The Economics of the Family

Chapter 5: Empirical issues for the collective model

Martin Browning
Department of Economics, Oxford University

Pierre-André Chiappori
Department of Economics, Columbia University

Yoram Weiss
Department of Economics, Tel Aviv University

August 2010

1 What are the objects of interest?

We have seen above that various approaches can be used to describe household behavior, from the unitary setting to noncooperative approaches and the collective model. Ultimately, the choice between these various frameworks will rely on particular considerations. First, general methodological principles may favor one approach over the others. For instance, one can argue that the unitary framework is not totally faithful to methodological individualism, a cornerstone of micro theory that postulates that individuals, not groups, are the ultimate decision makers. A second requirement is the model's ability to generate testable predictions for observable behavior, that can be taken to data using standard techniques. Standard consumer theory fares pretty well in this respect. Utility maximization under a linear budget constraint yields strong predictions (adding-up, homogeneity, Slutsky symmetry and negative semidefiniteness and income pooling) and adequate methodologies have been developed for testing these properties. Finally, a crucial criterion is the fruitfulness of the approach, particularly in terms of normative analysis and policy recommendations. A remarkable feature of standard consumer theory is that individual preferences can be uniquely recovered from demand functions (if these satisfy the Slutsky conditions); it is therefore possible to analyze welfare issues from the sole knowledge of observed behavior. This is a particular case of the general requirement that the model be *identifiable*, that is, that it should be possible to recover the underlying structure from observed behavior.

The first line of argument, concerning methodological individualism, has been evoked earlier. In this chapter, we concentrate on the remaining two aspects, namely testability and identifiability of preferences and processes from observed behavior. Most of the existing knowledge for non-unitary models concerns the cooperative framework, and especially the collective model. The testability requirement, *per se*, is not problematic. The idea that a model should generate predictions that can be taken to data belongs to the foundations of economics (or any other science!). Identifiability is more complex and it is useful to define more precisely what is meant by 'recovering the underlying structure'. The structure, in our case, is the (strictly convex) preferences of individuals in the group and the decision process. In the collective setting, because of the efficiency assumption, the decision process is fully summarized (for any particular coordinatization of individual utilities) by the Pareto weight corresponding to the outcome at stake. The structure thus consists in a set of individual preferences (with a particular coordinatization) and a Pareto weight - which, as we should remember, can be (and generally is) a function of prices, incomes and distribution factors.

The structure cannot be directly observed; instead we observe the outcomes of the interactions between preferences, constraints and the decision process. Often we observe only aggregate outcomes and not individual outcomes. In addition, the 'observation' of, say, a demand function is a complex process, that entails specific difficulties. For instance, one never observes a (continuous) function, but only a finite number of values on the function's graph. These values are measured with some errors, which raises problems of statistical inference. In some cases, the data are cross-sectional, in the sense that different groups are observed in different situations; specific assumptions have to be made on the nature and the form of (observed and unobserved) heterogeneity between the groups. Even when the same group is observed in different contexts (panel data), other assumptions are needed on the dynamics of the situation - for example, on the way past behavior influences present choices. All these issues lay at the core of what is usually called the *inference* problem.[1]

A second and different aspect relates to what has been called the *identifiability* problem, which can be defined as follows: when is it the case that the *(hypothetically) perfect knowledge* of a smooth demand function uniquely defines the underlying structure within a given class? This abstracts from the econometrician's inability to exactly recover the form of demand functions - say, because only noisy estimates of the parameters can be obtained, or even because the functional form itself (and the stochastic structure added to it) have been arbitrarily chosen. These econometric questions have, at least to some extent, econometric or statistical answers. For instance, confidence intervals can be computed for the parameters (and become negligible when the sample size grows); the relevance of the functional form can be checked using specification tests; etc. The non-identifiability problem has a different nature: even if a *perfect* fit to *ideal* data was feasible, it might still be impossible to recover the underlying structure from this ideal data.

In the case of *individual* behavior, as analyzed by standard consumer theory, identifiability is an old but crucial result. Indeed, it has been known for more than a century that an individual demand function uniquely identifies the underlying preferences. Familiar as this property may have become, it remains one of the strongest results in microeconomic theory. It implies, for

[1] In the original Koopmans discussion of identification, the step from sample information to inferences about population objects (such as demand functions) is referred to as identification. Here we follow modern terminology and refer to it as the inference step.

instance, that assessments about individual well-being can unambiguously be made based only on the observation of demand behavior with sufficiently rich (and 'exogenous') variation in prices and total expenditures; a fact that opens the way to all of applied welfare economics. It is thus natural to ask whether this classical identifiability property can be extended to more general approaches.[2]

Finally, it should be remembered that identifiability is only a necessary condition for identification. If different structures are observationally equivalent, there is no hope that observed behavior will help to distinguish between them; only ad hoc functional form restrictions can do that. Since observationally equivalent models may have very different welfare implications, non-identifiability severely limits our ability to formulate reliable normative judgments: any normative recommendation based on a particular structural model is unreliable, since it is ultimately based on the purely arbitrary choice of one underlying structural model among many. Still, whether an *identifiable* model is econometrically *identified* depends on the stochastic structure representing the various statistical issues (measurement errors, unobserved heterogeneity,...) discussed above. After all, the abundant empirical literature on consumer behavior, while dealing with a model that is always identifiable, has convinced us that identification crucially depends on the nature of available data.

The main properties of the collective model have been described in the previous chapter. However, which empirical test can actually be performed obviously depends on the nature of available data. Three different contexts can be distinguished. In the first context, individual demand can be estimated as a function of income and possibly distribution factors; this approach is relevant when no price variation is observed, for instance because data are cross-sectional and prices are constant over the sample. We then allow that we also observe price variation so that we can estimate a complete demand system. The analysis of labor supply raises specific issues that are considered in the third section. The final half of this chapter presents a review of empirical analysis using non-unitary models (including the results of applying the tests of the first half of the chapter).

[2] Note, however, that only one utility function is identifiable in the standard case. In a 'unitary' framework in which agents are characterized by their own utility function (see chapter 3, subsection 5.8) but the household behaves as a single decision unit, it is typically *not* possible to identify the individual utility functions.

