NBER WORKING PAPER SERIES

INTRAHOUSEHOLD INEQUALITY

Pierre-André Chiappori
Costas Meghir

Working Paper 20191
http://www.nber.org/papers/w20191

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
June 2014

Pierre-André Chiappori gratefully acknowledges financial support from the NSF (grant 1124277). Costas Meghir is grateful for financial support by the Cowles foundation and the Institution for Social and Policy Studies at Yale. Moreover, we thank Tony Atkinson, Francois Bourguignon and Marc Fleurbaey for useful comments on a previous version. The usual disclaimer applies.The views expressed herein are those of the authors and do not necessarily reflect the views of the National Bureau of Economic Research.

Intrahousehold Inequality
Pierre-André Chiappori and Costas Meghir
NBER Working Paper No. 20191
June 2014
JEL No. D1,D11,D12,D13,D31,D6,H41

ABSTRACT

Studies of inequality often ignore resource allocation within the household. In doing so they miss an important element of the distribution of welfare that can vary dramatically depending on overall environmental and economic factors. Thus, measures of inequality that ignore intra household allocations are both incomplete and misleading. We discuss determinants of intrahousehold allocation of resources and welfare. We show how the sharing rule, which characterizes the within household allocations, can be identified from data on household consumption and labor supply. We also argue that a measure based on estimates of the sharing rule is is inadequate as an approach that seeks to understand how welfare is distributed in the population because it ignores public goods and the allocation of time to market work, leisure and household production. We discuss a money metric alternative, that fully characterizes the utility level reached by the agent. We then review the current literature on the estimation of the sharing rule based on a number of approaches, including the use of distribution factors as well as preference restrictions.

Pierre-André Chiappori
Department of Economics
Columbia University
1009A International Affairs Building
420 West 118th St.
New York, NY 10027
pc2167@columbia.edu

Costas Meghir
Department of Economics
Yale University
37 Hillhouse Avenue
New Haven, CT 06511
and IZA
and also NBER
c.meghir@yale.edu

1 Introduction

1.1 Inequality between individuals

Consider an economy with identical couples, each of which has a total income of 100. Individuals privately consume a unique, perfectly divisible commodity; there exist neither externalities nor economies of scale, so that, in each couple, the sum of individual consumptions equals the couple's total income. Inequality, as measured in a standard way, is nil. Assume, now, that some of these couples divorce, and that after divorce husbands each receive an income of 75, while each wife gets 25. The new income distribution, again by standard criteria, is now unequal; in particular, the presence of lower income singles (the divorced wives) increases both inequality and poverty.

From a deeper perspective, however, the conclusion just stated is far from granted. It entirely relies on an implicit assumption – namely, that the pre-divorce distribution of income *within households* was equal. Most of the time, such an assumption has little or no empirical justification; and from a theoretical viewpoint, it is actually quite unlikely to hold – few serious models of household behavior would predict an equal distribution of income while married if the post-divorce allocation is highly skewed. Still, it is crucial. Assume, for the sake of the argument, that the distribution of resources within married couples simply mimics what it would be in case of divorce (he gets 75, she gets 25) – not an unreasonable assumption, given that in our (admittedly simplistic) structure this is the only individually rational allocation. Then the claim that inequality increased after the wave of divorces is simply wrong. Inequality, at least across individuals, has not changed; each agent has exactly the same income, consumption and welfare than before. And the surge in measured poverty is just as spurious. There are exactly as many poor women after than there were before; it is just that, in the pre-divorce situation, the standard measures missed them.

The previous example, extreme as it may be, illustrates a basic point that the present paper will try to emphasize – namely, that any attempt at measuring inequality (or its evolution over time) that ignores allocation of resources within the family is unreliable at best, and deeply flawed at worst, especially when the basic demographics regarding family composition evolve over the period under consideration. This point had already been emphasized in the literature; for instance, Haddad and Kanbur (1990) have showed, on Philippine data, that standard measures of inequality in calorie adequacy would be understated by 30 to 40 percent if intrahousehold inequality was ignored. As a more recent illustration, consider the graph below (Figure 1) due to Lise and Seitz (2011), that plots the evolution of inequality across households, within households and across individuals in the UK over the last decades, as estimated from a collective model of labor supply. The main conclusion is that the standard approach, based on adult equivalence scales, underestimates the initial level of cross-sectional consumption inequality by 50%. Moreover, it gives a deeply flawed picture of the evolution of inequality over the last decades. While the usual story - a large surge in inequality between 1970 and 2000 - applies to inequality across household, it is compensated by a considerable reduction of intra-household inequality - so that total inequality (across *individuals*) remains more or less constant over the period.[1]

[1]This conclusion must be qualified in view of the population under consideration. Indeed, the sample excludes all households with children, all persons aged under 22 or over 65, all persons who were self-employed, and the top1 per cent of the earnings distribution (which is, in any case, not well covered by the Family Expenditure Survey).; so the conclusions are only valid for that particular, subpopulation. Still, it is suggestive of the general claim that ignoring intra-household allocation may severely bias our views regarding inequality.

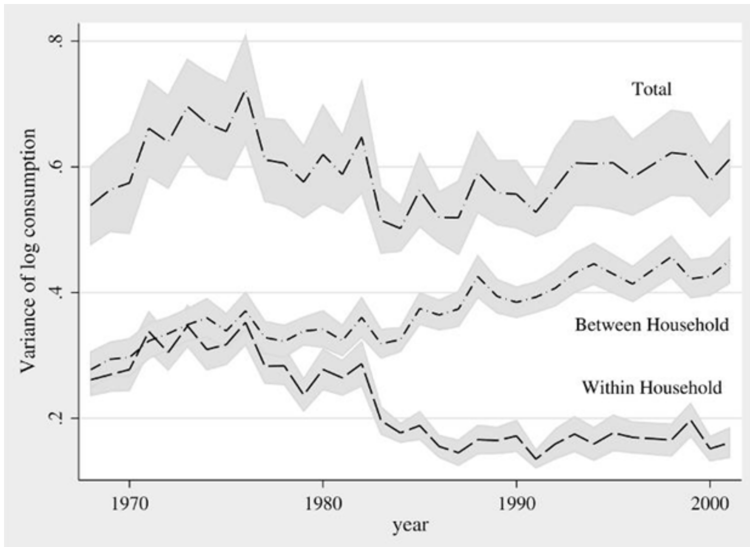

Figure 1: Trends in the variance of log consumption, UK. Source: Lise and Seitz (2011).

All this strongly suggests, at the very least, that much more attention should be paid to intrahousehold inequality, from both a theoretical and an empirical viewpoint. Analyzing intrahousehold inequality, however, raises a host of specific problems. Some are of a conceptual nature. A large fraction of household expenditures relate to public commodities - i.e., goods that are jointly consumed by the household, without exclusion restrictions; moreover, in many cases these public commodities are internally produced within the household. Spouses may have different preferences regarding public goods; therefore, the fraction of household expenditures devoted to public consumption has a potentially important impact on intrahousehold inequality, that cannot be disregarded. Similar questions arise for intrahousehold production, with the additional twist that time spent by each spouse should also be taken into account. How should such public productions and consumptions be taken into account in our inequality measures? While the impact of public goods on inequality is by no means a new problem, it is particularly stringent in our context, if only because public goods and domestic production are among the main (economic) reasons for the existence of the household.

As we shall see, these conceptual issues affect the standard notion of inequality in two ways. Besides shedding a new light on its measurement, they also revive some old discussions about its foundations. In particular, the role of public goods raises questions about which type on inequality should we concentrate on: income? (private) consumption? utility? The problem is far from innocuous: in the presence of public goods, it is relatively easy to generate examples in which a change in prices and incomes results in a decrease in a person's private consumption and an increase in the spouse's, whereas utilities evolve in the opposite way (welfare declines for the person whose private consumption increases and conversely). In such a context, the impact of the change on intrahousehold inequality is not clearly defined – it all depends on what exactly we are interested in.

Empirical problems are equally challenging. As always in economics, preferences are not directly observed, and have to be recovered from observable data (demand, labor supply). But, in addition, the allocation of resources within the household cannot (in general) be directly observed; it has to be recovered from the household's (aggregate) behavior. It follows that when deciding which aspect of inequality should be considered, one cannot abstract from identification issues: there is little interest in concentrating on a notion that is not identifiable in practice. An interesting paradox, in this respect, is provided by a standard result of household economics – namely, that in some circumstances, a continuum of different models generate the same observable behavior (so they are observationally indistinguishable). In some cases, these models correspond to different intrahousehold allocations of resources, but to the same allocation of utility (in the language of the theory, the indeterminacy is welfare-irrelevant). In other words, the main justification for concentrating on inequality in income or consumption rather than in utility – namely, the fact that the former are observed, but not the latter – is sometimes partially reversed.

These questions obviously arise whenever inequality is assessed on a utilitarian basis. However, even an alternative approach in term of capabilities could hardly disregard them. Issues related to individual preferences for public goods would be less problematic in that case; what matters, from a capabilities perspective, is more an individual's potential access to the public goods than the utility the individual actually derives from their consumption. But the difficulties in recovering individual private consumptions (especially when it comes to nutrition or other fundamental needs) would become all the more crucial. All in all, the problems raised by intrahousehold allocation should be central to any analysis of inequality, even though specific aspects may be more damaging for some approaches than for others.

What recent developments in the literature clearly indicate, however, is that while these problems are serious, they are by no means insuperable. Although intra-household allocation is not (fully) observable, it can be recovered using specific, identifying assumptions that will be discussed later; that is the path borrowed by Lise and Seitz, but also by Chiappori, Fortin and Lacroix (2002), Dunbar, Lewbel and Pendakur (2013), Browning, Chiappori and Lewbel (2013) and many others in the literature. Clear progress has been made on this front over the last decades. One goal of the present chapter is to briefly review these advances.

A first step is to adopt an explicit model of household decision making that clarifies the notion of inequality within the household. Obviously, such models must explicitly recognize that household members each have their own preferences - if only because omitting individuals does not seem a promising way of analyzing inequality between them. An additional requirement is empirical tractability. To be usable, a model of household behavior should fulfill a double requirement: testability (i.e., it should generate a set of empirically testable restrictions that fully characterize the model, in the sense that any given behavior is compatible with the model if and only if these conditions are satisfied) and identifiability (it should be feasible, possibly under additional assumption, to recover the structure of the model – in our case, individual preferences and the decision process – from the sole observation of household behavior). Lastly, the model should provide (or be compatible with) an 'upstream' theory of the generation of intrahousehold inequality; i.e., we need to explain, and ideally predict, how the intrahousehold distribution of resources - and ultimately of power - responds to changes in the household's socio-economic environment.

Most of the recent advances use one particular class of models, based on the collective approach.[2] While other (non-unitary) perspectives have been adopted in the literature, none of the alternatives has (so far) convincingly addressed the double requirement of testability and identifiability just evoked.

1.2 Modeling household decision making: the collective model

The basic axiom of the collective approach is Pareto efficiency: whatever decision the household is making, no alternative choice would have been prefered by *all* members. While this assumption is undoubtedly restrictive, its scope remains quite large. It encompasses as particular cases many models that have been proposed in the literature, including:

- 'unitary' models, which posit that the household behaves like a single

[2] For a more detailed presentation, the reader is refered to Browning, Chiappori and Weiss (2013)

decision maker; this includes simple dictatorship (possibly by a 'benevolent patriarch', as in Becker, 1974) to the existence of some household welfare function (as in Samuelson 1956),

- models based on cooperative game theory, and particularly bargaining theory (at least in a context of symmetric information), as pioneered by Manser and Brown (1980) and McElroy and Horney (1981),
- model based on market equilibrium, as analyzed by Grossbard-Shechtman (1993), Gersbach and Haller (2001), Edlund, and Korn (2002) and others.
- more specific models, such as Lundberg and Pollak's 'separate spheres' (1993) framework.

On the other hand, the collective framework excludes models based on non cooperative game theory (at least in the presence of public good), such as those considered by Ulph (2006), Browning, Chiappori and Lechene (2009), Lechene and Preston (2011) and many others, as well as models of inefficient bargaining a la Basu (2006).

The efficiency assumption is standard in many economic contexts, and has often been applied to household behavior. Still, it needs careful justification. Within a static context, this assumption amounts to the requirement that married partners will find a way to take advantage of opportunities that make both of them better off. Because of proximity and durability of the relation, both partners are in general aware of the preferences and actions of each other. They can act cooperatively by reaching some binding agreement. Enforcement of such agreements can be achieved through mutual care and trust, by social norms and by formal legal contracts. Alternatively, the agreement can be supported by repeated interactions, including the possibility of punishment. A large literature in game theory, based on several 'folk theorems', suggests that in such situations, efficiency should prevail.[3] At the very least, efficiency can be considered as a natural benchmark.

Another potential issue with a collective approach to inequality issues is of a more conceptual nature. By definition, the collective approach is axiomatic; it assumes specific properties of the outcome (efficiency), and leaves aside the specific process by which this outcome has been generated. It has sometimes been argued that one should judge differently situations that generate the same allocations (and the same utility levels) but which are reached by different processes. In that case, the collective approach has to be further specialized - and this may be (and has been) done in several directions.[4]

Finally, an obvious but crucial advantage of the collective model is that it has been by now fully characterized. We have a set of necessary and sufficient conditions for a demand function to stem from a collective framework (Chiappori, Ekeland 2006); and exclusion restrictions have been derived under which individual preferences and the decision process (as summarized by the Pareto weights) can be recovered from the sole observation of household behavior ((Chiappori, Ekeland 2009a,b). To the best of our knowledge, this is the only model of the household for which similar results have been derived.[5]

The next section describes the basic model. We then discuss the conceptual issues linked with intrahousehold inequality, first in the case where all commodities are privately consumed, then in the presence of public goods, finally for the case of domestic production. Finally, we discuss issues related to identification.

[3]Note, however, that folks theorems essentially apply to infinitely repeated interactions.

[4]See Browning, Chiappori and Weiss (2014).

[5]Browning, Chiappori and Lechene (2009) and Lechene and Preston (2011) provide a set of necessary conditions for non cooperative models. However, whether these conditions are sufficient is not known; moreover, no general identification result has been derived so far.

2 The collective model: concepts, definitions, axioms

In what follows, we consider a K-person household that can consume several commodities; these may include standard consumption goods and services, but also leisure, future or contingent goods, etc. Formally, N of these commodities are publicly consumed within the household. The market purchase of public good j is denoted Q_j; the N-vector of public goods is given by Q. Similarly, private goods are denoted q_i with the n-vector q. Each private goods bought is divided between the members so that member a $(a = 1, ..., K)$ receives q_i^a of good i, with $\sum_a q_i^a = q_i$. The vector of private goods that a receives is q^a, with $\sum_a q^a = q$. An *allocation* is a $N + Kn$-vector $\left(Q, q^1, ..., q^K\right)$. The associated market prices are given by the N-vector P and the n-vector p for public and private goods respectively.

