eoj.12385>.
- Chiappori, Pierre-André, Bernard Fortin, and Guy Lacroix. 2002. “Marriage Market, Divorce Legislation, and Household Labor Supply.” *Journal of Political Economy* 110 (1). UChicago Press: 37–72. <https://doi.org/10.1086/324385>.

- D'Aspremont, Claude, and Rodolphe Dos Santos Ferreira. 2014. "Household Behavior and Individual Autonomy: an Extended Lindahl Mechanism." *Economic Theory* 55 (3): 643–64. <http://www.jstor.org/stable/43562947>.
- Deaton, Angus S, and John Muellbauer. 1980. "An Almost Ideal Demand System." *American Economic Review* 70 (3). JSTOR: 312–26. <http://www.jstor.org/stable/10.2307/1805222>.
- Del Boca, Daniela, and Christopher Flinn. 2012. "Endogenous household interaction." *Journal of Econometrics* 166 (1). Elsevier B.V.: 49–65. <http://linkinghub.elsevier.com/retrieve/pii/S0304407611001187>.
- Del Boca, Daniela, and Christopher J Flinn. 2014. "Household Behavior and the Marriage Market." *Journal of Economic Theory* 150: 515–50.
- Doepke, Matthias, and Michèle Tertilt. 2017. "Does Female Empowerment Promote Economic Development?" Northwestern University, Department of Economics.
- Duflo, Esther. 2003. "Grandmothers and Grandaughters: Old Age Pension and Intra-Household Allocation in South Africa." *World Bank Economic Review* 17 (1): 1–25.
- Dunbar, Geoffrey, Arthur Lewbel, and Krishna Pendakur. 2013. "Children's Resources in Collective Households: Identification, Estimation, and an Application to Child Poverty in Malawi." *American Economic Review* 103 (1): 438–71. <https://doi.org/10.1257/aer.103.1.438>.
- Hausman, Jerry A., Whitney K. Newey, Hidehiko Ichimura, and James L Powell. 1991. "Identification and estimation of polynomial errors-in-variables models." *Journal of Econometrics* 50 (3): 273–95. [https://doi.org/10.1016/0304-4076\(91\)90022-6](https://doi.org/10.1016/0304-4076(91)90022-6).
- Hidrobo, Melissa, Amber Peterman, and Lori 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. <https://doi.org/10.1257/app.20150048>.
- LaFave, Daniel, and Duncan Thomas. 2017. "Extended families and child well-being." *Journal of Development Economics* 126 (July 2016). Elsevier: 52–65. <https://doi.org/10.1016/j.jdeveco.2016.11.006>.
- Lechene, Valérie, and Ian Preston. 2011. "Noncooperative household demand." *Journal of Economic Theory* 146 (2). Elsevier: 504–27. <http://www.sciencedirect.com/science/article/pii/S0022053110001407>.
- Leuthold, JH. 1968. "An Empirical Study of Formula Income Transfers and the Work Decision of the Poor." *Journal of Human Resources* 3 (3): 312–23. <https://doi.org/10.2307/145103>.
- Lise, J, and S Seitz. 2011. "Consumption Inequality and Intra-household Allocations." *The Review of Eco-*

nomie Studies 78 (1). Oxford University Press: 328–55. <http://restud.oxfordjournals.org/content/78/1/328.shortpapers2://publication/uuid/B27FE653-3ADD-482E-B141-069DD18ADB87>.

Lundberg, Shelly, and Robert A Pollak. 1993. “Separate Spheres Bargaining and the Marriage Market.” *Journal of Political Economy* 101 (6). The University of Chicago Press: 988–1010. <http://www.jstor.org/stable/10.2307/2138569>.

Lundberg, Shelly, Robert A Pollak, and Terence J Wales. 1997. “Do Husbands and Wives Pool Their Resources? Evidence From the UK Child Benefits.” *Journal of Human Resources* 32 (3): 463–80. <http://www.jstor.org/stable/146179>.

Naidoo, Jesse. 2015. “Essays on the Economics of the Family.” Ph.D. Thesis, University of Chicago. <http://search.proquest.com/docview/1705542406?accountid=14657>.

Pollak, Robert A. 2019. “How Bargaining in Marriage Drives Marriage Market Equilibrium.” *Journal of Labor Economics* 37 (1): 297–321. <https://doi.org/10.1086/698901>.

Rangel, Marcos A. 2006. “Alimony Rights and Intrahousehold Allocation of Resources: Evidence from Brazil.” *Economic Journal* 116 (513): 627–58. <https://doi.org/10.1111/j.1468-0297.2006.01104.x>.

Schaner, Simone. 2015. “Do Opposites Detract? Intrahousehold Preference Heterogeneity and Inefficient Strategic Savings.” *American Economic Journal: Applied Economics* 7 (2): 135–74. <https://doi.org/10.1257/app.20130271>.

Thomas, Duncan. 1990. “Intra-Household Resource Allocation: An Inferential Approach.” *Journal of Human Resources* 25 (4). JSTOR: 635–64. <http://www.jstor.org/stable/10.2307/145670>.

Vermeulen, Frederic. 2005. “And the winner is... An empirical evaluation of unitary and collective labour supply models.” *Empirical Economics* 30 (3). Springer-Verlag: 711–34. <http://link.springer.com/10.1007/s00181-005-0258-4>.