2 Data without price variation.

2.1 Necessary and sufficient conditions for a collective model.

In this section we consider testing and identification in the absence of price variation as is often the case with cross-sectional data. We begin with the case in which we observe only household (aggregate) demand of each good. Let x denote the household's total expenditures and let $\mathbf{z}$ be a K-vector of distribution factors. Recall that distribution factors, by definition, influence neither preferences nor the budget constraint. In a unitary setting, they have no impact on demand. In the collective framework, on the contrary, household behavior can be described by a program of the following form:

$$\begin{gathered} \max \mu\left(x, \mathbf{z}\right) u^{a}\left(\mathbf{g}\right)+u^{b}\left(\mathbf{g}\right) \\ \text{subject to } \mathbf{e}^{\prime} \mathbf{g} \leq x \end{gathered} \tag{1}$$

where $\mathbf{g}$ is the vector $\left(\mathbf{q}^{a}, \mathbf{q}^{b}, \mathbf{Q}\right)$ and quantities are normalized so that the price vector is a vector of ones, $\mathbf{e}$. The resulting vector of *collective demand functions* can be written $\mathbf{g} = \tilde{g}(x, \mu\left(x, \mathbf{z}\right))$ with a corresponding observable demand functions $\hat{\mathbf{g}}\left(x, \mathbf{z}\right)$.

An alternative demand formulation which is useful for empirical work (see below) can be formulated if there is at least one good (good j, say) that is strictly monotone in one distribution factor (z_1, say); that is, $g_j\left(x, \mathbf{z}\right)$ is strictly monotone in z_1. This demand function can be inverted on the first factor to give:

$$z_1 = \zeta(x, \mathbf{z}_{-1}, g_j)$$

where $\mathbf{z}_{-1}$ is the vector of distribution factors without the first element. Now substitute this into the demand for good i:

$$g_i = \hat{g}_i(x, z_1, \mathbf{z}_{-1}) = \hat{g}_i[x, \zeta(x, \mathbf{z}_{-1}, g_j), \mathbf{z}_{-1}] = \theta_i^j\left(x, \mathbf{z}_{-1}, g_j\right).$$

Thus the demand for good i can be written as a function of total expenditure, all distribution factors but the first and the demand for good j. To distinguish this conditioning from the more conventional conditional demands used in the demand literature, we shall refer to them as *z-conditional* demands.[3]

We now address the issue of what restrictions a collective model imposes upon observable demands. Bourguignon, Browning and Chiappori (2009)

[3] In the unitary setting, distribution factors cannot influence demand, so that z-conditional demands are not defined in this case.

provide a complete characterization of these conditions. Specifically they prove that the following equivalent conditions are necessary consequences of the collective model:

1. there exist real valued functions $\tilde{g}_1, ..., \tilde{g}_n$ and μ such that :
$$\hat{g}_i(x, \mathbf{z}) = \tilde{g}_i[x, \mu(x, \mathbf{z})] \qquad \forall i = 1, ..., n \tag{2}$$

2. household demand functions satisfy the *proportionality condition*:
$$\frac{\partial \hat{g}_i / \partial z_k}{\partial \hat{g}_j / \partial z_k} = \frac{\partial \hat{g}_i / \partial z_1}{\partial \hat{g}_j / \partial z_1} \qquad \forall i = 1, .., n;\ j = 1, .., n;\ k = 2, .., K \tag{3}$$

3. for any good j such that $\partial \hat{g}_j / \partial z_1 \neq 0$, the z-conditional demands satisfy:
$$\frac{\partial \hat{\theta}_i^j (x, \mathbf{z}_{-1}, g_j)}{\partial z_k} = 0 \qquad \forall i \neq j, k = 2, .., K \tag{4}$$

The intuition for this result relates to the discussion provided in earlier chapters. Again, the basic idea is that, by definition, distribution factors do not influence the Pareto set. They may affect consumption, but only through their effect upon the *location* of the final outcome on the Pareto frontier or, equivalently, upon the *respective weighting* of each member's utility that is implicit in this location. The key point is that this effect is one-dimensional (see chapter 4, subsection 1.3). This explains why restrictions appear only in the case where there is more than one distribution factor. Whatever the number of such factors, they can only influence consumption through a single, real-valued function μ. Conditions (2) and (3) are direct translations of this remark. By the same token, if we compute q_i as a z-conditional function of $(x, \mathbf{z}_{-1}, g_j)$, it should not depend on $\mathbf{z}_{-1}$. The reason is that, for any given value of x, whenever distribution factors $(z_1, \mathbf{z}_{-1})$ contain some information that is relevant for intra-household allocation (hence for household behavior), this information is one-dimensional and can be *fully* summarized by the value of g_j. Once we condition on g_j, $\mathbf{z}_{-1}$ becomes irrelevant. This is the meaning of condition (4).

The conditions (2)-(4) are also sufficient for the collective model: if they are satisfied for the observable demands $\hat{g}(x, \mathbf{z})$, then one can find utility functions and Pareto weights which rationalize the observed demands (see Bourguignon *et al* (2009)). An important implication of these conditions is that in the absence of price variation, proportionality is the *only* testable implication of the collective model. This means that if we have only one distribution factor, then we can never reject the hypothesis of collective rationality. Any extra restrictions for a collective model require that *additional assumptions* be made on the form of individual preferences. For instance, restrictions exist even for a single distribution factor when some goods are private and/or are consumed exclusively by one member of the household. It may surprise readers that in the absence of price variation, proportionality is the full empirical content of the collective model. Recall, however, that in the unitary model, *without price variation*, any demands as a function of total expenditure are compatible with utility maximization.

This result provides two distinct ways of testing for efficiency. Condition (3) leads to tests of cross-equation restrictions in a system of unconditional demand equations. An alternative method, implied by (4), tests for exclusion restrictions in a conditional demand framework. Empirically, the latter is likely to be more powerful for at least two reasons. First we can employ single equation methods (or even non-parametric methods). Second, single equation exclusion tests are more robust than tests of the equality of parameters across equations. Both tests generalize easily to a framework in which domestic goods are produced by the household. Adding a domestic production function that relates market inputs and domestic labor to goods actually consumed by household members does not modify the above tests on household demands for market goods.

As discussed in Chapter 3, the bargaining version of the collective model has attracted lot of attention. A bargaining framework should be expected to impose *additional* restrictions to those discussed above. Indeed, an easy test can be described as follows. Assume that some distribution factors, which are part of a K'-sub-vector $\mathbf{z}'$, are known to be positively correlated with member b's threat point, while others, constituting a K''-sub-vector $\mathbf{z}''$, are known to favor a. Then in program (1) μ should decrease with distribution factors in $\mathbf{z}''$ and increase with those in $\mathbf{z}'$. This property can readily be tested; it implies that,

$$\frac{\partial \hat{g}_i / \partial z'_k}{\partial \hat{g}_i / \partial z''_m} = \frac{\partial \hat{g}_j / \partial z'_k}{\partial \hat{g}_j / \partial z''_m} \leq 0 \text{ for } i, j = 1, .., n; k = 1, ..., K'; m = 1, ..., K''$$

Should one be willing to go further and assume, for instance, that only the ratio z'_1 / z''_2 of distribution factors matters, then we have in addition:

$$\frac{\partial \hat{g}_i}{\partial \ln(z'_1)} + \frac{\partial \hat{g}_i}{\partial \ln(z''_2)} = 0 \qquad \forall i = 1, .., n$$

This is simple to test and easy to interpret.