We assume that each married person has her or his own preferences over the allocation of family resources. The most general version of the model would consider utilities of the form $U^a\left(Q, q^1, ..., q^K\right)$, implying that a is concerned directly with all members' consumptions. Here, however, tractability requires additional structure. In what follows, we therefore assume that preferences are of the *caring* type. That is, each individual a has a felicity function $u^a\left(Q, q^a\right)$; and a's utlity takes the form:

$$U^a\left(Q, q^1, ..., q^K\right) = W^a\left(u^1\left(Q, q^1\right), ..., u^K\left(Q, q^K\right)\right), \tag{1}$$

where $W^a(., .)$ is a monotone increasing function. The weak separability of these 'social' preferences represents an important moral principle; a is indifferent between bundles $\left(q^b, Q\right)$ that b consumes whenever b is indifferent. In this sense caring is distinguished from paternalism. Caring rules out direct externalities between members because a's evaluation of her private consumption q^a does not depend directly on the private goods that b consumes.

Lastly, a particular but widely used version of caring is *egotistic* preferences, whereby members only care about their own (private and public) consumption; then individual preferences can be represented by felicities (i.e., utilities of the form $u^a(Q, q^a)$).[6] Note that such egotistic preferences *for consumption* do not exclude non economic aspects, such as love, companionship or others. That is, a person's utility may be affected by the *presence* of other persons, but not by their *consumption*. Technically, the 'true' preferences are of the form $F^a(u^a(Q, q^a))$, where F^a may depend on marital status and on the spouse's characteristics. Note that the F^as will typically play a crucial role in the decision to marry and in the choice of a partner. However, it is irrelevant for the characterization of married individuals' preferences over consumption bundles.

Efficiency has a simple translation - namely, the household behaves as if it was maximizing a weighted sum of utilities of its members. Technically, the program is thus (assuming egotistic preferences):

$$\max_{(Q, q^1, \ldots, q^K)} \sum_a \mu^a u^a(Q, q^a) \tag{P}$$

under the budget constraint:

$$\sum_i P_i Q_i + \sum_j p_j \left(q_j^1 + \ldots + q_j^K\right) = y^1 + \ldots + y^K = y$$

where y^a denotes a's (non labor) income. Here, μ^a is the *Pareto weight* of member a; one may, for instance, adopt the normalization $\sum_a \mu^a = 1$. In the particular case where μ^a is constant, the program above describes a *unitary* model, since household behavior is described by the maximization of some

[6] Throughout the chapter, we assume, for convenience, that utility functions $u^a(.)$, $a = 1, K$ are continuously differentiable and strictly quasi-concave.

(price independent) utility. In general, however, μ^a may vary with prices and individual incomes; the maximand in (P) is therefore price-dependent, and we are not in a unitary framework in general.

This program can readily be extended to caring preferences - one must simply replace $u^a(Q, q^a)$ with $W^a\left(u^1\left(Q, q^1\right), ..., u^K\left(Q, q^K\right)\right)$ in (P). In what follows, however, (P) plays a very special role, because any allocation that is efficient for caring preferences must be efficient for the underlying, egotistic felicities, as stated by the following result:

Proposition 1 *Assume that some allocation is Pareto-efficient for the caring utilities* $W^1, ..., W^K$*. Then it solves (P) for some* $\left(\mu^1, ..., \mu^K\right)$*.*

Proof. *Assume not, then there exists an alternative allocation that gives a larger value to* u^a *for* all $a = 1, ..., K$*. But then that allocation also gives a higher value to all* W^a*s, a contradiction.* ■

The converse is not true, because a very unequal solution to (P) may fail to be Pareto efficient for caring preferences: transferring resources from well endowed but caring individuals to the poorly endowed ones may be Pareto improving. Still, any property of the solutions to a program of the form (P) must be satisfied by any Pareto-efficient allocation with caring preferences.

A major advantage of the formulation (P) is that the Pareto weight has a natural interpretation in terms of respective decision *powers*. The notion of 'power' in households may be difficult to define formally, even in a simplified framework like ours. Still, it seems natural to expect that when two people bargain, a person's gain increases with the person's power. This somewhat hazy notion is captured very effectively by the Pareto weights. Clearly, if μ^a in (P) is zero then a has no say on the final allocation, while if μ^a is large then a effectively gets her way. A key property of (P) is precisely that increasing μ will result in a move along the Pareto set, in the direction of higher utility for a. If we restrict ourselves to economic considerations, we may thus consider that the Pareto weight μ^a 'reflects' a's power, in the sense that a larger μ^a corresponds to more power (and better outcomes) being enjoyed by a.

If $\left(\bar{Q}(p,P,y), \bar{q}^1(p,P,y), ..., \bar{q}^K(p,P,y)\right)$ denotes the solution to (P), we define the *collective indirect utlity* of a as the utility reached by a at the end of the decision process; formally:

$$V^a(p,P,y) = u^a\left(\bar{Q}(p,P,y), \bar{q}^a(p,P,y)\right)$$

Note that, unlike the unitary setting, in the collective framework a member's collective indirect utility depends not only on the member's preferences but also on the decision process (hence the adjective 'collective'). This notion is crucial for welfare analysis, as we shall see below.

Finally, an important concept is the notion of *distribution factors*. A distribution factor is any variable that (i) does not affect preferences or the budget constraint, but (ii) may influence the decision process, therefore the Pareto weights. Think, for instance, of a bargaining model in which the agents' respective threat points may vary. A change in the threat point of one member will typically influence the outcome of the bargaining process, even if the household's budget constraint is unaffected. In particular, several tests of household behavior consider the income pooling property. The basic intuition is straightforward: in a unitary framework, whereby households behave like single decision makers (and maximize a unique, income-independent utility), only total household income should matter. Individual contributions to total income have no influence on behavior: they are pooled in the right hand side of the household's budget constraint. For instance, paying a benefit to the wife rather than the husband cannot possibly impact the household's demand. As we shall see later, this property has been repeatedly rejected by the data. The most natural interpretation for such rejections (although not the only one) is that individual incomes may impact the decision process (in addition to their aggregate contribution to the budget constraint). Technically, if $\left(y^1, ..., y^K\right)$ is the vector of individual incomes and $y = \sum_a y^a$, while total income y is *not* a distribution factor (it enters the budget constraints), the $(K-1)$ ratios $y^1/y, ..., y^{K-1}/y$ are.[7] Of course, such a setting by no means imply that each individual consumes exactly his or her income. On the contrary, empirical evidence strongly suggests that transfers between family members are paramount. Whether these transfers are progressive or regressive - i.e., whether they increase or decrease intra household inequality - is in the end an empirical question; whether it can be answered ultimately depends on the extend to which these transfers can be either observed or identified, an issue to which the end of this survey is dedicated.

In what follows, the vector of distribution factors will be denoted $z = (z_1, ..., z_S)$; Pareto weights and collective indirect utilities, therefore, have the general form $\mu^a(p, P, y, z)$ and $V^a(p, P, y, z)$.

3 Modeling household behavior: the collective model

3.1 Private goods only: the sharing rule

We first consider a special case in which all commodities are privately consumed. Then the household can be considered as a small economy without externalities or private goods. From the second welfare theorem, any Pareto efficient allocation can be decentralized by adequate transfers; formally, we have the following result:

[7] In practice, distribution factors must be *uncorrelated* with preferences, which, in the case of individual incomes, can generate subtle exogeneity problems. See Browning, Chiappori and Weiss (2014) for a detailed discussion.

Proposition 2 *Assume an allocation* $\left(\bar{q}^1, ..., \bar{q}^K\right)$ *is Pareto efficient. Then there exists* K *non-negative functions* $\left(\rho^1, ..., \rho^K\right)$ *of prices, total income and distribution factors, with* $\sum_k \rho^k (p, y, z) = y$*, such that agent* a *solves*

$$\max_{q^a} u^a (q^a) \tag{D}$$

under the budget constraint

$$\sum_{i=1}^{n} p_i q_i^a = \rho^a$$

Conversely, for any non-negative functions $\left(\rho^1, ..., \rho^K\right)$ *such that* $\sum_k \rho_k (p, y, z) = y$*, an allocation that solves (D) for all* a *is Pareto-efficient.*

In words: in a private goods setting, any efficient decision can be described as (or as if) a two-stage process. In the first stage, agents jointly decide on the allocation of household aggregate income y between agents (and agent a gets ρ^a); in stage two, agents freely spend the share they have received. The decision process (bargaining, for instance) takes place in the first stage; its outcome is given by the functions $\left(\rho^1, ..., \rho^K\right)$, which are called the *sharing rule* of the household. From a welfare perspective, there exists a one-to-one, increasing correspondence between Pareto weights and the sharing rule, at least when the Pareto set is strictly convex: when prices and incomes are constant, increasing the weight of one individual (keeping the other weights unchanged) always results in a larger share for that individual and conversely. Finally, the collective, indirect utility takes a simple form; namely:

$$V^a (p, y) = v^a (p, \rho^a (p, y))$$

where v^a is the standard, indirect utility of agent a. We therefore have the following result:

Proposition 3 *When all commodities are privately consumed, then for any given price vector there exists a one-to-one correspondence between the sharing rule and the indirect utility*

In particular, a member's collective indirect utility can be directly computed from the knowledge of that person's preferences and sharing rule; given the preferences, the sharing rule is a sufficient statistic for the entire decision process.

Regarding the issue of intrahousehold inequality, the key remark is that the sharing rule contains all the information required: since all agents face the same prices, the sharing rule fully summarizes intrahousehold allocation of resources. As such, it is directly relevant for intrahousehold inequality. Specifically, let $I(y_1, ..., y_n)$ be some inequality index (as a function of individual incomes). Then the intrahousehold index of inequality is:

$$I_I(p, y) = I\left(\rho^1(p, y), ..., \rho^K(p, y)\right)$$

3.2 Public and private commodities

Convenient as the previous notion may be, it still relies on a strong assumption - namely that all commodities are privately consumed. Relaxing this assumption is obviously necessary, if only because the existence of public consumption is one of the motives of household formation.

Different notions have been considered in the literature. The notion of *conditional sharing rule* (CSR), initially introduced by Blundell, Chiappori and Meghir (2005), refers to a two-stage process, whereby in stage one the household decides the consumption of public goods and the distribution of remaining income between members, while in stage two members all spend their alloted amount on private consumptions so as to maximize individual utility *conditional on* the level of public consumption decided in stage 1. As before, any efficient decision can be represented as stemming from a two stage process of that type. The converse, however, is not true: for any given level of public consumptions, almost all conditional sharing rules lead to inefficient allocations. Moreover, the monotonic relationship between sharing rule and Pareto weights is lost. In particular, increasing a's weight does not necessarily result in a larger value for a's conditional sharing rule; the intuition being that more weight to one agent may result in a different allocation of public expenditures, which may or may not result in an increase in the agent's private consumption. Lastly, and more importantly for our purpose, the conditional sharing rule may give a biased estimate of intrahousehold inequality, because it simply disregards public consumption. That this pattern could be problematic is easy to see. Assume that one spouse (say the wife) cares a lot for a public good, while her husband cares very little. If the structure of household demand entails a significant fraction of expenditures being devoted to that public good, one can expect this pattern to have an impact on any inequality measure within the household. Disregarding public consumption altogether is therefore not an adequate approach.

A second approach relies on an old result in public economics, stating that in the presence of public goods, any efficient allocation can be decentralized using personal (or Lindahl) prices for the public good. This result establishes a nice duality between private and public goods: for the former, agents face identical prices and purchase different quantities (the sum of which is the household's aggregate demand), whereas for the latter the quantity is the same for all but prices are individual-specific (and add up to market prices).[8] Again, the household behaves as if it was using a two stage process. In stage one, the household chooses a vector of individual prices for the public goods and an allocation of total income between members; in stage two members all spend their income

[8] See Chiappori and Ekeland (2009b) for a general presentation. For applications, see for instance Donni (2009), and Cherchye et al (2007) for a revealed preferences perspective.

on private and public consumptions under a budget constraint entailing their Lindahl prices. Formally, member a solves

$$\max_{q^a} u^a (Q, q^a) \tag{DP}$$

under the budget constraint

$$\sum_{i=1}^{n} p_i q_i^a + \sum_{j=1}^{N} P_j^a Q_j^a = \rho^{*a}$$

where P_j^a is the Lindahl price of good j for agent a. The vector $\rho^* = (\rho^{*1}, ..., \rho^{*K})$, with $\sum_a \rho^{*a} = y$, defines a *generalized sharing rule* (GSR).

From an inequality perspective, this notion raises interesting issues. One could choose to adopt ρ^* as a description of intrahousehold inequality; indeed, agents now maximize utility under a budget constraint in which ρ^* describes available income. In particular, ρ^* is a much better indicator of the distribution of resources than the conditional sharing rule $\tilde{\rho}$, because it takes into account both private and public consumptions.

However, the welfare of agent a is *not* fully described by ρ^{*a}; one also needs to know the vector P^a of a's personal prices. Technically, the collective indirect utility of a is:

$$V^a (p, P, y, z) = v^a (p, P^a, \rho^{*a} (p, P, y, z))$$

which depends on both ρ^{*a} and P^a. This implies that the sole knowledge of the GSR is not sufficient to recover the welfare level reached by a given agent, even if her preferences are known; indeed, one laso need to know the prices, which depend on *all* preferences. In particular, we believe that the level of inequality within the household cannot be analyzed from the sole knowledge of the generalized sharing rule. Agents now face different personal prices, and this should be taken into account. Of course, this conclusion was expected; it simply reflects a basic but crucial insight - namely that if agents 'care differently' about the public goods (as indicated by personal prices, which reflect individual marginal willingnesses to pay), then variations in the quantity of these public goods have an impact on intrahousehold inequality.

Finally, Chiappori and Meghir (2014) have recently proposed the concept of *Money Metric Welfare Index* (MMWI). Formally, the MMWI of agent a, $m^a(p, P, y, z)$, is defined by:

$$v^a(p, P, m^a(p, P, y, z)) = V^a(p, P, y, z)$$

Equivalently, if c^a denotes the expenditure function of agent a,then:

$$m^a(p, P, y, z) = c^a(p, P, V^a(p, P, y, z))$$

In words, m^a is the monetary amount that agent a would need to reach the utility level $V^a(p, P, y)$, *if she was to pay the full price of each public good* (i.e., if she faced the price vector P instead of the personalized prices P^a). The basic intuition is simple enough. The index is defined as the monetary amount that would be needed to reach the same utiity level, at some reference prices; a natural benchmark is to use the current market price for all goods, private and public. In particular, there exists a direct relationship between the MMWI and the standard notion of *equivalent income*,[9] although to the best of our knowledge, equivalent income has exclusively been applied so far to private goods. Both approaches rely on the notion that referring to a common price vector can facilitate interpersonal comparisons of welfare.

Unlike the GSR, the Money Metric Welfare Index fully characterizes the

[9]See f.i. the Chapter by Fleurbaey, Decancq and Schokkaert in this *Handbook*.

utility level reached by the agent. That is, knowing an agent's preferences, there is a one-to-one relationship between her utility and her MMWI, and this relationship does *not* depend on the partner's characteristics. In the pure private goods case, the MMWI coincides with the sharing rule; it generalizes this notion to a general setting without losing its main advantage, namely the one-to-one relationship with welfare. Finally, it can readily be extended to allow for labor supply and domestic production; the reader is referred to Chiappori and Meghir (2014) for a detailed presentation.

3.3 An example

The previous concepts can be illustrated on a very simple example, borrowed from Chiappori and Meghir (2014). Assume two agents a and b, two commodities - one private q, one public Q - and Cobb-Douglas preferences:

$$\begin{aligned} u^a &= \frac{1}{1+\alpha}\log q^a + \frac{\alpha}{1+\alpha}\log Q \\ u^b &= \frac{1}{1+\beta}\log q^b + \frac{\beta}{1+\beta}\log Q \end{aligned}$$

corresponding to the indirect utilities:

$$\begin{aligned} v^a &= \log y - \frac{\alpha}{1+\alpha}\log P - \log(1+\alpha) + \frac{\alpha}{1+\alpha}\log\alpha \\ v^b &= \log y - \frac{\beta}{1+\beta}\log P - \log(1+\beta) + \frac{\beta}{1+\beta}\log\beta \end{aligned}$$

Let μ be b's Pareto weight; then the couple's consumption is given by:

$$\begin{aligned} q^a &= \frac{1}{(1+\alpha)(1+\mu)}y, q^b = \frac{\mu}{(1+\beta)(1+\mu)}y \\ \text{and } Q &= \frac{\alpha(1+\beta)+\mu\beta(1+\alpha)}{(1+\alpha)(1+\beta)(1+\mu)}\frac{y}{P} \end{aligned}$$

generating utilities equal to:

$$
\begin{aligned}
V^a &= \log y - \frac{\alpha}{1+\alpha}\log P - \log\left((1+\alpha)(1+\mu)\right) + \frac{\alpha}{1+\alpha}\log\left(\frac{\alpha(1+\beta)+\mu\beta(1+\alpha)}{1+\beta}\right) \\
V^b &= \log y - \frac{\beta}{1+\beta}\log P - \log(1+\beta)(1+\mu) + \frac{1}{1+\beta}\log\mu + \frac{\beta}{1+\beta}\log\left(\frac{\alpha(1+\beta)+\mu\beta(1+\alpha)}{1+\alpha}\right)
\end{aligned}
$$

In this context, straightforward calculations allow to see that:

1. The conditional sharing rule coincides with private consumption:

$$\tilde{\rho}^a = \frac{1}{(1+\alpha)(1+\mu)}y, \tilde{\rho}^b = \frac{\mu}{(1+\beta)(1+\mu)}y$$

2. Lindahl prices are

$$
\begin{aligned}
P^a &= \frac{\alpha(1+\beta)}{\alpha(1+\beta)+\mu\beta(1+\alpha)}P \\
P^b &= \frac{\mu\beta(1+\alpha)}{\alpha(1+\beta)+\mu\beta(1+\alpha)}P
\end{aligned}
$$

and the generalized sharing rule is

$$
\begin{aligned}
\rho^{*a} &= \frac{y}{1+\mu} \\
\rho^{*b} &= \frac{\mu y}{1+\mu}
\end{aligned}
$$

3. The two MMWIs are given by:

$$
\begin{aligned}
m^a &= \left(\frac{\alpha(1+\beta)+\mu\beta(1+\alpha)}{\alpha(1+\beta)}\right)^{\frac{\alpha}{1+\alpha}} \frac{y}{1+\mu} = \left(\frac{\alpha(1+\beta)+\mu\beta(1+\alpha)}{\alpha(1+\beta)}\right)^{\frac{\alpha}{1+\alpha}} \rho^{*a} \\
m^b &= \left(\frac{\alpha(1+\beta)+\mu\beta(1+\alpha)}{\mu\beta(1+\alpha)}\right)^{\frac{\beta}{1+\beta}} \frac{\mu y}{1+\mu} = \left(\frac{\alpha(1+\beta)+\mu\beta(1+\alpha)}{\mu\beta(1+\alpha)}\right)^{\frac{\beta}{1+\beta}} \rho^{*b}
\end{aligned}
$$

Assume, now, that $\mu = 1$ but agents have different preferences for the public good. For instance, $\alpha = 2$ while $\beta = .5$, implying that the wife (husband) puts two third of the weight on the public (private) consumption. In this setting, we can analyze intrahousehold inequality using three possible indicators.

1. If we concentrate on private consumption (or equivalently on the conditional sharing rule), we find that

$$\tilde{\rho}^a = \frac{1}{6}y, \quad \tilde{\rho}^b = \frac{1}{3}y$$

and we conclude that member b is much better off than a.

2. This conclusion is clearly unsatisfactory, because it disregards the fact that half the budget is spent on the public good, which benefits a more than b. Indeed, the GSR is

$$\rho^{*a} = \frac{y}{2} = \rho^{*b}$$

and we conclude that for this indicator, the household is perfectly equal: the benefits of public expenditures exactly compensate differences in private consumptions.

3. The later conclusion is however too optimistic, since it omits the fact that a 'pays' twice as much for the public good than b does (here, $P^a = \frac{2}{3}P$ while $P^b = \frac{1}{3}P$). Taking this last aspect into account, the respective

MMWIs are:

$$m^a = .655y, \quad m^b = .72y$$

Again, b is better off than a (although by much less than with the first measure). In addition, one may note that

$$m^a + m^b = 1.375y$$

Individual MMWIs add up to more than total income, reflecting the gain generated by the publicness of one commodity.

3.4 Domestic production

Finally, the previous analysis can readily be extended to domestic production. Here, we only consider the case where all commodities are privately consumed; for a more general presentation along similar lines, the reader is referred to Chiappori and Meghir (2014). The household production technology is thus described by a production function that gives the possible vector of outputs $q = f(x, \tau)$ that can be produced given a vector of market purchases x and the time $\tau = (\tau^a, a = 1, K)$ spent in household production by each of the members.

We first disregard the time spent by each member on domestic production. This setting is thus identical to the general model of household production of Browning, Chiappori and Lewbel (2003).[10] Pareto efficiency translates into the program:

$$\max \sum \mu^a u^a (q^a)$$

[10]For empirical applications, these authors use a linear technology a la Barten.

$$\begin{aligned} \sum_a q_i^a &= f_i\left(x^i\right), \\ p'\left(\sum_i x^i\right) &= y \end{aligned}$$

where

$$\begin{aligned} q^a &= \left(q_i^a\right), i = 1, n \\ x^i &= \left(x_j^i\right), j = 1, k, \end{aligned}$$

As before, this program can be decentralized, although decentralization now requires specific (shadow) prices for the produced goods. Specifically, let η_i, λ be the respective Lagrange multipliers of the production constraints in (2), and define

$$\pi_i = \frac{\eta_i}{\lambda}$$

Let $((q^{a*}), a = 1, ..., K, x^*)$ denote the solutions, and define the sharing rule by

$$\rho^a = \pi' q^{a*}$$

Then the program is equivalent to a two stage process, in which q^{a*} solves

$$\max u^a\left(q^a\right)$$

under the budget constraint

$$\pi' q^a = \rho^a$$

and x^* solves the profit maximization problem:

$$\max \sum_i \pi_i f_i \left(x^i\right) - \sum_{i,j} p_j x_j^i$$

or equivalently the cost minimization one:

$$\min p' x$$

$$f(x) = \sum_a q^{a*}$$

In that case, again, individual welfare is adequately measured by the sharing rule.

Extending this model to domestic labor supply is straightforward. The Pareto program is now:

$$\max \sum \mu^a u^a \left(q^a, L^a\right)$$

$$\begin{aligned} \sum_a q_i^a &= f_i \left(x^i, \tau_i\right) \\ p' \left(\sum_i x^i\right) + \sum_a w_a \left(L^a + \sum_i \tau_i^a\right) &= y + \sum_a w_a T = Y \end{aligned}$$

where

$$\tau_i = \left(\tau_i^a\right), a = 1, K$$

Prices for internally produced goods are defined as before; the sharing rule is now:

$$\rho^a = \pi' q^{a*} + w_a L^{a*}, a = 1, K$$

where L^{a*} denotes a's optimal leisure. The program can be decentralized as follows: for each a, (q^{a*}, L^{a*}) solve

$$\max u^a\left(q^a, L^a\right)$$

$$\pi' q^a + w_a L^a = \rho^a$$

and x^*, τ^{a*} solves

$$\max \sum_i \pi_i f_i\left(x^i, \tau_i\right) - \sum_{i,j} p_j x_j^i - \sum_{i,a} w_a \tau_i^a$$

or equivalently:

$$\min \sum_{i,j} p_j x_j^i + \sum_{i,a} w_a \tau_i^a$$

under

$$f_i\left(x^i, \tau_i^a\right) = \sum_a q_i^{a*}, i = 1, n$$

In practice, several variants of this basic framework can be considered, depending on whether the internally produced goods are marketable, and whether (market) labor supplies are at an interior or a corner solution. These technical issues are not without importance. For instance, a standard issue in family economics is whether a change in the respective powers of the various members has an impact on the intrahousehold allocation of domestic work. In the model just described, if the produced commodities are marketable and all individuals work on the market, then the πs and the ws must coincide with market prices and wages; they are therefore exogenous, and individual, domestic labor supplies are fully defined by the program (3.4), which does not depend on Pareto weights. We conclude that, in that case, powers have no impact on domestic work, which is fully determined by efficiency considerations. Clearly, this argument must be modified when either the πs or the ws are endogenous (as will be the case if, respectively, the commodity is not marketable or a person does not participate to the labor market); the reader is referred to Browning, Chiappori and Weiss (2014) for a precise discussion, as well as to the Chapter on Gender Inequality in this Handbook.

4 The determinants of intrahousehold allocation

The second task assigned to theory is to explain the allocation of powers, hence of resources, within the household. As such, it must address issues related to household formation and dissolution, as well as the interaction between the household and its environment - i.e., which external factors may impact the intrahousehold decision process. In what follows, we concentrate on two types of approaches, respectively based on cooperative bargaining and matching or search theory. In a sense, this distinction reflects the classic dichotomy between partial and general equilibrium. Bargaining models analyze, for a given household, how the particular situation of each member may affect the household decision; much emphasis is put on individual 'threat points', generally considered as exogenous. Matching and search models, on the other hand, describe a global equilibrium on the 'market for marriage' as a whole; while the decision process may in some cases entail bargaining (in search models, or in matching with a finite set of agents), the crucial distinction is that the threat points are now *endogenous* - their determination is part of the equilibrium conditions.

4.1 Bargaining models

Any bargaining model requires a specific setting: in addition to the framework described above (K agents, with specific utility functions), one has to define a *threat point* T^a for each individual a. Intuitively, a person's threat point describes the utility level this person could reach in the absence of an agreement with the partner. Typically, bargaining models assume that the outcome of the decision process is Pareto efficient and individually rational, in the sense that individuals never receive less than their threat point. Bargaining theory is used to determine how the threat points influence the location of the chosen point on the Pareto frontier. Clearly, if the point $T = \left(T^1, ..., T^K\right)$ is outside of the Pareto set, then no agreement can be reached, since at least one member would lose by agreeing. However, if T belongs to the interior of the Pareto set so that all agents can gain from the relationship, the model picks a particular point on the Pareto utility frontier. Note that the crucial role played by threat points - a common feature of all bargaining models - has a very natural interpretation in terms of distribution factors. Indeed, *any variable that is relevant for threat points only is a potential distribution factor*. For instance, the nature of divorce settlements, the generosity of single parent benefits or the probability of re-marriage do not directly change a household's budget constraint (as long as it does not dissolve), but may affect the respective threat points of individuals within it. Then bargaining theory implies that they will influence the intrahousehold distribution of power in households and, ultimately, household behavior. Equivalently, one could say that these variables are distribution factors that affect the Pareto weights.

In practice, models based on bargaining must make a number of basic choices. One is the bargaining concept to be used. While most frameworks refer to Nash bargaining, some works either adopt Kalai-Smorodinski or refer to a non cooperative bargaining model. Second, one must choose a relevant threat point. This part is crucial; indeed, a result due to Chiappori, Donni and Komunjer (2010) states that *any* Pareto efficient allocation can be derived as the Nash bargaining solution for an *ad hoc* definition of the threat points.

Hence any additional information provided by the bargaining concepts (besides the sole efficiency assumption) must come from specific hypotheses on the threat points - that is, on what is meant by the sentence: ‘no agreement is reached’. Several ideas have been used in the literature. One is to refer to divorce as the ‘no agreement’ situation. Then the threat point is defined as the maximum utility a person could reach after divorce. Such an idea seems well adapted when one is interested, say, in the effects of laws governing divorce on intrahousehold allocation. It is probably less natural when minor decisions are at stake: choosing who will walk the dog is unlikely to involve threats of divorce.[11] Another interesting illustration would be public policies such as single parent, or the guaranteed employment programs that exist in some Indian states; Kanbur and Haddad (1992) convincingly argued that the main impact of the program was to change the opportunities available to the wife outside marriage.

A second idea relies on the presence of public goods and the fact that non-cooperative behavior typically leads to inefficient outcomes. The idea, then, is to take the non-cooperative outcome as the threat point: in the absence of an agreement, both members provide the public good(s) egotistically, not taking into account the impact of their decision on the other member’s welfare. This version captures the idea that the person who would suffer more from this lack of cooperation (the person who has the higher valuation for the public good) is likely to be more willing to compromise in order to reach an agreement. A variant, proposed by Lundberg and Pollak (1993), is based on the notion of ’separate spheres’. The idea is that each partner is assigned a set of public goods to which they alone can contribute; this is their ‘sphere’ of responsibility or expertise. These spheres are determined by social norms. Then the threats consist of continued marriage in which the partners act non-cooperatively and

[11] An additional difficulty is empirical. The estimation of utility in case of divorce is delicate, since most data sets allow us to estimate (at best) an ordinal representation of preferences, whereas Nash bargaining requires a cardinal representation. See Chiappori (1991)

each chooses independently the level of public goods under their domain.

Finally, it must be reminded that assumptions on threat points tend to be strong, not grounded on strong theoretical arguments, and often not independently testable. This suggests that models based on bargaining should be used parsimoniously and with care.

4.2 Equilibrium models

Alternatively, one can consider the 'market for marriage' as a whole from a general perspective. Two types of models can be found in the literature, that make opposite assumptions on the role of frictions in the matching game. Specifically, models based on matching (with transferable or imperfectly transferable utility, TU and ITU respectively) assume away frictions and consider perfectly smooth markets, while models based on search emphasize the importance of frictions in the emergence of marital patterns. In what follows, and for the sake of brevity, we concentrate on matching models; while search-based approaches use a different technology, their scope and outcomes are largely similar for what we are concerned with here. Moreover, we shall only discuss models based on transferable utility (TU). The non transferable utility (NTU) framework, which assumes away any transfer between members, is not relevant here; and although more general approaches, based on imperfectly transferable utility (ITU), have recently been developed (see Chiappori, 2012), the distinction between TU and ITU can basically be disregarded for our current purpose.