2.2 Identifiability.

A more difficult issue arises when we consider identifiability. That is, when is it possible to recover the underlying structure from the sole observation of household behavior? Note, first, that the nature of the data strongly limits what can be recovered. For instance, one cannot hope to identify utility functions in the absence of price variations. 'Identifiability', in this context, essentially means recovering individual Engel curves (that is, demand as a function of income) and the decision process, as summarized by the Pareto weights or (in the private good case) by the sharing rule, again as functions of income and distribution factors only.

With these precautions in mind, we start with some mathematical results concerning integrability that are useful in the current context. Suppose we have a smooth unknown function $f(x, y)$ with non-zero partials f_x and f_y. Suppose first that we observe:

$$h^1(x, y) = f_x(x, y) \text{ and } h^2(x, y) = f_y(x, y) \tag{5}$$

If $f(.)$ is twice continuously differentiable, these two functions must satisfy the cross derivative restriction $h^1_y(x, y) = h^2_x(x, y)$. In general, these conditions can be translated into empirical tests of the hypothesis that $h^1(.)$ and $h^2(.)$ are indeed partials of the same function. Moreover, if this symmetry condition is satisfied, then $f(.)$ is identifiable up to an additive constant. That is, if $\bar{f}$ is a solution of (5), then any alternative solution must be of the form $\bar{f}(x, y) + k$ where k is an arbitrary constant.

Suppose now that rather than observing the partials themselves we only observe their ratio:

$$h(x, y) = \frac{f_y}{f_x} \tag{6}$$

Given $h(x, y)$, $f(x, y)$ is identifiable 'up to a strictly monotone transformation'. That is, we can recover some $\bar{f}(x, y)$ such that any solution is of the form $f(x, y) = G\left(\bar{f}(x, y)\right)$ where $G(.)$ is an arbitrary strictly monotone function.

In general, when $f(.)$ has more than two arguments, $f(x_1, ..., x_n)$, assume that we observe $m < n - 1$ ratios of partials, say those involving the $m+1$ first partials of f : $\frac{f_2}{f_1}, ..., \frac{f_{m+1}}{f_1}$. Then f is identifiable up to a function of the other variables. That is, we can identify some $\bar{f}(x_1, ..., x_{m+1})$ such that any solution is of the form

$$f(x_1, ..., x_n) = G\left(\bar{f}(x_1, ..., x_{m+1}), x_{m+2}..., x_n\right)$$

where $G(.)$ is an arbitrary function. In particular:

- if we observe only one ratio of partials, say $h(x_1, ..., x_n) = f_1/f_2$, then $f(.)$ is identifiable up to a function of the other variables $(x_3, ..., x_n)$.
- if we observe all ratios of partials, then $f(.)$ is identifiable up to an arbitrary, strictly monotone transformation.

Note, as well, that whenever we observe more than one ratio of partials, testable restrictions are generated. These generalize the previous cross-derivative conditions.

We can now return to the identifiability problem for the collective model. Even in the most general case (no identifying restriction beyond efficiency), some (but by no means all) of the structure can be recovered from the observation of demand functions. To see why, note that by equation (3) we have:

$$\frac{\partial \hat{g}_i/\partial z_k}{\partial \hat{g}_i/\partial z_1} = \frac{\partial \mu/\partial z_k}{\partial \mu/\partial z_1} = \frac{\mu_k}{\mu_1} = \kappa_k \text{ for all } i \text{ and } k \tag{7}$$

The left hand side expression is potentially observable so that we can identify the ratio of partials of $\mu(x, \mathbf{z})$ with respect to distribution factors. Since the right hand side does not depend on the good, the ratio on the left hand side must be the same for all goods; this is the proportionality condition. Given the ratio of partials of the Pareto weight, we can recover $\mu(.)$ up to some function of x. That is, we can recover a particular Pareto weight $\bar{\mu}$ such the true Pareto weight μ must be of the form:

$$\mu(x, \mathbf{z}) = m(x, \bar{\mu}(x, \mathbf{z})) \tag{8}$$

for some unobserved function $m(.)$.

The ratio κ_k in equation (7) has a natural interpretation in terms of *power compensation*. Assume, for instance, that $\mu_1 > 0$ and $\mu_k < 0$ so that z_1 favors a while z_k serves b. If z_k is increased by some infinitesimal quantity dz_k then $\kappa_k dz_k$ is the increase in z_1 required to offset the change and maintain the same balance of power. Power compensations may be important for welfare analysis, whenever a 'shift of power' has to be compensated. The good news is that even in the most general version of the collective model, they can be directly recovered from observed demands. Furthermore, the proportionality condition (3) imposes that the estimation of the power compensation ratio does not depend on the particular commodity chosen. An alternative and important interpretation of this result is that the model always behaves 'as if' there were only one factor, $\bar{\mu}$, influencing the individual's relative powers. Whatever the actual number of distribution factors, they always operate through the index $\bar{\mu}$. Moreover, this index is identifiable.

What is not identifiable in the general case is the exact impact of the index on the actual Pareto weight; an impact that will in general depend on the level of total expenditures.

2.3 Private consumption.

Although useful, recovery of the Pareto weight up to a strictly monotone function that also depends on total expenditure is far short of what is needed for some important purposes. Is it possible to recover more? To achieve this, we need either better data or more theory restrictions. As an example of the latter, consider the particular but useful case in which all commodities are privately consumed and preferences are either egoistic or caring. As we have seen in chapter 4, efficiency is then equivalent to the existence of a sharing rule in which a receives $\rho(x, \mathbf{z})$ and b receives $(x - \rho(x, \mathbf{z}))$. Individual a solves:

$$\max v^a(\mathbf{q}^a) \text{ subject to } \mathbf{e}'\mathbf{q}^a = \rho(x, \mathbf{z}) \tag{9}$$

and similarly for b. It follows that the household aggregate demand for commodity i takes the form:

$$q_i(x, \mathbf{z}) = q_i^a(\rho(x, \mathbf{z})) + q_i^b(x - \rho(x, \mathbf{z}))$$

where q_i^s is s's demand for good i. The question is: what can be said about q_i^a, q_i^b and ρ from the observation of household demands q_i for $i = 1, ...n$.