Consider two populations (men and women); each individual is defined by a vector of characteristics, denoted $x \in X$ for women and $y \in Y$ for men; each set is endowed with a finite measure, denoted μ_X and μ_Y respectively. When matched, Mrs. x and Mr. y jointly generate a surplus $s(x, y)$, which can be derived from a more structural framework (e.g., a collective model). A *matching* is defined by (i) a measure μ on the set $X \times Y$, the marginals of which coincide with μ_X and μ_Y, and (ii) two functions $u(x)$ and $v(y)$ such that $u(x) + v(y) = s(x, y)$ on the support of μ. Intuitively, the measure μ defines who marries whom, while the functions determines how the surplus is divided within couples who are matched with positive probability - she gets $u(x)$, he gets $v(y)$. A matching is *stable* if (i) no married person would prefer being single, and (ii) no pair of currently unmarried persons would both prefer forming a new couple. Technically, this is equivalent to:

$$u(x) + v(y) \geq s(x, y) \text{ for } \textit{all } (x, y)$$

The functions $u(x)$ and $v(y)$ are crucial, since they fully determine the intrahousehold inequality. The key feature of matching models is that these functions are *endogenous*. They are determined (or constrained) as part of the equilibrium, and depend on the whole matching game structure; in particular, the allocation within any given couple depends on the entire distribution of characteristics in the two populations. In that sense, the model does provide an endogenous determination of intrahousehold inequality. Note, however, that in this abstract presentation, their exact interpretation is undetermined; depending on the framework, $u(x)$ can be a monetary amount, the consumption of some commodity, or the utility generated by the consumption of bundles of private and public commodities. For instance, the simple framework used by Chiappori and Weiss (2007) consider an economy with two commodities, one private and one public within the household, and agents with Cobb-Douglas preferences $u^a = q^a Q$; x and y are one-dimensional and denote male and female income. In this TU framework, any efficient allocation maximizes the *sum* of utilities; i.e., a (x, y) couple solves

$$\max_{q^1,q^2,Q} \left(q^1+q^2\right) Q \quad \text{under } q^1+q^2+Q=x+y$$

and the surplus $s(x, y)$ is the value of this program, namely $(x+y)^2/4$. Here, $u(x)$ and $v(y)$ are utilities, although there exists a one-to-one correspondance between utilities and transfers (since $Q=(x+y)/2$, we have that $q^1=2u(x)/(x+y)\,, q^2=2v(y)/(x+y)$).

From a mathematical point of view, a basic result states that if a matching is stable, then the corresponding measure maximizes total surplus over the set of measures whose the marginals coincide with $\mu_X$and μ_Y. That is, the measure μ must solve:

$$\max_{\mu} \int_{X\times Y} s(x,y)\, d\mu(x,y) \tag{2}$$

under the marginal conditions. This maximization problem is linear in its unknown μ. Therefore, it admits a dual, which can be written as:

$$\min_{u,v} \int_X u(x)\, d\mu_X(x) + \int_Y v(y)\, d\mu_Y(y)$$

under the constraints

$$u(x)+v(y) \geq s(x,y) \;\; \forall (x,y) \tag{3}$$

Here, functions u and v are the dual variables of the program. But, crucially, they can be interpreted as describing the utility reached by each individual at the optimal matching; in particular, they define the allocation of surplus between (matched) spouses. Note that conditions (3) of the dual program are exactly the stability conditions (4.2).

From standard, duality results, a solution to the dual exists if and only if the primal has a solution, and the values are then the same. It follows that the existence of a stable match - i.e., of functions $u$and v satisfying (4.2) - boils down to the existence of a solution to the linear maximization problem (2). This allows to establish existence under very general conditions; see for instance Chiappori, McCann and Nesheim (2010).

Regarding uniqueness, if the sets X and Y are finite, then the us and vs are not pinned down, although the equilibrium conditions generate constraints. However, with continuous, atomless populations, the functions are in general fully determined by the equilibrium conditions. The intuition is straightforward: in the continuous case, each individual has almost perfect substitutes, and (local) competition determines exactly the surplus sharing that must exist at equilibrium. Finally, stochastic versions of these models can be considered, in which some of the individual characteristics are unobserved (to the econometrician); see for instance the recent survey by Chiappori and Salanié (2013).

5 Identification

While the conceptual tools just presented help clarifying some of the issues involved, their empirical content must be very carefully considered. As said before, there is no point putting much emphasis on a concept that cannot possibly be identified from existing data. This section summarizes the main results obtained on this issue over the last two decades; for a detailed presentation, the reader is referred to Chiappori and Ekeland (2009).

We divide the presentation into three subsection. One considers the 'pure' identification problem. Assume that the entire demand function of a household can be observed; what can be recovered from such data (and such data only)? Next, we introduce additional identifying assumptions; broadly speaking, these postulate a relationship between an individual's preferences as a single and as part of a household; in other words, we admit that some information about spouses' utilities can be derived from the observation of the behavior of single persons. Lastly, we introduce a general, market-wide perspective, and ask whether (and how) equilibrium conditions on the marriage market can help identifying the intrahousehold allocation process.

5.1 'Pure' identification in the collective model

Identification issues in the collective model have been extensively studied during the recent years; the interested reader is refered to Chiappori and Ekeland (2009a, b) for an exhaustive presentation. In what follows, we briefly summarize some key findings.

We start with the basic framework described above, assuming egotistic preferences of the type $u^a(Q, q^a)$; also, for the sake of brevity, we assume only two persons ('spouses') in the household, although the generalization to any number is straightforward. In what follows, we assume that we observe the household's 'aggregate' demand, i.e. the vector $(q, Q) \in \mathbb{R}^{n+N}$ (where $q_i = \sum_a q_i^a, i = 1, ..., n$) as a function of prices (p, P) and total income y, plus possibly a vector of distribution factors z. Remember that the *collective indirect utility* of agent a is defined as the utility level a will reach at the end of the decision process, as a function of (p, P, y, z).

5.1.1 Main identification result

Assume, first, that we observe the demand function of some household. This demand is aggregated at the household level; i.e., what we observe is the household's total demand for any private commodity, together with its demand for public goods. However, in general, we are not able to observe the internal allocation of the private goods between household member. When is this information sufficient to recover the underlying structure - i.e., preferences and the decision process (as summarized by the Pareto weights)?

A first answer is provided by a result due to Chiappori and Ekeland (2009a). It states that, generically, all what is needed is one exclusion restriction per agent; i.e., for any agent a, there should be some commodity that a does *not* consume (and which does not enter a's egoistic utility). Then the *local* knowledge of the household demand allows to exactly (locally) identify each agent's collective indirect utility, irrespective of the number of private and public goods. Formally:

Theorem 4 *(Chiappori, Ekeland 2009) Assume* $N + n \geq 4$. *Consider a point* $\left(\bar{p}, \bar{P}, \bar{y}\right)$ *such that the conditional sharing rule satisfies the condition*

$$\frac{\partial \rho^a}{\partial y}\left(\bar{p}, \bar{Q}, \bar{y}\right) \neq 0, a = 1, 2$$

where $\bar{Q} = Q\left(\bar{p}, \bar{P}, \bar{y}\right)$. *Assume that for each member, there exists at least one good not consumed by this member (but consumed by the other). Then generically there exists an open neighborhood of* $\left(\bar{p}, \bar{P}, \bar{y}\right)$ *on which the indirect collective utility of each member is exactly (ordinally) identifiable from household demand. For any cardinalization of indirect collective utilities, the Pareto weights are exactly identifiable.*

Proof. *For a precise proof, see Chiappori and Ekeland 2009a. The underlying intuition is that if commodity* i *is not consumed by agent* y, *then any impact of its price on that agent's behavior can only operate through the decision process, i.e. the Pareto weights. The resulting conditions, which are reminiscent of separability restrictions in standard consumer theory, are sufficient in general to fully recover the (ordinal) indirect collective utility of each member, as well as, for any choice of cardinalization, the corresponding Pareto weights.* ■

The specific nature of the identification result can be simply illustrated on a Cobb-Douglas example, described below. Before considering it, a few remarks are in order. First, the identification result stated in Theorem 4 is only local. This is important because additional constraints of a global nature (such as non negativity restrictions on consumption), which are not considered in this result, typically provide additional identification power; a precise illustration will be given below. Second, the result does not require distribution factors. Again, the later would allow a stronger identification result. Indeed, Chiappori and Ekeland show that in the presence of distribution factors, the exclusivity requirement can be relaxed; one only need either one excluded good (instead of two) or an *assignable* commodity.[12] Third, identification requires the observation of the household demand as a function of prices and income; in particular, price variations are crucial. While this fact is not surprising - even in standard consumer theory, preferences cannot be recovered from demand without price variations - it has important empirical applications, since data entailing significant (and credibly exogenous) price variations are not easy to find. However, recent approaches relax this requirement by imposing additional structure on the decision process; they will be discussed below.

Fourth, the identification result above is only generic: it may fail to hold in particular cases, although such cases are not robust to 'small variations'. Quite interestingly, one of the situations in which identification does not obtain is the unitary model. To see why, consider program (P) above, and assume that the Pareto weights μ^a are all constant. For one thing, we are in a unitary context: the household maximizes the sum $\sum_a \mu^a u^a (Q, q^a)$, which is a *price- and income-independent* utility. More importantly, Hicks's aggregation theorem applies. If we define U by

[12] A good is assignable when it is consumed by both members, and the consumption of each member is independently observed.

$$U(Q,q) = \max_{\sum_a q^a = q} \sum_a \mu^a u^a (Q, q^a) \qquad \text{(UHicks)}$$

then the household maximizes U under the budget constraint. By standard integration, U can be recovered from the household demand. However, this is not sufficient to identify individual preferences: there exists a continuum of different sets of individual utilities that generate the same U by (UHicks). The paradox, here, is that the unitary model, which used to be the dominant framework for empirical works on household behavior, belongs to the small (actually non generic) class of frameworks for which individual welfare cannot be identified from household demand.

Lastly, it is important to note that what is identified is the *indirect collective utility* of each member. From a welfare perspective, this is the only relevant concept, since it fully characterizes the utility reached by each agent. However, the inequality measures described above require more - namely, an assessment of the intrahousehold allocation of income. We now consider to what extent the latter can be recovered from the indirect collective utility.

5.1.2 Private goods and the sharing rule

We start with the case in which all commodities are private. In that case, the various concepts (conditional sharing rule, generalized sharing rule, money metric welfare index) coincide with the sharing rule, and the collective indirect utility takes the form:

$$V^a(p,y) = v^a(p, \rho^a(p,y))$$

where, as above, v^a is a's indirect utility and ρ is the sharing rule. If we assume that the first (respectively the second) good is exclusively consumed by the second (first) agent, the collective indirect utility of each agent is identified (as always, up to some increasing transform).

Local identification A first result states that the sharing rule is *not* fully identified from the knowledge of the collective indirect utility, at least locally; identification only obtains up to an additive function of the prices of the non exclusive goods. Formally, assume that one observes the functions $(q_1, ...q_n)$ of (p, y), with $p \in \mathbb{R}^n$ and:

$$\begin{aligned} q_1(p,y) &= \chi_1^a(p,\rho(p,y)) \\ q_2(p,y) &= \chi_2^b(p,y-\rho(p,y)) \\ q_i(p,y) &= \chi_i^a(p,\rho(p,y)) + \chi_i^b(y-\rho(p,y)), \quad i=3,...,n \end{aligned} \tag{4}$$

where the functions χ_i^s and ρ are unknown. Then:

Proposition 5 *(Chiappori, Ekeland 2009) Assume* $n \geq 3$ *, and let* $\left(\bar{\chi}_1^a, ..., \bar{\chi}_n^b, \bar{\rho}\right)$ *solve (4). For any other solution* $\left(\chi_1^a, ..., \chi_n^b, \rho\right)$*, there exist a* $\phi : \mathbb{R}^{n-2} \to \mathbb{R}$ *such that:*

$$\begin{aligned} \rho(p,y) &= \bar{\rho}(p,y) + \phi(p_3,...,p_n) \\ \chi_i^a(\rho) &= \bar{\chi}_i^a(\rho - \phi(p_3,...,p_n)) \\ \chi_j^b(\rho) &= \bar{\chi}_j^b(\rho + \phi(p_3,...,p_n)) \end{aligned} \tag{5}$$

Moreover, overidentifying restrictions are generated.

The basic conclusion is thus that the sharing rule is identified up to an additive function, which cannot be pinned down unless either all commodities are assignable or individual preferences are known (for instance, from data on singles). To see why, consider the simple case of three private commodities;

two of these are exclusive (for members a and b respectively), while the third is consumed by both. Individual consumptions of commodity 3 are not observed, and its price is taken as numeraire. In practice, we observe two demand functions q_1^a and q_2^b that satisfy:

$$q_1^a(p_1, p_2, y) = \tilde{q}^a(p_1, \rho(p_1, p_2, y)) \tag{6}$$

$$q_2^b(p_1, p_2, y) = \tilde{q}^b(p_2, y - \rho(p_1, p_2, y)) \tag{7}$$

where $\tilde{q}^s$ denotes the Marshallian demand by person s. Now, for some constant K, define ρ_K, u_K^a and u_K^b by:

$$\begin{aligned} \rho_K(p_1, p_2, y) &= \rho(p_1, p_2, y) + K \\ u_K^a(q_1^a, q_3^a) &= u_K^a(q_1^a, q_3^a - K) \\ u_K^b\left(q_2^b, q_3^b\right) &= u_K^b\left(q_2^b, q_3^b + K\right) \end{aligned}$$

It is easy to check that the Marshallian demands derived from ρ_K, u_K^a and u_K^b satisfy (6) and (7). The intuition is illustrated in Figure 2 in the case of a. Switching from ρ and u^a to ρ_K and u_K^a does two things. First, the sharing rule and the intercept of the budget constraint are shifted downward by K. Second, all indifference curves are also shifted downward by the same amount. When only demand for commodity 1 (on the horizontal axis) is observable, these models are empirically indistinguishable. Lastly, with several, non exclusive goods, this construct is still possible, and the constant may in addition vary with non exclusive prices in an arbitrary way.

Two remarks can be made about this result. One is that the indetermination is not welfare relevant; one can easily check that the different solutions correspond to the same collective indirect utilities for each agent. This is the paradox evoked in introduction. Unlike standard consumer theory, there is no longer an

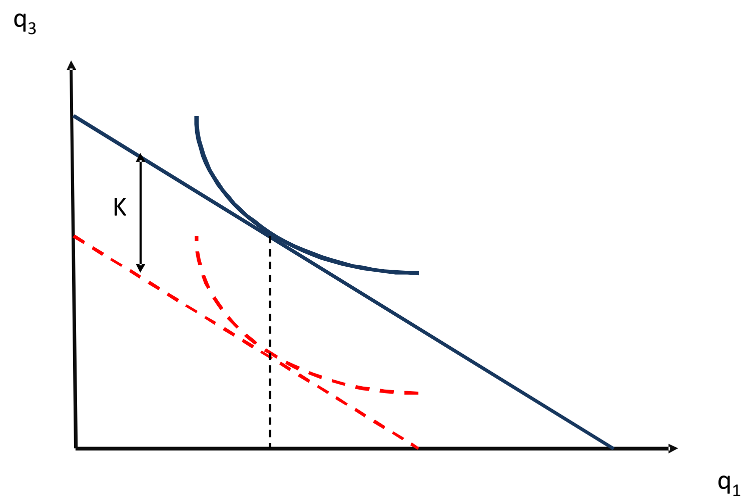

Figure 2:

equivalence between identifying direct and indirect utilities. Indirect utilities are identified as soon as the exclusion restrictions are satisfied; but they may correspond to various, welfare equivalent direct utilities, each of them associated with a specific sharing rule.

Global restrictions The second remark is that the non identification result is only local. In particular, it disregards additional, global restrictions such as non negativity constraints. If these are added, then more precise identification obtains. For instance, consider (5), and add the restrictions that

$$\rho(p, 0) = 0 \quad \forall p$$

which stems from non negativity of consumption at very low income levels. Then ϕ is exactly pinned down:

$$\phi(p_3, ..., p_n) = -\bar{\rho}(p, 0)$$

and additional, overidentifying restrictions are generated (e.g., $\partial\bar{\rho}(p,0)/\partial p_i = 0$ for $i = 1, 2$).