Equation (7) has an equivalent in this context:

$$\frac{\partial q_i / \partial z_k}{\partial q_i / \partial z_1} = \frac{\partial \rho / \partial z_k}{\partial \rho / \partial z_1} \text{ for all } k \tag{10}$$

This result remains valid in the presence of public goods, provided that the sharing rule is taken to be conditional on public goods (as described in subsection 5.2 of Chapter 4). The potential observability of the left hand side of equation (10) means that we can recover the sharing rule up to an arbitrary monotone function of total expenditures x. In other words, we can recover some $\bar{\rho}(x, \mathbf{z})$ such that the true sharing rule must be of the form $\rho(x, \mathbf{z}) = G(\bar{\rho}(x, \mathbf{z}), x)$ for some mapping G. And, as above, instead of analyzing the impact of each distribution factor independently, we may just consider the impact of the 'index' $\bar{\rho}$. Consequently we can always consider the case of a unique distribution factor; no loss of generality results.

2.4 Assignability.

Up until now we have considered the case where we only observe aggregate household demands. In some cases, we can observe the consumption of a particular good by each partner. That is, for some goods we observe q_i^a and q_i^b. We refer to such a good as being *assignable.* The most widely used example of an assignable good is clothing: in expenditure surveys we always see a distinction made between men and women's clothing. An alternative terminology is that each of the clothing commodities is an *exclusive* good.[4] That is, an exclusive good is one that is consumed by a unique person in the household.

Suppose that we observe the individual consumption of the first good and estimate $\hat{q}_1^a(x, \mathbf{z})$ and $\hat{q}_1^b(x, \mathbf{z})$. Assuming, without loss of generality, that there is only one distribution factor, the collective demands $\tilde{q}_1^s$ are related to the observable demands $\hat{q}_1^s$ by:

$$\hat{q}_1^a(x, z) = \tilde{q}_1^a(\rho(x, z)) \tag{11}$$

$$\hat{q}_1^b(x, z) = \tilde{q}_1^a(x - \rho(x, z)) \tag{12}$$

Thus:

$$\begin{aligned} \frac{\partial \hat{q}_1^a / \partial x}{\partial \hat{q}_1^a / \partial z} &= \frac{\rho_x}{\rho_z} \\ \frac{\partial \hat{q}_1^b / \partial x}{\partial \hat{q}_1^b / \partial z} &= -\frac{1 - \rho_x}{\rho_z} \end{aligned} \tag{13}$$

Thus the two ratios ρ_x / ρ_z and $(1 - \rho_x) / \rho_z$ are identifiable. There is a unique solution to these two equations for (ρ_x, ρ_z) if and only if:

$$\Gamma = \frac{\partial \hat{q}_1^a}{\partial x} \frac{\partial \hat{q}_1^b}{\partial z} - \frac{\partial \hat{q}_1^b}{\partial x} \frac{\partial \hat{q}_1^a}{\partial z} \neq 0 \tag{14}$$

If this condition holds, we can identify the partials of ρ:

$$\begin{aligned} \rho_x &= \frac{1}{\Gamma} \frac{\partial \hat{q}_1^a}{\partial x} \frac{\partial \hat{q}_1^b}{\partial z} \\ \rho_z &= \frac{1}{\Gamma} \frac{\partial \hat{q}_1^a}{\partial z} \frac{\partial \hat{q}_1^b}{\partial z} \end{aligned} \tag{15}$$

[4] In general, individual consumptions of an assignable good have the same price, whereas exclusive goods have different prices. The distinction is ineffective in the present context, but will become important when price variations are considered.

By the result before (6), knowing the partials allows us to identify the function itself, up to an additive constant: $\rho = \rho(x, z) + k$. Thus we can learn everything about the sharing rule from observing the assignment of a single good, *except its location.* One good is sufficient because the same Pareto weight function appears in all goods; see equation (2). Moreover, new restrictions are generated, since

$$\frac{\partial}{\partial z}(\rho_x) = \frac{\partial}{\partial x}(\rho_z)$$

This provides a test for assignability of any particular good within the collective setting.

Finally, what about the individual Engel curves of the two spouses? First, for any value of the constant k, (11) and (12) identify individual demands for commodity 1. Consider, now, commodity i; remember that, in general, i is neither exclusive nor assignable. Still, from:

$$\hat{q}_i(x, \mathbf{z}) = \tilde{q}_i^a(\rho(x, \mathbf{z})) + \tilde{q}_i^b(x - \rho(x, \mathbf{z})) \tag{16}$$

we have:

$$\begin{aligned} \frac{\partial \hat{q}_i}{\partial x} &= \frac{d\tilde{q}_i^a}{d\rho}\rho_x + \frac{d\tilde{q}_i^b}{d\rho}(1 - \rho_x) \\ \frac{\partial \hat{q}_i}{\partial z} &= \left(\frac{d\tilde{q}_i^a}{d\rho} - \frac{d\tilde{q}_i^b}{d\rho}\right)\rho_z \end{aligned} \tag{17}$$

Since the left hand side is observed and we have (ρ_x, ρ_z) we invert (so long as $\rho_z \neq 0$) and identify $\tilde{q}_a^i$ and $\tilde{q}_b^i$ up to an additive constant. We conclude that the presence of an assignable good is sufficient to identify (up to additive constant) the sharing rule and individual demands for each commodity, including the non assignable ones.

We thus get a great deal of mileage from the presence of one assignable (or two exclusive) goods. Can we do without? Surprisingly enough, the answer is positive. Bourguignon, Browning and Chiappori (2009) prove the following strong result: if we observe household demand (as a function of total expenditures x and a distribution factor z) for at least three commodities, then we can recover individual demands and the sharing rule up to the same additive constants as before *and* (this is the only twist) up to a permutation of a and b.[5] This result arises from equation (2) and follows

[5] Identifiability, here, is only 'generic'. It is indeed possible to construct examples in which it does not hold, but these examples are not robust. For instance, if individual demands and the sharing rule are all linear, identification does not obtain. However, adding quadratic terms is sufficient to guarantee identification except maybe for very specific values of the coefficients.

since we have three demands that depend on the one Pareto weight function. For the technical details, see Bourguignon *et al* (2009). The result requires observation of cross partial terms involving x and z; since these are are often difficult to pin down in empirical work, this route for identifying the sharing rule is less robust than using assignability. It is important to note that the identification here does require the existence of at least one distribution factor. Without a distribution factor no information concerning the preferences or the sharing rule can be recovered.

3 Observing price responses.

3.1 Testing the collective model

3.1.1 The basic result

We now turn to the situation in which we observe variation in prices as well as in income and distribution factors. This would be the case, for instance, if we have panel data, or if the cross sectional data exhibit important and *exogenous* fluctuations in prices. Then strong tests are available. Moreover, the model can be proved to be identifiable under reasonably mild exclusion conditions.

Again, we consider a two person household for expositional convenience. Tests of the most general form of the collective model are based on the fundamental SNR1 condition demonstrated in Chapter 4. Namely, the Slutsky matrix S (which can be derived from estimated demand functions) must be of the form:

$$S = \Sigma + R \tag{18}$$

where Σ is symmetric, negative and R is of rank at most one.

Direct tests of (18) are not straightforward, because the theorem simply says that *there exists* such a decomposition. To construct a testable implication of the symmetry of Σ, consider the matrix M defined by:

$$M = S - S'$$

where S' is the transpose of S. Since Σ is symmetric:

$$M = R - R'$$

and since R is of rank (at most) 1, M is of rank (at most) 2. This property is easy to test, using either standard rank tests or more specific approaches.

Note, however, that five commodities (at least) are needed for that purpose. The reason is that neither M nor S are of full rank. Indeed, a standard result of consumer theory, stemming from homogeneity and adding up, states that

$$\boldsymbol{\pi}' S = S\boldsymbol{\pi} = \mathbf{0}$$

where $\boldsymbol{\pi}$ denotes the price vector. It follows that $M\boldsymbol{\pi} = \mathbf{0}$, and M cannot be invertible. Moreover, M is antisymmetric (equal to minus its transpose); hence its rank must be even. With four commodities, M is a 4×4, antisymmetric, non-invertible matrix, so that its rank can never exceed 2 anyway.

Negative semidefiniteness of Σ, on the other hand, can be directly tested on the Slutsky matrix. Indeed, among the eigenvalues of S, one is zero (reflecting non invertibility); among the others, one (at most) can be positive. Therefore, while symmetry of Σ cannot be tested from less than five goods, three are sufficient to test negativeness. In practice, such a test may however not be very powerful. An alternative approach is to use revealed preference techniques; following an early discussion in Chiappori (1988), Cherchye, De Rock and Vermeulen (2007) and (2008) provide a complete characterization of the revealed preference approach to collective models.

3.1.2 Distribution factors

Distribution factors can be readily introduced for parametric approaches. Using equation (**??**) in Chapter 4, Browning and Chiappori (1998) prove the following result. Take any distribution factor k, and compute the vector $v' = \left(\frac{\partial \hat{q}_1}{\partial z_k}, ..., \frac{\partial \hat{q}_n}{\partial z_k}\right)$. Then replacing any column (or any row) of M with v should not increase the rank. It is relatively simple to devise an empirical test for this; see Browning and Chiappori (1998) for details.

3.1.3 Some extensions

Finally, a similar investigation has been conducted for other, non-unitary models of household behavior. Lechene and Preston (2009) analyze the demand function stemming from a non cooperative model (involving private provision of the public goods) similar to that discussed in Chapter 4. They show that, again, a decomposition of the type (18) holds. However, the rank conditions on the 'deviation' matrix R are different; specifically, Lechène and Preston show that the rank of R can take any value between 1 and the number of public goods in the model. Recently, d'Aspremont and Dos Santos Fereira (2009) have introduced a general framework that provides a continuous link between the cooperative and the non cooperative solutions. In their setting, couples are characterized by a pair of parameters that indicate how 'cooperatively' each agent behaves. Again, they derive a (18) decomposition; however, the rank of matrix R can now take values between 1 and twice the number of public goods. On the empirical front, Del Bocca and Flinn ($$$) have proposed models in which agents may cooperate at some coordination cost; the decision to cooperate (or not) is then endogeneously derived from the model.

3.2 Identifying the collective model

In the presence of price variation, the identifiability problem can be stated in full generality; indeed, when price effects are observable it may be possible to recover individual preferences and demand functions (not only the Engel curves). Clearly, identifying assumptions are necessary; in its most general version (with general preferences $u^a\left(\mathbf{q}^a, \mathbf{q}^b, \mathbf{Q}\right)$ and $u^b\left(\mathbf{q}^a, \mathbf{q}^b, \mathbf{Q}\right)$), there exists a continuum of different structural models generating the same demand function. For instance, Chiappori and Ekeland (2006) show that any function satisfying SNR1 (see equation (18)) can be generated as the Pareto efficient demand of a household in which all consumption is public, and also of an (obviously different) household in which all consumption is private. Therefore, we assume in this subsection that preferences are egoistic ($u^a\left(\mathbf{q}^a, \mathbf{Q}\right)$ and $u^b\left(\mathbf{q}^b, \mathbf{Q}\right)$) , although our results have implications for caring preferences as well. We also assume that the econometrician knows which goods are private and which are public.

Even with egoistic preferences, however, the collective structure cannot in general be fully identified from demand data. To give a simple counterexample, assume for a moment that all goods are publicly consumed and consider two pairs of utility functions, $\left(u^a\left(\mathbf{Q}\right), u^b\left(\mathbf{Q}\right)\right)$ and $\left(\tilde{u}^a\left(\mathbf{Q}\right), \tilde{u}^b\left(\mathbf{Q}\right)\right)$ with

$$\begin{aligned} \tilde{u}^a &= F\left(u^a, u^b\right) \\ \tilde{u}^b &= G\left(u^a, u^b\right) \end{aligned}$$

for two arbitrary, increasing functions F and G. It is easy to check that any allocation that is Pareto efficient for $\left(\tilde{u}^a, \tilde{u}^b\right)$ must be Pareto efficient for $\left(u^a, u^b\right)$ as well; otherwise one could increase u^a and u^b without violating the budget constraint, but this would increase $\tilde{u}^a$ and $\tilde{u}^b$, a contradiction. It follows that any demand that can be rationalized by $\left(\tilde{u}^a, \tilde{u}^b\right)$ can also be rationalized by $\left(u^a, u^b\right)$ (of course, with different Pareto weights), so that the two structures are empirically indistinguishable. Since F and G are arbitrary, we are facing a large degree of indeterminacy.

A negative result of this type has a simple meaning: additional identifying hypotheses are required. If there are at least four commodities, then Chiappori and Ekeland (2009) prove the following results.

- If for each household member there is a commodity that this member does not consume and is consumed by at least one other member, then generically one can exactly recover the collective indirect utility function[6] of each member (up to an increasing transform). For any cardinalization of these utility functions, Pareto weights can be recovered. If there are only two persons in the household then this exclusion restriction is equivalent to an *exclusivity* condition that each member has one good that only they consume; with at least three members, exclusion is weaker than exclusivity.

- If all commodities are publicly consumed, identifying collective indirect utility functions is equivalent to identifying individual utilities. With private consumptions, on the contrary, any given pair of collective indirect utilities is compatible with a continuum of combinations of individual utilities and (conditional) sharing rules. However, all these combinations are welfare equivalent, in the sense that they generate the same welfare conclusions. For instance, if a given reform is found to increase the welfare of a while decreasing that of b under a specific combination of individual utilities, the same conclusion will hold for all combinations.

- Finally, if there is at least one distribution factor, the exclusivity restriction can be relaxed and identifiability obtains with one *assignable* good only.