This result should be related to recent work on the estimation of the sharing rules based on a revealed preference approach (see for instance Cherchye et al 2012, Cherchye et al. 2014). Since the revealed preference approach is global by nature, it can generate bounds on the sharing rule, which can actually be quite narrow. In all cases, the global restrictions are generated at one end of the distribution of expenditures, so their use for identifying the sharing rule outside this range should be submitted to the usual caution. Still, they tend to considerably reduce the scope of the non identification conclusion.

5.1.3 Public goods only

We now consider the opposite polar case, in which all commodities (but the exclusive ones) are public. That is, utilities are now of the form:

$$U^a(Q_1, Q_3, ...Q_N) \text{ and } U^b(Q_2, Q_3, ...Q_N)$$

Note that the exclusive commodities 1 and 2 can be considered as either public or private.

In that case, the collective indirect utility has a simple form; namely,

$$\begin{aligned} V^a(P,y) &= U^a(Q_1, Q_3, ...Q_N) \\ V^b(P,y) &= U^b(Q_2, Q_3, ...Q_N) \end{aligned}$$

The crucial remark is that the demands for public goods (as functions of prices and total income) are empirically observed. An important consequence is that, in general, the knowledge of indirect collective utilities is equivalent to that of direct utilities. To see why, normalize y to be 1 (by homogeneity), and take a point at which the Jacobian matrix $D_P(Q_1, Q_2, ...Q_N)$ is of full rank. By the implicit function theorem, we can locally invert the function, thus defining P as a function of Q; but then:

$$\begin{aligned} U^a(Q_1, Q_3, ...Q_N) &= V^a(P_1(Q_1, Q_2, ...Q_N), ..., P_N(Q_1, Q_2, ...Q_N), 1) \\ U^b(Q_2, Q_3, ...Q_N) &= V^b(P_1(Q_1, Q_2, ...Q_N), ..., P_N(Q_1, Q_2, ...Q_N), 1) \end{aligned}$$

which proves identification. In addition, overidentifying restrictions are generated. In particular, we see that, in this context, Lindahl prices for all goods - therefore the MMWIs - are exactly identified. Somewhat paradoxically, the pure public good case appears to be the one in which identification is least problematic...

5.1.4 The general case

Finally, the general case is a direct generalization of the two particular cases just described. The exclusion restrictions described above guarantee identification of the collective indirect utility of each agent. Then the exact intra household allocation is locally identified up to an additive function of the prices of the non exclusive private goods. Moreover, global restrictions (e.g. non negativity) allow exact identification in general. The interested reader is referred to Chiappori and Ekeland (2009 a, b) for detailed statements.

5.1.5 A Cobb-Douglas example

The previous discussions can be illustrated on a simple example, borrowed from Chiappori and Ekeland (2009a). Consider individual preferences of the LES type:

$$U^s(q^s, Q) = \sum_{i=1}^{n} \alpha_i^s \log\left(q_i^s - c_i^s\right) + \sum_{j=n+1}^{N} \alpha_j^s \log\left(Q_j - C_j\right), \quad s = a, b$$

where the parameters α_i^s are normalized by the condition $\sum_{i=1}^{N} \alpha_i^s = 1$ for all s, whereas the parameters c_i^s and C_j are unconstrained. Here, commodities 1 to n are private while commodities $n+1$ to N are public. Also, given the LES form, it is convenient to assume that the household maximizes the weighted sum $\mu U^a + (1-\mu) U^b$, where the Pareto weight μ has the simple, linear form:

$$\mu = \mu^0 + \mu^y y + \mu^z z, \quad s = a, b$$

Household demand The couple solve the program:

$$\begin{aligned} &\max \left(\mu^0 + \mu^y y + \mu^z z\right) \left(\sum_{i=1}^{n} \alpha_i^a \log\left(q_i^a - c_i^a\right) + \sum_{j=n+1}^{N} \alpha_j^a \log\left(Q_j - C_j\right) \right) \\ &+ \left(1 - \left(\mu^0 + \mu^y y + \mu^z z\right)\right) \left(\sum_{i=1}^{n} a_i^b \log\left(q_i^b - c_i^b\right) + \sum_{j=n+1}^{N} \alpha_j^b \log\left(Q_j - C_j\right) \right) \end{aligned}$$

under the budget constraint. Individual demands for private goods are given by:

$$\begin{aligned} p_i q_i^a &= p_i c_i^a + \alpha_i^a \left(\mu^0 + \mu^y y + \mu^z z\right) \left(y - \sum_{i,s} p_i c_i^s - \sum_j P_j C_j \right) \\ p_i q_i^b &= p_i c_i^b + \alpha_i^b \left[1 - \left(\mu^0 + \mu^y y + \mu^z z\right)\right] \left(y - \sum_{i,s} p_i c_i^s - \sum_j P_j C_j \right) \end{aligned}$$

generating the aggregate demand:

$$p_i q_i = p_i c_i + \left[\alpha_i^a \left(\mu^0 + \mu^y y + \mu^z z\right) + \alpha_i^b \left(1 - \left(\mu^0 + \mu^y y + \mu^z z\right)\right)\right] Y \tag{8}$$

and for public goods:

$$P_j Q_j = P_j C_j + \left[\alpha_j^a \left(\mu^0 + \mu^y y + \mu^z z\right) + \alpha_j^b \left(1 - \left(\mu^0 + \mu^y y + \mu^z z\right)\right)\right] Y$$

where $c_i = c_i^a + c_i^b$ and $Y = \left(y - \sum_{i,s} p_i c_i^s - \sum_j P_j C_j\right)$. The household demand is thus a direct generalization of the standard LES, with additional quadratic terms in y^2 and cross terms in yp_i and yP_j, plus terms involving the distribution factor z.

A first remark is that c_i^a and c_i^b cannot be individually identified from group demand, since the latter only involves their sum c_i. As a consequence, the various generalizations of the sharing rule will only be identified up to one additive constant, a result mentioned earlier. Also, the constant is welfare irrelevant; indeed, the collective indirect utilities of the wife and the husband are (up to an increasing transform):

$$\begin{aligned} W^a(p, P, y, z) &= \log Y + \log\left(\mu^0 + \mu^y y + \mu^z z\right) \\ &\quad - \sum_i \alpha_i^a \log p_i - \sum_j \alpha_j^a \log P_j \\ W^b(p, P, y, z) &= \log Y + \log\left(1 - \left(\mu^0 + \mu^y y + \mu^z z\right)\right) \\ &\quad - \sum_i \alpha_i^b \log p_i - \sum_j \alpha_j^b \log P_j \end{aligned}$$

which does not depend on the c_i^s. Secondly, the form of aggregate demands is such that private and public goods have exactly the same structure. We therefore simplify our notations by defining

$$\xi_i = q_i \quad \text{for } i \leq n,\ \xi_i = Q_i \quad \text{for } n < i \leq N$$

and similarly

$$\gamma_i = c_i \quad \text{for } i \leq n,\ \gamma_i = C_i \quad \text{for } n < i \leq N$$

$$\pi_i = p_i \quad \text{for } i \leq n,\ \ \pi_i = P_i \quad \text{for } n < i \leq N$$

so that the group demand has the simple form:

$$\pi_i \xi_i = \pi_i \gamma_i + \left[\alpha_i^a \left(\mu^0 + \mu^y y + \mu^z z\right) + \alpha_i^b \left(1 - \left(\mu^0 + \mu^y y + \mu^z z\right)\right)\right] Y \tag{9}$$

leading to collective indirect utilities of the form:

$$\begin{aligned} W^a(p, P, y, z) &= \log Y + \log\left(\mu^0 + \mu^y y + \mu^z z\right) - \sum_i \alpha_i^a \log \pi_i \\ W^b(p, P, y, z) &= \log Y + \log\left(1 - \left(\mu^0 + \mu^y y + \mu^z z\right)\right) - \sum_i \alpha_i^b \log \pi_i \end{aligned}$$

It is clear, on this form, that the distinction between private and public goods can be ignored. This illustrates an important remark: while the *ex ante* knowledge of the public versus private nature of each good is necessary for the identifiability result to hold in general, for many parametric forms it is actually not needed.

Identifiability: the general case The question, now, is whether the empirical estimation of the form (9) allows us to recover the relevant parameters - namely, the α_i^s, the γ^i, and the μ^α. We start by rewriting (9) as:

$$\pi_i \xi_i = \pi_i \gamma_i + \begin{pmatrix} \alpha_i^b + \left(\alpha_i^a - \alpha_i^b\right) \mu^0 \\ + \left(\alpha_i^a - \alpha_i^b\right) \left(\mu^y y + \mu^z z\right) \end{pmatrix} \left(y - \sum_m \pi_m \gamma^m\right) \tag{10}$$

The right hand side of (10) can in principle be econometrically identified; we can thus recover the coefficients of the variables, namely $y,\ y^2,\ yz$, the π_m and the products $y\pi_m$ and $z\pi_m$. For any i and any $m \neq i$, the ratio of the coefficient of y by that of π_m gives γ^m; the γ^m are therefore vastly overidentified. However, the remaining coefficients are identifiable only up to an arbitrary choice of two of them. Indeed, an empirical estimation of the right hand side of (10) can only recover for each j the respective coefficients of y, y^2 and yz, that is the three expressions:

$$\begin{aligned} K_y^j &= \alpha_j^b + \left(\alpha_j^a - \alpha_j^b\right)\mu^0 \\ K_{yy}^j &= \left(\alpha_j^a - \alpha_j^b\right)\mu^y \\ K_{yz}^j &= \left(\alpha_j^a - \alpha_j^b\right)\mu^z \end{aligned} \tag{11}$$

Now, pick up two arbitrary values for μ^0 and μ^y, with $\mu^y \neq 0$. The last two expressions give $\left(\alpha_j^a - \alpha_j^b\right)$ and μ^z; the first gives α_j^b therefore α_j^a.

As expected, a continuum of different models generate the same aggregate demand. Moreover, these differences are welfare relevant, in the sense that the individual welfare gains of a given reform (say, a change in prices and incomes) will be evaluated differently by different models. In practice, the collective indirect utilities recovered above are not invariant across the various structural models compatible with a given aggregate demand.

A unitary version of the model obtains when the Pareto weights are constant: $\mu^y = \mu^z = 0$. Then $K_{yz}^j = 0$ for all j (since distribution factors cannot matter), and $K_{yy}^j = 0$ for all j (demand must be linear in y, since a quadratic term would violate Slutsky). We are left with $K_y^j = \alpha_j^b + \left(\alpha_j^a - \alpha_j^b\right)\mu^0$, and it is obviously impossible to identify independently α_j^a, α_j^b and μ^0; as expected, the unitary framework is not identifiable.

Identification under exclusion We now show that in the non-unitary version of the collective framework, an exclusion assumption per member is sufficient to exactly recover all the (welfare-relevant) coefficients. Assume that member a does not consume commodity 1 and member b does not consume commodity 2; that is, $\alpha_1^a = \alpha_2^b = 0$. Then equations (11) give:

$$\alpha_1^b \left(1 - \mu^0\right) = K_y^1,\ -\alpha_1^b \mu^y = K_{yy}^1,\ -\alpha_1^b \mu^z = K_{yz}^1$$

and:

$$\alpha_2^a \mu^0 = K_y^2,\ \alpha_2^a \mu^y = K_{yy}^2,\ \alpha_2^a \mu^z = K_{yz}^2$$

Combining the first two equations of each block and assuming $\mu^y \neq 0$, we get:

$$\frac{1-\mu^0}{\mu^y} = -\frac{K_y^1}{K_{yy}^1} \quad \text{and} \quad \frac{\mu^0}{\mu^y} = \frac{K_y^2}{K_{yy}^2}$$

therefore, assuming $K_y^2 K_{yy}^1 - K_y^1 K_{yy}^2 \neq 0$

$$\frac{1-\mu^0}{\mu^0} = -\frac{K_y^1 K_{yy}^2}{K_y^2 K_{yy}^1} \text{ and } \mu^0 = \frac{K_y^2 K_{yy}^1}{K_y^2 K_{yy}^1 - K_y^1 K_{yy}^2}$$

It follows that

$$\mu^y = \frac{K_{yy}^2}{K_y^2} \mu^0 = \frac{K_{yy}^2 K_{yy}^1}{K_y^2 K_{yy}^1 - K_y^1 K_{yy}^2}$$

and all other coefficients can be computed as above. It follows that the collective indirect utility of each member can be exactly recovered, which allows for unambiguous welfare statements. As mentioned above, identifiability is only generic in the sense that it requires $K_y^2 K_{yy}^1 - K_y^1 K_{yy}^2 \neq 0$. Clearly, the set of parameters values violating this condition is of zero measure. Also, identifiability requires $\mu^y \neq 0$; in particular, *it does not hold true* in the unitary version, in which $\mu^y = \mu^z = 0$. Indeed, the same exclusion restrictions property as above only allow to recover $\alpha_1^b \left(1 - \mu^0\right) = K_y^1$ and $\alpha_2^a \mu^0 = K_y^2$; this is not sufficient to identify μ^0, let alone the α_j^i for $j \geq 3$. This confirms that the unitary version of the model is not identified even under the exclusivity assumptions that guarantee generic identifiability in the general version.

Finally, one can readily check the previous claim that the MMWIs are not identified. Indeed, the MMWI m^s of s is defined by:

$$v^s(\pi, m^s) = \log\left(m^s - \sum_k \pi_k \gamma_k^s\right) - \sum_i \alpha_i^s \log \pi_i = W^s(\pi, y, z)$$

where

$$v^s(\pi, P, y) = \log\left(y - \sum_k \pi_k \gamma_k^s\right) - \sum_{i=1}^n \alpha_i^s \log \pi_i$$

and

$$W^s(\pi, z) = \log\left(y - \sum_{i,k} \pi_i \gamma_i^k\right) + \log\left(\mu^0 + \mu^y y + \mu^z z\right) - \sum_i \alpha_i^s \log \pi_i$$

This gives

$$m^s(\pi, y, z) = \left(\mu^0 + \mu^y y + \mu^z z\right)\left(y - \sum_i \pi_i \left(\sum_k \gamma_i^k\right)\right) + \sum_i \pi_i \gamma_i^s$$

For any private commodity i, the sums $\sum_k \gamma_i^k$ are identified, but the individual γ_i^s are not; therefore m^s is identified up to an additive function of the prices of private, non exclusive goods.

5.2 Comparing different family sizes

A second approach enlarges the set of usable information by allowing comparisons between families of different composition. A first idea is to assume some relationship between individual preferences when married and single. In that sense, the 'pure' approach just described relies on an extreme version, since it does not postulate *any* link between utilities when married and single; hence, knowledge of an individual's preferences when single brings no information about her tastes within the household. At the other extreme, some models assume that preferences are unaffected by marital status, at least ordinally. This means that if $u_S^a(Q, q^a)$ denotes a's utility when single, then her utility when married takes the form:

$$u^a(Q, q^a) = F(u_S^a(Q, q^a))$$

where F is an increasing transform. Thus marriage can directly affect a person's utility *level*, but not the person's marginal rates of substitution between various commodities. Note that if we assume preferences are unaffected by marital status, then the MMWI defined above has a natural interpretation; namely, it is the level of income that would be needed by the individual, *if single*, to reach the same utility level as what she currently gets within marriage. It must however be stressed that the assumption of constant preferences across marital status is not needed for the *definition* of the index, but only for this particular interpretation.