In the literature the traditional choice for exclusive goods for husband and wife is men and women's clothing respectively. There is a subtle but important difference between the notion of exclusivity and that of assignability. In both cases, we observe consumptions at the individual level. But exclusive goods have different prices, whereas under assignability we observe individual consumptions of the same good - so there is only one price. Therefore, when considering clothing as two exclusive goods we have to assume they have different prices. In practice prices for men and women's clothing tend to be very colinear and we have to treat clothing as an assignable good.

[6] See section 2.2 of chapter 4 for the definition of the collective indirect utility function.

Two remarks are in order at that point. First, the identifiability result just presented is, by nature, non parametric, in the sense that it does not rely on the choice of a specific functional form for either preferences or Pareto weights.[7] Under an explicitly parametric approach, stronger identification results may obtain; for instance, it may be the case that one exclusive good only is sufficient to identify all the relevant parameters. Clearly, these additional properties are due to the specific functional form under consideration. Second, the result is generic, in the sense that it holds for 'almost all' structures. An interesting remark is that (non-generic) exceptions include the case in which Pareto weights are constant; in such a case, the collective indirect utilities are not identifiable in general.[8] To see why, simply note that, in that case, the household maximizes a collective utility of the form:

$$U\left(\mathbf{q}^a, \mathbf{q}^b, \mathbf{Q}\right) = \mu u^a\left(\mathbf{q}^a, \mathbf{Q}\right) + u^b\left(\mathbf{q}^b, \mathbf{Q}\right) \tag{19}$$

subject to the budget constraint. Since μ is a constant, standard results in consumer theory guarantee that we can recover U from observed (household) demand. However, for any given U there exists a continuum of u^a and u^b such that (19) is satisfied. For instance, take any such u^a and u^b that are strongly increasing and concave, pick up any smooth function ϕ, and define $\bar{u}^a$ and $\bar{u}^b$ by:

$$\begin{aligned}
\bar{u}^a\left(\mathbf{q}^a, \mathbf{Q}\right) &= u^a\left(\mathbf{q}^a, \mathbf{Q}\right) + \varepsilon\phi\left(\mathbf{Q}\right) \\
\bar{u}^b\left(\mathbf{q}^b, \mathbf{Q}\right) &= u^b\left(\mathbf{q}^b, \mathbf{Q}\right) - \mu\varepsilon\phi\left(\mathbf{Q}\right)
\end{aligned}$$

Then $\mu\bar{u}^a + \bar{u}^b = U$ and (19) is satisfied; moreover, on any compact set, $\bar{u}^a$ and $\bar{u}^b$ are concave and increasing for ε small enough.

Ironically, the case of a constant Pareto weight corresponds to the Samuelson justification of the unitary setting, in which a single, price-independent welfare index is maximized. From an identification viewpoint, adopting a unitary framework is thus a very inappropriate choice, since *it rules out the identification of individual welfares.*

[7] This notion of 'non parametric', which is used for instance by econometricians, should be carefully distinguished from the perspective based on revealed preferences - which, unfortunately, is also often called 'non parametric'. In a nutshell, the revealed preferences approach does not require the observability of a demand *function*, but only of a finite number of points; it then describes relationship that must be satisfied for the points to be compatible with the model under consideration. This view will be described in subsection 3.4.

[8] This case is 'non generic' in the sense that in the set of continuous functions, constant functions are non-generic.

Our general conclusion is that welfare relevant structure is indeed identifiable in general, provided that one can observe one exclusive consumption per member (or one overall with a distribution factor). However, identifiability fails to obtain in a context in which the household behaves as a single decision maker.

3.3 A simple example

The previous results can be illustrated by the following example, directly borrowed from Chiappori and Ekeland (2009). 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:

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

3.3.1 Household demand

The group solves the program:

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

under the budget constraint:

$$\mathbf{p}'\left(\mathbf{q}^a + \mathbf{q}^b\right) + \mathbf{P}'\mathbf{Q} = x$$

where one price has been normalized to 1. Individual demands for private goods are given by:

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

generating the aggregate demand:

$$p_i q_i = p_i c_i + \left[\alpha_i^a \left(\mu^0 + \mu^x x + \mu^z z\right) + \alpha_i^b \left(1 - \left(\mu^0 + \mu^x x + \mu^z z\right)\right)\right] Y \quad (20)$$

$$P_j Q_j = P_j C_j + \left[\alpha_j^a \left(\mu^0 + \mu^x x + \mu^z z\right) + \alpha_j^b \left(1 - \left(\mu^0 + \mu^x x + \mu^z z\right)\right)\right] Y \quad (21)$$

where $c_i = c_i^a + c_i^b$ and $Y = \left(x - \sum_{i,s} p_i c_i^s - \sum_j P_j C_j\right)$. The household demand is thus a direct generalization of the standard LES, with additional quadratic terms in x^2 and cross terms in xp_i and xP_j, plus terms involving the distribution factor z; one can readily check that it does not satisfy Slutsky symmetry in general, although it does satisfy SNR1.

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 discussed above, this indeterminacy is however welfare irrelevant, because the collective indirect utilities of the wife and the husband are, up to an additive constant:

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

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

$$\begin{aligned}
&\xi_i = q_i \quad \text{for } i \leq n,\ \xi_i = Q_i \quad \text{for } n < i \leq N \\
&\quad \text{and similarly} \\
&\gamma_i = c_i \quad \text{for } i \leq n,\ \gamma_i = C_i \quad \text{for } n < i \leq N \\
&\pi_i = p_i \quad \text{for } i \leq n, \quad \pi_i = P_i \quad \text{for } n < i \leq N
\end{aligned}$$

so that the group demand has the simple form:

$$\pi_i \xi_i = \pi_i \gamma_i + \left[\alpha_i^a \left(\mu^0 + \mu^x x + \mu^z z\right) + \alpha_i^b \left(1 - \left(\mu^0 + \mu^x x + \mu^z z\right)\right)\right] Y \quad (22)$$

leading to collective indirect utilities of the form:

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

It is clear, from 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.

3.3.2 Identifiability

The general case The question, now, is whether the empirical estimation of the form (22) allows us to recover the parameters of interest - namely, the α_i^s, the γ^i, and the μ^α. We start by rewriting (22) as:

$$\pi_i \xi_i = \pi_i \gamma_i + \left(\alpha_i^b + \left(\alpha_i^a - \alpha_i^b\right) \mu^0 + \left(\alpha_i^a - \alpha_i^b\right) \left(\mu^x x + \mu^z z\right)\right) \left(x - \sum_m \pi_m \gamma^m\right) \quad (23)$$

The right hand side of (23) can in principle be econometrically identified; we can thus recover the coefficients of the variables, namely x, x^2, xz, the π_m and the products $x\pi_m$ and $z\pi_m$. For any i and any $m \neq i$, the ratio of the coefficient of x 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 (23) can only recover for each j the respective coefficients of x, x^2 and xz, that is, the three expressions

$$\begin{aligned} K_x^j &= \alpha_j^b + \left(\alpha_j^a - \alpha_j^b\right)\mu^0 \\ K_{xx}^j &= \left(\alpha_j^a - \alpha_j^b\right)\mu^x \\ K_{xz}^j &= \left(\alpha_j^a - \alpha_j^b\right)\mu^z \end{aligned} \tag{24}$$

Now, pick up two arbitrary values for μ^0 and μ^x, with $\mu^x \neq 0$. 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: $\mu^x = \mu^z = 0$. Then $K_{xz}^j = 0$ for all j (since distribution factors cannot matter[9]), and $K_{xx}^j = 0$ for all j (demand must be linear in x, since a quadratic term would violate Slutsky). We are left with $K_x^j = \alpha_j^b + \left(\alpha_j^a - \alpha_j^b\right)\mu^0$, 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 coefficients. Assume that member a does not consume commodity 1 and b does not consume good 2; that is, $\alpha_1^a = \alpha_2^b = 0$. Then equations (24) gives:

$$\alpha_1^b\left(1-\mu^0\right) = K_x^1,\ -\alpha_1^b\mu^x = K_{xx}^1,\ -\alpha_1^b\mu^z = K_{xz}^1$$

and:

$$\alpha_2^a\mu^0 = K_x^2,\ \alpha_2^a\mu^x = K_{xx}^2,\ \alpha_2^a\mu^z = K_{xz}^2$$

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

$$\frac{1-\mu^0}{\mu^x} = -\frac{K_x^1}{K_{xx}^1} \quad \text{and} \quad \frac{\mu^0}{\mu^x} = \frac{K_x^2}{K_{xx}^2}$$

[9] For a discussion of the role of distribution factor in a unitary context, see Browning, Chiappori and Lechene (2006).

therefore, assuming $K_x^2 K_{xx}^1 - K_x^1 K_{xx}^2 \neq 0$:

$$\frac{1-\mu^0}{\mu^0} = -\frac{K_x^1 K_{xx}^2}{K_x^2 K_{xx}^1} \text{ and } \mu^0 = \frac{K_x^2 K_{xx}^1}{K_x^2 K_{xx}^1 - K_x^1 K_{xx}^2}$$

It follows that

$$\mu^x = \frac{K_{xx}^2}{K_x^2} \text{ and } \mu^0 = \frac{K_{xx}^2 K_{xx}^1}{K_x^2 K_{xx}^1 - K_x^1 K_{xx}^2}$$

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 $K_x^2 K_{xx}^1 - K_x^1 K_{xx}^2 \neq 0$. Clearly, the set of parameters values violating this condition is of zero measure.

Finally, it is important to note that this conclusion requires $\mu^x \neq 0$; in particular, *it does not hold true* in the unitary version, in which $\mu^x = \mu^z = 0$. Indeed, the same exclusion restrictions as above only allow us to recover $\alpha_1^b \left(1 - \mu^0\right) = K_x^1$ and $\alpha_2^a \mu^0 = K_x^2$; 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.

3.4 The revealed preference approach

Up until now we have considered analysis that posits that we can estimate smooth demands and test for the generalized Slutsky conditions for integrability. An alternative approach to empirical demand analysis that has gained ground in the last few years is the *revealed preference* (RP) approach that derives from Afriat (1967) and Varian (1982). This style of analysis explicitly recognizes that we only ever have a finite set of observations on prices and quantities which cannot be used to directly construct smooth demand functions without auxiliary assumptions. The revealed preference approach instead identifies linear inequality conditions on the finite data set that characterize rational behavior. The most attractive feature of the Afriat-Varian approach is that no functional form assumptions are imposed. Moreover powerful numerical methods are available to implement the RP tests. The drawback of the RP approach is that even when the data satisfy the RP conditions, we can only set identify preferences; see Blundell *et al* (2008).

Generalizing the unitary model RP conditions to the collective setting was first achieved in Chiappori (1988) for a specific version of the collective model. The conditions for the general model have been established in Cherchye, De Rock and Vermeulen (2007), (2009a) and Cherchye, De Rock, Sabbe, and Vermeulen (2008); these papers provide a complete characterization of the collective model in a revealed preference context. This requires several significant extensions to the RP approach for the unitary model. Amongst these are these authors allow for non-convex preferences and develop novel (integer programming) methods since the linear programming techniques that work for the unitary model are not applicable for the collective model. The tests for 'collective rationality' require finding individual utility levels, individual marginal utilities of money (implying Pareto weights) and individual assignments for private goods and Lindahl prices for public goods. As in the unitary model, these methods can only set identify the preferences of the household members and the Pareto weight. Cherchye, De Rock and Vermeulen (2009b) apply these methods to a Russian expenditure panel.

4 The case of labor supply

4.1 Egoistic preferences and private consumption

A large part of the empirical literature on household behavior is devoted to labor supply. The theory has been presented in Section 4 of Chapter 4; here we concentrate on the empirical implications. Most empirical works consider the simple setting with egoistic preferences and private consumption; see subsection 4.2 of Chapter 4. In this framework, results have been established by Chiappori (1988, 1992) and Chiappori, Fortin and Lacroix (2002). Regarding testability, strong implications can be derived, even in this simple setting. Even more remarkable is the fact that the observation of individual labor supplies, as functions of wages, non labor income and distribution factors, allows us to identify the sharing rule up to an additive constant. We start from the two leisure demand equations:

$$l^a\left(w^a, w^b, y, z\right) = \tilde{l}^a\left(w^a, \rho\left(w^a, w^b, y, z\right)\right) \tag{25}$$

$$l^b\left(w^a, w^b, y, z\right) = \tilde{l}^b\left(w^b, y - \rho\left(w^a, w^b, y, z\right)\right) \tag{26}$$

where $\tilde{l}^a$ denotes the Marshallian demand for leisure by person a, y is full income and $\rho(.)$ is a's share of full income (see equations (42) and (43) of chapter 4). We assume that both partners shares are increasing in full income, $0 < \partial\rho/\partial y < 1$, and that the distribution factor is 'meaningful', $\partial\rho/\partial z \neq 0$.