Various, intermediate approaches can be found in the literature. One, mostly used in a labor supply context, only assumes that some preference parameters are common to singles and households, and can therefore be estimated separately on a sample of singles. In general, this is sufficient to identify (or calibrate) the remaining parameters (relevant for marriage-specific preferences and the Pareto weights) on observed labor supplies of men and women in a sample of couples. This approach has been adopted in a series of papers recently published in the *Review of Economics of the Household* (Bargain et al., 2006; Beninger et al., 2006; Myck et al., 2006; Vermeulen et al., 2006). For instance, consider a model of labor supply in a couple in which the utility of agent a takes the form:

$$u^a\left(q^a, L^a, L^b\right) = \alpha^a \ln\left(q^a - \bar{q}^a\right) + \beta^a \ln\left(L^a - \bar{L}^a\right) + \gamma^a \ln\left(L^a - \bar{L}^a\right) \ln\left(L^b - \bar{L}^b\right)$$

where L denotes leisure; note that this form is more general than the ones considered above, since it allows for (positive) externalities of leisure within the couple.[13] The α and β parameters are assumed to be independent of marital status, and are therefore identified from a sample of singles; the γs and the Pareto weights are then calibrated from data on households.

An intermediate approach, that relies on the notion of domestic production, has recently been proposed by Browning, Chiappori and Lewbel (2013). It posits that agents, when they get married, keep the same preferences but can access a different (and generally more productive) technology. That is, while the basic rates of substitution between *consumed* commodities remains unaffected by marriage (or cohabitation), the relationship between purchases and consumptions is not; therefore, the structure of demand, including for exclusive commodities (consumed only by one member) is different from what it would be for singles. More generally, one can, following Dunbar, Lewbel and Pendakur (2012), only assume that preferences are unaffected by family composition; e.g., that parents' preferences regarding their own consumption does not depend on the number of children. These approaches are described in the next section.

5.3 Identifying from market equilibrium

Lastly, a series of recent contributions are aimed at taking to data the equilibrium approaches described above. The basic, theoretical intuition is quite straightforward: the equilibrium conditions on the marriage market (with or without search frictions, but with intrahousehold transfers) either constrain or

[13] Equivalently, this approach considers both leisures as public goods within the household.

exactly pin down intrahousehold allocations. Several papers propose an empirical implementation of this idea. A first set of works only consider matching patterns; the marriage market equilibrium is then exclusively characterized by a matrix of intermarriages between various categories, which can be defined by age, education, income or any combination of these. On the matching front, following the initial contribution by Choo and Siow (2006), Chiappori, Salanié and Weiss (2012) have shown how a structural, parametric model can be (over)identified from such patterns, under the assumption that, while the surplus generated by marriage may (and does) vary over time, its supermodularity (which drives the extent of assortative matching in the population) is constant.[14] According to their estimate, while the gain from marriage have globally decreased over the last decades, the decline has been much smaller for educated couples. Moreover, the share of household resources received has increased for college educated wives, resulting in a strong increase in their 'marital college premium' (defined as the additional gain provided by university education on the marriage market). This is compatible with the theoretical analysis of Chiappori, Iyigun and Weiss (2009), who argued that the asymmetry between male and female marital college premium could explain (at least in part) the higher demand for university training by women. Alternatively, Jacquemet and Robin (2013) and Goussé (2013) analyze marital patterns from a search perspective.

A clear limitation of these approaches is that the sole observation of marital patterns conveys only limited information on the form of the marital surplus (therefore on distribution). For instance, knowing that matching is assortative only tells us that the surplus is supermodular. The previous approaches, therefore, must rely on strong and largely untestable assumptions on the precise form

[14]For a general presentation of the econometrics of matching models, see the survey by Chiappori and Salanié (2014).

of the heterogeneity distribution across couples. Adding information on transfers would greatly enhance the identification power of these models. But such information is precisely what collective models can provide, based on observed behavior. The intuition, here, is that the observation of, say, labor supply patterns of married couples (which reflects intrahousehold transfers), together with that of marital patterns, should allow to fully identify a general matching model in a very robust way. This line of research is pursued, in a series of paper, by Chiappori, Costa Dias and Meghir (2014a, b).

6 Empirical findings

In this section we review some empirical work based on the collective model and emphasizing the identification of the sharing rule.

The first generation models used information on private and assignable goods such as consumption of clothing or individual leisure to identify the sharing rule up to a constant. These models adopt mainly two approaches for identification. The first approach refers to what we called 'pure' identification; i.e., it recovers the derivatives of the sharing rule with no further information than observed consumption bundles of the household. As discussed above, while some identifying conditions can be relaxed by using distribution factors, these models cannot identify separately the level of sharing (how much goes to each household member) from preferences. There exist a continuum of allocations of resources, each associated to a utility function for each household member, that fit the data equally well; across these allocations, income inequality within the household is different, although the allocation of welfare to each member remains the same.

To identify the way overall resources are allocated and thus measure inequality, one needs more information, either in terms of identifying assumptions on the behavior of the sharing rule (such as non-negativity conditions discussed earlier) or assumptions on preferences. One possibility is to compare the behavior of married and single individuals by making assumptions on the way preferences change with marriage. Other approaches involve specific restrictions on preferences. We show how some of these approaches have been used in the literature. Finally we also consider the information content of revealed preference restrictions. These extend the revealed preference arguments for individual choice to the case of collective households. Clearly this is a much more complicated setup than standard revealed preference restrictions for individuals or for unitary households because the aggregate household does not necessarily behave like a rational single agent. We discuss what can be learned from revealed preference in this context.

However, the issue of identification of the sharing rule is deeper than what is suggested by the use of the restrictions above and has to do with the way people make agreements at the point of marriage and the level of commitment associated with these agreements. In other words, fundamentally the sharing rule is identified from behavior without having to impose possibly ad hoc restrictions. Identification requires extending the model to include marital decisions in an equilibrium context. Indeed a marriage market equilibrium will define the sharing rule and conditions in the marriage market can allow us to identify it. This effectively introduces dynamics, which then allows one to delve deeper into the extent of commitment and what this means about within household inequality. Characterizing the theoretical and empirical power of using marriage market data to understand better intrahousehold allocations is is a relatively new and active area of research, particularly when limited commitment is allowed for.

Before we discuss the empirical literature we need to introduce a distinction between the concept of identifiably of preferences and the sharing rule on the one hand and econometric identification on the other. The identifiability results discussed above relate to our ability to recover individual preferences and the sharing rule given *we know* the household level demand functions exactly. Empirical analysis is concerned with estimating these household demands from empirical data so as to be able to then recover the sharing rule. This issue brings forth all the standard econometric concerns, such as the role of unobserved heterogeneity, the endogeneity of wages, prices and income, corner solutions (particularly in labor supply) etc. One of the hardest issues concerns the way that unobserved heterogeneity enters household demands, particularly if such unobservables are correlated with observables. The specific issue arises from the fact that, in general, unobserved heterogeneity in preferences will imply unobservables in the sharing rule. In most specifications this will mean that unobservables are non separable from observables, with implications for econometric identification. For example, Blundell, Chiappori, Magnac and Meghir (2007) used linearity to bypass the difficulties implied by unobserved heterogeneity in preferences. Here we are not offering any general solution to the problem, but we need to point out that, before we even consider identification of the sharing rule, an empirical approach would have to solve the standard econometric identification issues, which in this context may be severe.[15]

6.1 'Pure' identification of the sharing rule

In this first generation of collective models we can point to three main empirical studies. The first is by Browning, Bourguignon, Chiappori and Lechene (1994, BBCL); the second is by Chiappori, Fortin and Lacroix (2002, CFL) and the third by Blundell, Chiappori, Magnac and Meghir (2007, BCMM). All three share a similar approach to identification: they assume efficiency and an assignable good. However, the details of the empirical approach differ.

[15] For recent attempts in this direction (including a discussion of the specific difficulties it raises), see for instance Lewbel and Pendakur (2013) and Chiappori and Kim (2013).

In BBCL the authors use a sample of couples drawn from the Canadian FAMEX and estimate a model for the demand of mens' and women's clothing, and identify the sharing rule, up to a constant. Identification relies on two assumptions: first clothing is an assignable good, which effectively means that we can observe male and female clothing and that only the person using the clothing derives any utility from it. In other words clothing does not include a public good element. Second they assume that the distribution of partner's income does not affect preferences, but may enter the sharing rule, reflecting bargaining positions. Given these assumptions they identify a sharing rule as a function of the age difference of the partners, total household expenditure (thus allowing wealth effects in the way resources are shared) and most importantly the share of income attributable to the female partner. It turns out that the effect of the way resources are distributed between couples is not very sensitive to the proportion of income for which they are accountable. For example going from a share of income of 25% to 75% raises the share of household expenditure by a significant but small 2.3%. The age difference and the level of expenditure also matter with relatively older individuals gaining more and wealthier households allocating more to the wife.

The BBCL paper shows the potential of the approach and the richness of the empirical results that can be obtained by judicious use of information reflecting bargaining power of households. However, the main determinant of female bargaining power in their model is the relative magnitude of female income. A higher share of income may reflect her relative skills or alternatively it may reflect her decision to forgo leisure and work more; in other words, this distribution factor is indeed endogenous. In principle, this fact does not harm identification provided that labor supply is separable from consumption: controlling for total expenditures, individual consumption should then be independent of labor supply (therefore of labor income). However, separability is a strong assumption, that has been empirically criticized. The next two papers address exactly this issue by endogenizing labor supply.

Specifically CFL set up a model of collective labor supply for couples with no public goods. The key feature of their model is the use of distribution factors. They use the sex-ratio (males/females) in the relevant state and a set of indicators describing the nature of divorce laws. Within that first generation paper the sex ratio and the divorce laws that favor women at the dissolution of marriage are viewed as factors that will improve the share of women.[16]

The empirical relevance of the discussion above for within household inequality and allocation of resources is illustrated by CFL. They use data from the PSID to estimate a collective labor supply model, where the sharing rule is identified (up to a constant) based on distribution factors. These include the sex ratio in the state as measured by the 1990 census as well as as dummy variables indicating the nature of divorce laws. Measuring the sex-ratio is of course very tricky, both because we need to define the relevant labor market and because timing may matter. In a full commitment model for example the sex ratio at the time of marriage is what is going to matter. However the sex-ratio is unlikely to change vastly over time and it is probably a good idea to define marriage markets quite broadly rather than too narrowly. The authors also report using the county level sex-ratio with the state level as an instrument, which had little impact on their results. In their model labor supply is evaluated over one whole year and they consider a sample where both are working. So the relevant group are individuals with sufficient attachment to the labor market to want to work at least some part of the year. In their model the sharing rule is allowed to be

[16]The intuition underlying the CFL paper - that a relative scarcity of women and/or more favorable divorce laws should improve the wife's Pareto weight - can be supported by an explicit matching model, with some nuances (e.g., changes in divorce laws affect differently women already married and women getting married after the change). On these issues, see Chiappori, Iyigun, Lafortune and Weiss (2013).

a function of the wages of both partners, nonlabor income and the distribution factors. Allowing both wages to enter is important: it has been empirically observed that both wages matter when estimating family labor supply (see for example Blundell and Walker, 1986) a fact that in a unitary context has been interpreted as nonseparability in household preferences between male and female leisure. Here this non-separability is interpreted as being driven by the impact of the sharing rule on individual labor supply in a collective setting. The fact that the restrictions from the collective model are not rejected strengthens this interpretation.

The results suggest that marriage and labor market conditions can lead to large differences in the allocation of household resources within a couple. For example a \$1 increase in the female hourly wage rate leads to a transfer at the means of \$1,600 to the husband, implying that most of the extra income goes to him. However a \$1 increase in his wage rate leads only to \$600 transfer to her, implying he keeps the lion share and does not behave as altruistically (to use the authors' words) as she does. The wage effects are of particular interest because changes in wages and in male/female wage differentials may be a key driver of within household allocation of resources. Unfortunately these results are not precisely estimated; we revisit this issue below, in our discussion of BCMM and of Lise and Seitz (2011). Anyhow, a result that stands out in CFL is the impact of the sex ratio. Based on this result an increase of one percentage point in the sex ratio leads to \$2,160 transfer to the wife. Noting that the range of the sex ratio in their data is 0.46-0.57 the implication is that from the least favorable to the most favorable labor market the transfer can differ by as much as \$23,000. Of course this does not all translate into an increase in consumption because the income effect on labor supply will imply a change in the amount of hours worked, with women living in marriage markets more favorable for them, working less. To obtain a summary of divorce laws the authors constructed an index ranging from 1-4 and indicating the extent to which the divorce laws are favorable to women. Here the effects are particularly strong as well. A one point increase in the index leads to a transfer of $4,310 to the wife, which again is shared between consumption and leisure.

These results are important because they show the extent to which within household allocation of resources can be sensitive to external conditions affecting the bargaining power of the members of the couple. Noting for example that average household income in this data is $48,000 the change that can be induced just because of (admittedly extreme) changes in the sex ration can amount to almost half of household income.

However, there are a number of empirical issues that were not addressed by the papers already discussed. First, we need to be concerned that the allocation of women across states with different sex ratios is not random with respect to their unobserved preferences for labor supply. This can bias the results if women who live in areas abundant with men tend to have lower labor market attachment. Second, we need to address the issue of precision in the estimation of wage effects, an issue that persists in the BCMM paper we will discuss below. CLF instrument wages but the instruments are necessarily quite week: they rely on a polynomial in age and education as an instrument while (correctly) controlling for the level of education and for age in the labor supply function. This leaves higher order nonlinearity in the profile of wages with respect to age and education to act as an instrument which is both difficult to justify theoretically and at the same time is not very informative. To solve these empirical issues we will require exogenous events that change wages and the marriage market, something that a newer generation of collective models is now addressing, such as the paper by Attanasio and Lechene (2014) who use the experiment as an exogenous shifter in female bargaining power.

Beyond these difficulties there is one further important issue that the papers we discussed fail to address, namely non-participation of women. Given that many women do not work allowing for this possibility and understanding how resources are allocated despite the fact she is not producing in the formal market is a key concern. The BCMM paper addresses the question of identification and estimation of a collective labor supply model with both male and female non-work. In addition it considers the case where the male can only choose to work or not, rather than having a choice of hours of work. This restriction is imposed to accommodate the fact that in the UK (where the data is drawn from) the male hours of work distribution seems discontinuous between zero and about 35 hours per week, with the entire mass of workers concentrated in the full time range. This restriction is not entirely satisfactory, but it may do better justice to the data than assuming hours of work are freely chosen. Thus the resulting paper is where females make choices both on the intensive and the extensive margin, while males choose only on the extensive margin. The authors prove identification of the sharing rule; however this is only identified (non-parametrically) if at least one of the two household members work. Parametric restrictions provide the rest. In the empirical implementation BCMM deal with the endogeneity of the wage rate by exploiting the changes in wage inequality across cohorts and education groups. Econometric identification relies on the assumption that while the structure of wages changed across education groups and cohorts - a testable assumption, preferences remained unchanged. This implies that changes in work behavior across cohorts and education groups can be attributed to changes in the incentive structure., which is the identification strategy employed by Blundell, Duncan and Meghir (1998).

The empirical analysis is conducted on a sample of married couples, observed between 1978-2001 in the UK Family Expenditure Survey. The assumptions imposed for identification (over and above efficiency) required only private goods and one assignable good. The assignable good is leisure. Since expenditures on children are not separately observable in the data and since these are effectively public the authors exclude all couples with children and then assume that the observed aggregate household consumption reflects the sum of private consumption of each of the two members of the household.

In this model the authors estimate two different sharing rules depending on whether the husband works or not. They differ by a monotonic transformation, which in their empirical specification acts as an attenuation factor, implying that the husband only gets a fraction of transfers when he is not working. This fraction is 0.71, implying that the derivatives of the sharing rule (as well as the level) are attenuated by that amount when he is not working. Their empirical approach does not use any distribution factors that can be excluded from preferences: the sharing rule depends on male wages, female wages and unearned income as well as education and age. It turns out that empirically the effect of the female wage on the sharing rule is not well identified. However the effect of the male wage is precisely estimated. It implies that 88% of an increase in male market earnings translates into a transfer to the husband if he is working. Since there is no intensive margin for the male decision this translates to a direct impact on his consumption, if he continues to work. If he does not work the same change in potential earnings translates to a transfer equal to 62% of the potential increase (0.71×0.88). These results imply that when the earnings of a working husband increase the resulting increase in the consumption of the wife is only small; if potential earnings increase (and he is not working) her consumption declines substantially and he enjoys more of the household resources. Finally, the wife keeps 73% of increases in unearned income. Nevertheless, unearned income is a relatively low fraction of household income.

These results again illustrate that external factors (here the relative wages) can influence the allocation of resources substantially. Unfortunately BCMM does not provide precise estimates of the effects of female wages and this hinders an understanding of how the change in the wage structure affected within household allocations. The source of lack of precision is the relatively small sample size where the man does not work. Moreover, allowing both wages and non-labor income to be endogenous, while important for obtaining consistent estimates that make sense, does affect precision substantially. The paper does demonstrate that one does not need (in principle) distribution factors for identification. However, looking at the empirical problem from the perspective of CFL, other environmental factors may be very important in determining allocations and if they are omitted they could bias the results. On the other hand if included they can be allowed to affect preferences as well. Identification does not require they affect the sharing rule alone.

This first generation of models showed the potential of the collective model for identifying allocations of resources within the black box of the household. However there are key issues that had not been dealt with. First, taxes and welfare were ignored. At one level this is an empirical specification issue because ignoring taxes can bias the estimates of the preference parameters. But at a more fundamental level by not taking into account the tax and welfare system we omit one of the most important factors affecting (and sometimes designed to affect) within household allocations. Estimating models that allow for taxes and welfare can then explain how changes in the policy and the market environment can affect the allocation of resources.

The next fundamental issue is that the models described above can only identify the the derivatives of the sharing rule, i.e. how sharing changes when distribution factors, prices and unearned income change. This precludes any discussion of the the levels of inequality of resources and hence does not allow us to put into perspective the implication of changes that occur over time. It also does not allow us to obtain a complete picture of the distribution of welfare in the economy.

Adding taxes and welfare does not pose any important conceptual problems. In practice it involves allowing for more complex budget sets and solving the model to take into account nonlinear budget sets. An interesting issue is that the welfare and tax system may create a further interdependence in the decisions of husband and wife, over and above that induced by the sharing rule. These issues are considered for instance in Donni (2003), who uses a 'pure' identification strategy of the type just described, and by Beninger et al. (2006), Myck et al. (2006) and Vermeulen et al. (2006) who use information from singles and couples.

Extending the model to allow identification of the level of the sharing rule does however pose conceptual problems. Fundamentally, the sharing rule is identified by the equilibrium in the marriage market. However, barring the use of a complete marriage market equilibrium model one can obtain information on the level of inequality with alternative auxiliary assumptions. One possibility is to use information on singles. This involves restricting the way preferences change with marriage. This is an approach used by Lise and Seitz in an early version of their paper. Another possibility is to assume something about the sharing rule at one point of the wage space. For example that all resources are shared equally when wages are equal, which is the assumption made in the published version of Lise and Seitz (2011). Finally, one can make assumptions about the functional forms of demand, as in Dunbar, Lewbel and Pendakur

(2013). We now look into these empirical studies.

6.2 Intrahousehold inequality over time and the sharing rule - Lise and Seitz (2011)

Lise and Seitz (2011) use the collective model to first estimate overall consumption inequality (at the individual level) and to then decompose this to between household and within household. The important economic fact is that the distribution of wages in the UK changed dramatically over the period they consider [1968-2001] both within and between education groups (see Gosling, Machin and Meghir, 2000). Moreover the structure of the marriage market has also changed with increased degrees of marital sorting over time. They thus set up a model of male and female labor supply with many (but discrete) choices of hours worked for both members of the household. Hours can take values from 0-65 in five hour intervals. In many ways their empirical framework is similar to that of BCMM: they use couples with no children drawn from the UK Family Expenditure Survey over many years. However they depart in a number of important ways: first, they allow for taxes and account for the impact of joint taxation over the period that this was in effect in the UK (up to 1989); they allow for a richer choice set for the male and they impose further structure so as to identify the level of the sharing rule as well as its derivatives; they account for public goods when they define consumption, although they are taken as separable from private consumption and leisure.

While the logic underlying the identification of the derivatives of the sharing rule is similar to that of BCMM, identification of the location (level) of the sharing rule empirically is based on the identifying assumption that when individuals have the same potential earnings they share resources equally. In earlier versions of the paper it was instead assumed that preferences of married and single individuals are identical; both these assumptions can identify the model. The point at which one pins down the sharing rule is welfare irrelevant, because the preference specification adapts to leave welfare unchanged when the location of sharing is fixed. In principle just normalizing the location parameter will not cause any bias, but will of course lead to a specific level of inequality. On the other hand using information from singles has the advantage that it uses a restriction grounded in some explicit assumption on preferences (marriage does not affect marginal utilities) but if wrong it will bias all results.

Over the period considered in the paper (1968-2001) earnings inequality increased rapidly; there has been a steady increase in both the potential earnings and actual earnings share of women relative to men and a decline in male employment while female employment increased at the start of the period later remaining constant. Consumption inequality increased rapidly in the period between 1980 and 1990, but was basically stable the rest of the time. When Lise and Seitz interpret these results under the prism of their collective model they uncover some interesting facts: while between household inequality of consumption increases, within household inequality of consumption declines to such an extent that the overall inequality of consumption remains more or less the same over time. When they consider a different measure of resources, namely full consumption, which includes the value of leisure enjoyed by each member, they find similar but less stark results: first between household inequality still increases, but much less dramatically because the decline in consumption for those households who have workless members is compensated by the value of leisure; second again they find that within household inequality declines as before, but much less. Obviously none of these consumption measures is ideal and a money-metric measure of welfare may be better. However, these results illustrate exactly the potential importance of finding credible ways to understand inequality (and poverty) within households. This is more so given that who marries whom is endogenous and in part drives the way that within household inequality is determined and has implications for between household inequality is determined.

6.3 Intrahousehold inequality and children

While intrahousehold inequality may be of general interest because it tells us about allocation of resources within a household and can reveal hidden poverty and inequality, the whole issue acquires special importance when it comes to allocations of consumption to children. Thus is because child consumption and more generally investments in children have long term implications for the intergenerational transmission of poverty. Yet little or no empirical work had been done to understand how resources are allocated to children and the extent to which reallocations of income from the male spouse to the female can affect the shares directed to children. A theoretical framework for the analysis of this question has been developed by Blundell, Chiappori and Meghir (2005). In a recent important paper Dunbar, Lewbel and Pendakur (2013) address this issue empirically, using data from Malawi. In their model each child is represented as having their own utility function. This creates a very special difficulty regarding the assumption, used for identification in studies such as Browning, Chiappori and Lewbel (2007), that preferences of singles and married individuals are the same. Here, such a strategy is no longer available because children are never seen living as singles. Moreover, in data from Malawi that the authors use there is not enough price variation - another requirement of the Browning, Chiappori and Lewbel approach. Thus identification is obtained by making assumptions on the structure and shape of the Engel curves.

The identification strategy first requires either one assignable, private good

good or one exclusive good per person. Remember that an exclusive good is exclusively consumed by one household member type (for example child clothing is consumed only by children), while an assignable good is such that each member's consumption of this good is observable.[17] Of course, there can be many other purely private goods (such as food) for which we do not observe the amounts of individual consumption - this fact does not hamper identification. Beyond the presence of one assignable good all other goods can be private or public or partially private.

The assignability assumption is not sufficient to identify the share of resources of each household member; additional assumptions are therefore needed. Dunbar et al (2012) assume, first, that resource shares are invariant to total expenditure. In addition, they make two alternative assumptions on preferences: either the demand for goods is similar across household types (i.e. households with one, two or more children) or they are similar across types of goods within a household type. An extreme form of the assumption is that preferences do not vary across types of household; since shadow prices vary across households because of the partially public nature of goods this extreme assumption is essentially equivalent to assuming that the assignable good used for identification is irresponsive to prices. Another extreme form of this assumption is that preferences over the assignable good are identical across different household member types (male, female and children). However, Dunbar et al. (2012) show that identification only requires that some aspect of the demand functions be the same either across household member types or across household types. Thus in one case they assume that all household members share the same shape of Engel curves for the assignable good. In another case they assume that prefer-

[17]It should be stressed that a good is private when its consumption does not alter the *preferences* of other household members over goods consumption: as the authors put it, smoking by one household member may annoy the others but it can still be taken as private if it does not in itself alter their consumption of goods.

ences are the same across types of household (number of children) conditional on a deflator of income. This deflator reflects the different shadow prices that different sized households face and is the way that preferences for the assignable good are allowed to vary across types. The key point is that the authors need to define similarity so that identification is delivered without sacrificing theoretical consistency (integrability) of the demand functions.

Dunbar et al. (2012) estimate their model on data from Malawi, probably in one of the first such studies with development data. In a sense their framework is very well adapted to this context because wages and/or prices, which are at the heart of some other identification strategies are often not observed. Their approach relies on measuring expenditures and having an assignable good for which they use clothing and footwear. The results they obtain are both astounding and an excellent illustration of the importance of looking within the household. They find that the male obtains about 45-50% of household resources. His share seems to be insensitive to the number of children present. The mother's share declines with the second child, but then remains more or less constant, with the consumption share of children declining.

Even more pertinent are the implied poverty rates. Male poverty rates are at their highest in one child households and seem to decline in households with more children. However, the important result is that on poverty rates for women and children: compared to the male poverty rate of around 69%, there are 79% poor women and 95% poor children in one child households. In larger households the male poverty rate is about 55% while the female poverty rate is 89% and nearly all children are poor. Hence their approach not only offers a more complete picture of poverty but reveals the extent of child poverty, which is crucial to development. Without such an approach, child poverty would not be apparent to the extent that it is in reality. While the authors did not focus on gender differences between children, which may be another important dimension, this line of research can easily be extended in that direction; it offers an obvious mechanism for trying to understand how resources are allocated by gender.

A potential limitation of this approach is the fixed nature of the sharing rule. While the authors spend a lot of time explaining the upsides of not relying on distribution factors (essentially, they avoid having to take a position on whether they affect preferences or not), the absence of an underlying model of what the resource share should depend on and how it can be affected by exogenous driving forces may in some cases be problematic. In models where the sharing rule is allowed to depend on wages or institutional features we have some understanding of how policy can be used to target individuals. In the Dunbar et al model this aspect is missing. However, this is not an integral part of the approach and richer models can be identified.

6.4 Revealed Preference Restrictions and the identification of the Sharing Rule

The approach to the identification of the sharing rule has exploited the structure of the demand functions and the way that income affects observed outcomes when the collective model is true. This leads to a set of differential equations that when solved provide the derivatives of the sharing rule. However, the approach does not identify the level of the sharing rule.

A different approach is that of revealed preference. In the context of the single agent utility maximization model the axioms of revealed preference allow one to test nonparametrically whether a particular set of choices can be rationalized by utility maximization and if they can, to bound the underlying demand functions. Such an approach has been developed and implemented for the unitary model by Blundell, Browning and Crawford (2003, 2008) and is based on the original work of Afriat (1973) and Varian (1982) . In the collective framework the aggregate household demands will in general violate the revealed preference restrictions corresponding to the unitary model simply because as the budget constraint changes (wages, prices, incomes etc.) individuals make different choices and in addition the Pareto weights change. This insight was developed by Browning and Chiappori (1998) who showed that the aggregate household demands have to possess a Slutsky matrix that can be decomposed into a symmetric matrix plus a matrix with rank equal to the number of decision-makers (whose demands are aggregated) minus 1. The fact that the pattern of choices is restricted implies that there should also be revealed preference type restrictions - as noticed by Chiappori (1988), who provides an early example in a labor supply context. Indeed these restrictions have been fully developed by Cherchye, De Rock and Vermuelen (2007). In a further development Cherchye, Lewbel, De Rock and Vermuelen show how the revealed preference restrictions can be used to bound the sharing rule without imposing any restrictions other than Pareto efficiency of intra household allocations. The main result is based on the following principle: suppose that a set of observed demands are collectively rationalizable in the sense that the observed choices are consistent with the existence of admissible individual demand functions. Then it has to be that any alternative choices that could lead to a Pareto improvement within the household should be infeasible at current market prices and for any allocation of income within the household such that each person receives a non-negative share. More specifically, consider the set of demands of individual 1 that are revealed preferred to the current choice, based on all possible admissible demand functions for that person. They must cost more than person ones's share of total household income; similarly for person 2. The least costly bundle that would lead to a Pareto improvement provides the upper bound for a person's share. The adding up the shares to total income and the assumption that the shares cannot be negative determines the lower bound. The difficulty in implementing this principle is the fact that we need to search over all possible admissible individual demand functions.

This principle turns out to generate non-trivial upper and lower bounds for the sharing rule. Importantly, no restriction is needed for such bounds other than Pareto optimality: all or some goods may be either private, in part public and in part private or completely private. Moreover, we do not need to specify which goods (if any) are purely private, but if such information were to be available it can be used to tighten the bounds.

Cherchye et al. apply their approach to the PSID form 1999, when expenditures on individual consumption goods became available, until 2009. The sample consists of childless couples where both are working. Utility depends on leisure food and other goods which include health and transportation. Leisure is assumed assignable, but no assumption is made on the other goods. This is important because in this case, at least in general, neither the level nor the derivatives of the sharing rule are point identified.

To implement their approach they start by estimating three different versions of an aggregate household demand system: a non parametric system, the QUAIDS demand system (Banks, Blundell and Lewbel,) and a QUAIDS demand system where the substitution matrix is restricted to be symmetric plus rank one, which imposes that the demands are consistent with the collective model. Given this demand system they apply their algorithm to bound the sharing rule for different values of the full household income, wages and prices. Their empirical results are remarkable. First, the bounds are very narrow with the nonparametric demand system implying 12% median difference between upper and lower bounds and the fully restrictive demand system only 3%. Going from the non-parametric demand system to the unrestricted QUAIDS system the tightening is due to imposing the parametric restrictions that may or may not be valid - the authors provide no evidence on that. However, assuming the parametric restrictions are valid the further step of going from QUAIDS to restricted QUAIDS is just imposing restrictions that are implied by the problem and hence only serve to make the bounds sharp(er). Thus when Pareto efficiency is imposed the median difference between the upper and lower bound tightens from about 9% to 3%, a substantial improvement. It would have been useful to use a shape constrained nonparametric demand system (see Blundell, Horowitz and Parey) avoiding the parametric restrictions but using the Pareto constraints as implied by the model.

Using their bounds they establish that the female share is a normal good, i.e. as full household income grows so does the female share; interestingly, this finding confirms results previously derived in different contexts. Moreover they show that in percentage terms the average female share is very closely bounded around 50%, although there is substantial heterogeneity around that point. However, it is impressive how tightly bounded the sharing rule is throughout the distribution. In interpreting this result one needs to be careful because it is full income that is being shared equally. This measure of income includes both leisure and consumption. Thus the share of a woman with a high wage who does not work will include her leisure and her consumption; hence a 50% share may in certain cases hide very unequal levels of consumption of all other goods.

In the final part of the analysis the authors use their estimates to carry out a poverty analysis. The idea here is similar to that in Dunbar et al (2013) described earlier: they compare poverty rates implied by household level income and those implied by individual allocations. The household poverty line is 60% of median household income while the individual poverty line is set at half this amount. This of course is an income based and not a welfare based measure and ignores any household economies of scale. This point not withstanding the individual rates are higher: while household poverty is 11% individual poverty is bounded between 16% and 21%, the lower bound being above the household number. Interestingly the bounds do not differ by gender by any substantive amount.

The Cherchye et al study breaks new ground and shows the power of the collective approach. Specifically it reinforces the identifiability results substantially by showing not only that the levels of the sharing rule can be identified, but more importantly in our view, that the entire sharing rule can also be bounded without much more than within household Pareto efficiency. Nevertheless there is still a long and important agenda in this research. First, empirically we need to understand better how to deal with heterogeneity in preferences within such a non-parametric framework as well as with endogeneity of prices and wages. The entire analysis of Cherchye et al is based on the assumption that wages and prices are exogenous. This is internally consistent with the absence of heterogeneity and shocks, but is is broadly unsatisfactory. For example there is a vast labor supply literature dealing with endogenous wage rates. Moreover, prices of goods may not be exogenous if there are aggregate shocks to the demand functions. While these seem to be side issues as far as the central identifiability of the collective model is concerned they are important for the ultimate empirical credibility of the approach.

7 Conclusion

Understanding intrahousehold inequality and more broadly intrahousehold allocations is crucial for understanding the effects of policy and for targeting programs designed to alleviate poverty. The implications are far reaching and they span simple questions of who will benefit from certain programs to deeper questions about child poverty and even child development. It is now well understood that treating households as an individual unit does not just provide an incomplete picture of standards of living but can be seriously misleading when we try and understand behavior and its reactions to the environment. In our review we have discussed both the questions underlying the notion of intrahousehold inequality as well as the extent of our ability to identify what goes on in the household from typically observed data. In this context we have argued that it is important to be able to observe variables that shift the bargaining power of spouses without affecting preferences as well as other approaches to peeking inside the household black box. It is evident from this discussion that better data would be important; and nothing is more important that detailed consumption and time use data. A renewed emphasis on such data is called for, given the importance of the issues at hand. A better understanding of what may constitute distribution factors and indeed experimental evidence would be an important way to support research into intrahousehold allocations.

However beyond the above, research is now advancing into the dynamics of intrahousehold allocations and being linked to marriage markets. It is now becoming clear how the conditions at the time of marriage can affect intrahousehold allocations. Indeed, under full commitment, current distribution factors may have little to do with current allocations. On the other hand full commitment is a very strong and some may argue an implausible assumption. Thus research is also advancing in understanding how allocations are determined when commitment is limited. In such limited commitment environments changes in the institutional framework, may it be the structure of the welfare system or divorce laws will also have important implications for intrahousehold inequality as well as for the formation and dissolution of marriages. We thus are acquiring a rich theoretical and empirical framework that will allow us to better understand how individual welfare is determined within the context of the family. Important contributions in understanding the dynamics of intrahousehold allocations and of household formation include papers by Mazzocco (2007) and Voena (2013). We are convinced that this is a crucial direction for future research.

References

[1] Afriat, S. N. (1973): "On a System of Inequalities in Demand Analysis: An Extension of the Classical Method," International Economic Review, 14, 460-472

[2] Attanasio, O. and V. Lechene (2014) "Efficient Responses to Targeted Cash Transfers" Journal of Political Economy

[3] Banks, James, Richard W. Blundell and Arthur Lewbel (1997) Quadratic Engel Curves and Consumer Demand The Review of Economics and Statistics, Vol. 79, No. 4 (Nov., 1997), pp. 527-539

[4] Bargain, O., M. Beblo, D. Beninger, R. Blundell, R. Carrasco, M.-C. Chiuri, F. Laisney, V. Lechene, N. Moreau, M. Myck, J. Ruiz-Castillo and F. Vermeulen (2006), "Does the representation of household behavior matter for welfare analysis of tax-benefit policies? An introduction", *Review of Economics of the Household*, 4, 99-111.

[5] Basu K., 2006, "Gender and Say: A Model of Household Behavior with Endogenously-determined Balance of Power", Economic Journal, vol. 116, pp. 558-580.

[6] Becker, Gary S. (1974) A Theory of Social Interactions Journal of Political Economy, Vol. 82, No. 6 (Nov. - Dec., 1974), pp. 1063-1093

[7] Beninger, D., O. Bargain, M. Beblo, R. Blundell, R. Carrasco, M.-C. Chiuri, F. Laisney, V. Lechene, E. Longobardi, N. Moreau, M. Myck, J. Ruiz-Castillo and F. Vermeulen (2006), "Evaluating the move to a linear tax system in Germany and other European countries", *Review of Economics of the Household*, 4, 159-180.

[8] Blundell, Richard W. and Martin Browning and Ian A. Crawford (2003) Nonparametric Engel Curves and Revealed Preference Econometrica, Vol. 71, No. 1 (Jan., 2003), pp. 205-240

[9] Blundell, Richard W. , Martin Browning and Ian A. Crawford (2008)Best Nonparametric Bounds on Demand Responses Econometrica Vol. 76, No. 6 (Nov., 2008), pp. 1227-1262

[10] Blundell, Richard W., Pierre-André Chiappori, Thierry Magnac and Costas Meghir (2007) Collective Labour Supply: Heterogeneity and Non-Participation The Review of Economic Studies, Vol. 74, No. 2 (Apr., 2007), pp. 417-445

[11] Blundell, Richard W., Pierre André Chiappori, and Costas Meghir (2005) Collective Labor Supply with Children Journal of Political Economy, Vol. 113, No. 6 (December 2005), pp. 1277-1306

[12] Blundell, Richard W., Alan Duncan and Costas Meghir (1998) Estimating Labor Supply Responses Using Tax Reforms Econometrica, Vol. 66, No. 4 (Jul., 1998), pp. 827-861

[13] Blundell, Richard, W., Joel L. Horowitz and Matthias Parey (2012), "Measuring the price responsiveness of gasoline demand: Economic shape restrictions and nonparametric demand estimation", Quantitative Economics Volume 3, Issue 1, pages 2951, March 2012

[14] Blundell, Richard W., and Ian Walker (1986) A Life-Cycle Consistent Empirical Model of Family Labour Supply Using Cross-Section Data The Review of Economic Studies, Vol. 53, No. 4, Econometrics Special Issue (Aug., 1986), pp. 539-558

[15] Browning, Martin and P. A. Chiappori (1998) Efficient Intra-Household Allocations: A General Characterization and Empirical Tests Econometrica, Vol. 66, No. 6 (Nov., 1998), pp. 1241-1278

[16] Browning, Martin, P.A. Chiappori, and V. Lechene, 2008, "Distributional Effects in Household Models: Separate Spheres and Income Pooling", *Economic Journal* (forthcoming).

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

[18] Browning, Martin, P.A. Chiappori, and Y. Weiss, 2013, *Family Economics,* Cambridge University Press.

[19] Cherchye, Laurens, Bram de Rock and Frederic Vermeulen (2007) The Collective Model of Household Consumption: A Nonparametric Characterization Econometrica, Vol. 75, No. 2 (Mar., 2007), pp. 553-574

[20] Cherchye, L.; B. De Rock and F. Vermeulen (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), 1074-104.

[21] Cherchye, Laurens, J. H. ,B. de Rock, Arthur Lewbel and F. M. P. Vermeulen, (2012) , "Sharing Rule Identification for General Collective Consumption Models", 2012, Tilburg University, Center for Economic Research, Discussion Paper: 2012-041

[22] Chiappori, Pierre-André (1988) Rational Household Labor Supply Econometrica, Vol. 56, No. 1 (Jan., 1988), pp. 63-90

[23] Chiappori, Pierre-André (1991) Nash-Bargained Household Decisions: A Rejoinder International Economic Review, Vol. 32, No. 3 (Aug., 1991), pp. 761-762

[24] Chiappori, Pierre-André (1992) Collective Labor Supply and Welfare Journal of Political Economy, Vol. 100, No. 3 (Jun., 1992), pp. 437-467

[25] Chiappori, Pierre-André (2012) ”Modèles d'appariement en conomie: quelques avances rcentes”, Texte de la Confrence Jean-Jacques Laffont, Revue Economique, Vol. 63 , 2012/3, 437-452

[26] Chiappori, Pierre-André, O. Donni and I. Komunjer (2012) ”Learning from a Piece of Pi” Review of Economic Studies (2012) 79 (1): 162-195.

[27] Chiappori, Pierre-André, and I. Ekeland, 2009 a, "The Micro Economics of Efficient Group Behavior: Identification", *Econometrica*, 77 3, 763-99.

[28] Chiappori, Pierre-André, and I. Ekeland, 2009b, *The Economics and Mathematics of Aggregation*, Foundations and Trends in Microeconomics, Now Publishers, Hanover, USA

[29] Chiappori, Pierre-André, Bernard Fortin, and Guy Lacroix (2002) Marriage Market, Divorce Legislation, and Household Labor Supply Journal of Political Economy, Vol. 110, No. 1 (February 2002), pp. 37-72

[30] Chiappori, Pierre-André and Kim, J-H, 2013, "Identifying Heterogeneous Sharing Rules", mimeo, Columbia University

[31] Chiappori, P.-A., McCann, R., Nesheim, L. (2010) ”Hedonic price equilibria, stable matching, and optimal transport: equivalence, topology, and uniqueness," Economic Theory, 42(2):317–354.

[32] Chiappori, Pierre-André and Costas Meghir (2014) ”Intrahousehold Welfare”, mimeo

[33] Chiappori, Pierre-André and B. Salanié (2014) "The Econometrics of Matching Models", forthcoming *Journal of Economic Literature*

[34] Donni, Olivier. (2003). "Collective Household Labor Supply: Non-Participation and Income Taxation." *Journal of Public Economics* 87, 1179-1198.

[35] Dunbar, Geoffrey R., Arthur Lewbel, Arthur and Krishna Pendakur (2013) "Children's Resources in Collective Households: Identification, Estimation, and an Application to Child Poverty in Malawi" The American Economic Review, Volume 103, Number 1, February 2013 , pp. 438-471(34)

[36] Donni O., 2009, "A Simple Approach to Investigate Intra-household Allocation of Private and Public Goods", *Review of Economics and Statistics*, vol. 91, pp. 617-628.

[37] Edlund, Lena, and E. Korn, 2002, "A Theory of Prostitution", *Journal of Political Economy*, 2002, vol. 110, no. 1, 181-214.

[38] Fleurbaey, M., K. Decancq and E. Schokkaert (2014), Inequality, income and well-being, this Handbook

[39] Gersbach H. and H. Haller, 2001, "Collective Decisions and Competitive Markets", *Review of Economic Studies*, vol. 68, pp. 347–368.

[40] Goussé (2013) Marriage Market and Intra-Household Allocation : Evolution of preferences and transfers in the UK from 1991 to 2008, mimeo Sciences Po

[41] Grossbard-Shechtman Shoshana, *On the Economics of Marriage: A Theory of Marriage, Labor, and Divorce*, (Boulder: Westview Press, 1993).

[42] Haddad, L., and R. Kanbur (1990), 'How serious is the neglect of intra-household inequality', The Economic Journal, 100, 866-81.

[43] Haddad, L., and R. Kanbur (1992), 'Intrahousehold Inequality and the Theory of Targeting', European economic review, 36, 372-8

[44] Jacquemet, Nicolas and Jean-Marc Robin (2011) Marriage with Labor Supply, CES Working Papers 2011.50

[45] Lechene, Valérie and Ian Preston. 'Noncooperative Household Demand' (2011), Journal of Economic Theory, 146(2), 504-27.

[46] Lewbel, A., and Krishna Pendakur, 2013, "Unobserved Preference Heterogeneity in Demand Using Generalized Random Coefficients,"Mimeo, Boston College.

[47] Lise, Jeremy and Shannon Seitz (2011) Consumption Inequality and Intrahousehold Allocations The Review of Economic Studies, Vol. 78, No. 1 (January 2011), pp. 328-355

[48] Lundberg, Shelly and Robert A. Pollak (1993) Separate Spheres Bargaining and the Marriage Market Journal of Political Economy, Vol. 101, No. 6 (Dec., 1993), pp. 988-1010

[49] Manser, Marilyn and Murray Brown (1980), 'Marriage and Household Decision-Making: A Bargaining Analysis', *International Economic Review,* 21 (1980), 31-44.

[50] Mazzocco, M. (2007) "Household Inter Temporal Behaviour: A Collective Characterization and a Test of Commitment" *The Review of Economic Studies* Vol. 74, No. 3 (Jul., 2007) , pp. 857-895

[51] Mcelroy, Marjorie B. and Horney, Mary Jean, 'Nash-Bargained Household Decisions: Toward a Generalization of the Theory of Demand', *International Economic Review,* 22 (1981), 333-349.

[52] Myck, M., O. Bargain, M. Beblo, D. Beninger, R. Blundell, R. Carrasco, M.-C. Chiuri, F. Laisney, V. Lechene, E. Longobardi, N. Moreau, J. Ruiz-Castillo and F. Vermeulen (2006), "The Working Families' Tax Credit and some European tax reforms in a collective setting", *Review of Economics of the Household*, 4, 129-158.

[53] Ponthieux, S., and D. Meurs (2014), Gender Inequality, this Handbook

[54] Samuelson, Paul A. (1956) Social Indifference Curves The Quarterly Journal of Economics, Vol. 70, No. 1 (Feb., 1956), pp. 1-22

[55] Ulph D., 2006, "Un modèle non-coopératif de Nash appliqué à l'étude du comportement de consommation du ménage", *Actualité économique: revue d'analyse économique*, vol. 82, pp. 53-86.

[56] Varian, H. (1982): "The Nonparametric Approach to Demand Analysis," Econometrica, 50, 945-974

[57] Vermeulen, F., O. Bargain, M. Beblo, D. Beninger, R. Blundell, R. Carrasco, M.-C. Chiuri, F. Laisney, V. Lechene, N. Moreau, M. Myck and J. Ruiz-Castillo (2006), "Collective models of labor supply with nonconvex budget sets and nonparticipation: a calibration approach", *Review of Economics of the Household*, 4, 113-127.

[58] Voena, A. (2013) "Yours, Mine and Ours: Do Divorce Laws Affect the Intertemporal Behavior of Married Couples?", mimeo University of Chicago