Taking derivatives through (25):

$$\begin{aligned}
\frac{\partial l^a}{\partial w^b} &= \frac{\partial \tilde{l}^a}{\partial \rho}\frac{\partial \rho}{\partial w^b} \\
\frac{\partial l^a}{\partial y} &= \frac{\partial \tilde{l}^a}{\partial \rho}\frac{\partial \rho}{\partial y} \\
\frac{\partial l^a}{\partial z} &= \frac{\partial \tilde{l}^a}{\partial \rho}\frac{\partial \rho}{\partial z}
\end{aligned} \tag{27}$$

so that:

$$\frac{\partial l^a/\partial z}{\partial l^a/\partial y} = \frac{\partial \rho/\partial z}{\partial \rho/\partial y} \tag{28}$$

Similarly for b:

$$\begin{aligned}
\frac{\partial l^b}{\partial w^a} &= -\frac{\partial \tilde{l}^b}{\partial y^b}\frac{\partial \rho}{\partial w^a} \\
\frac{\partial l^b}{\partial y} &= \frac{\partial \tilde{l}^b}{\partial y^b}\left(1-\frac{\partial \rho}{\partial y}\right) \\
\frac{\partial l^b}{\partial z} &= -\frac{\partial \tilde{l}^b}{\partial y^b}\frac{\partial \rho}{\partial z}
\end{aligned} \tag{29}$$

so that:

$$\frac{\partial l^b/\partial z}{\partial l^b/\partial y} = -\frac{\partial \rho/\partial z}{1-\partial \rho/\partial y} \tag{30}$$

For notational simplicity, let F^s denote the fraction $\frac{\partial l^s/\partial z}{\partial l^s/\partial y}$ for $s = a, b$; note that F^s can in principle be observed (or estimated) as a function of $\left(w^a, w^b, y, z\right)$, and that $F^a = F^b$ would imply

$$\frac{\partial \rho/\partial z}{\partial \rho/\partial y} = -\frac{\partial \rho/\partial z}{1-\partial \rho/\partial y}$$

which is impossible if $\partial \rho/\partial z \neq 0$.

Now, (28) and (30) can be solved in $\partial \rho/\partial z$ and $\partial \rho/\partial y$ (since $F^b \neq F^a$):

$$\begin{aligned}
\frac{\partial \rho}{\partial y} &= \frac{F^b}{F^b - F^a} \\
\frac{\partial \rho}{\partial z} &= \frac{F^a F^b}{F^b - F^a}
\end{aligned}$$

We thus conclude that the partials of ρ with respect to income and distribution factor are identifiable.

Finally, the first two equations of (27) and of (29) give respectively:

$$\begin{aligned}
\frac{\partial \rho}{\partial w^b} &= \frac{\partial l^a/\partial w^b}{\partial l^a/\partial y}\frac{\partial \rho}{\partial y}=\frac{\partial l^a/\partial w^b}{\partial l^a/\partial y}\frac{F^b}{F^b-F^a} \text{ and} \\
\frac{\partial \rho}{\partial w^a} &= \frac{\partial l^b/\partial w^a}{\partial l^b/\partial y}\left(1-\frac{\partial \rho}{\partial y}\right)=-\frac{\partial l^a/\partial w^b}{\partial l^a/\partial y}\frac{F^a}{F^b-F^a} \qquad (31)
\end{aligned}$$

The conclusion is thus that *all partial derivatives of the sharing rule can be exactly recovered* from the observation of the two labor supply functions. From the sole observation of labor supplies, one can recover the impact of wages, non labor income and distribution factors on the sharing rule. Finally, the cross derivative restrictions generate additional testable predictions.

The sharing rule itself is identified up to an additive constant; that constant cannot be identified unless either all commodities are assignable or individual preferences are known (for instance, from data on singles). To see why, take labor supply functions l^a and l^b that satisfy (25) and (26) for some sharing rule ρ and some Marshallian demands $\tilde{l}^s$ derived from individual utilities $u^s, s=a,b$. Now, for some constant K, define ρ_K, u_K^a and u_K^b by:

$$\begin{aligned}
\rho_K\left(w^a,w^b,y,z\right) &= \rho\left(w^a,w^b,y,z\right)+K \\
u_K^a\left(l^a,C^a\right) &= u^a\left(l^a,C^a-K\right) \\
u_K^b\left(l^b,C^b\right) &= u^b\left(l^b,C^b+K\right)
\end{aligned}$$

It is easy to check that the Marshallian demands derived from ρ_K, u_K^a and u_K^b satisfy (25) and (26). The intuition is illustrated in Figure 1 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, therefore the intercept of the budget constraint, is shifted downward by K; second, all indifference curves are also shifted downward by the same amount. When only labor supply (on the horizontal axis) is observable, these models are empirically indistinguishable.

Note, however, that the models are also welfare equivalent (that is, the constant is 'irrelevant'), in the sense defined in section 3.3 of chapter 4: changing the constant affects neither the comparative statics nor the welfare analysis derived from the model. Technically, the collective indirect utility of each member is the same in both models; one can readily check that the two models generate the same *level* of utility for each spouse. In the end, the optimal identification strategy depends on the question under consideration. If one want to formulate welfare judgments, collective indirect

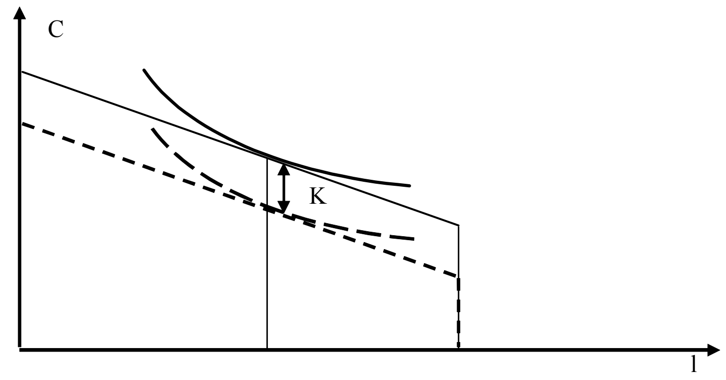

Figure 1:

utilities are sufficient, and they can be recovered without additional assumptions. If, on the other hand, the focus is on intrahousehold inequality, the basic model can identify the *changes* affecting intrahousehold inequality, but not its initial *level*; therefore additional assumptions may be needed. For instance, some empirical works assume that preferences are unchanged by marriage, therefore can be identified from the labor supply of singles; then the constant can also be recovered.

Finally, one should not conclude from the previous derivation that the presence of a distribution factor is needed for identifiability. This is actually not the case. The observation of individual labor supplies, as functions of wages and non labor income, are 'generically' sufficient to recover the sharing rule up to an additive constant (Chiappori 1988, 1992). However, identification is only generic in that case; moreover, it is arguably less robust, since it involves second derivatives of the labor supply functions.

4.2 Extensions

The model has been extended in various directions. First, while the assumption of a unique, Hicksian composite consumption good is standard in the labor supply literature, the model can address a more general framework. Chiappori (2008) consider a model with two leisures and many consumption goods that are privately (but not exclusively) consumed by the members. The context is cross-sectional, in the sense that there is variation in wages but not in prices. He shows that if one distribution factor (at least) is available, then it is possible to identify (again up to additive constants) not only the sharing rule but also the individual demands for all private commodities, as functions of wages and non-labor income. Chiappori concludes that in a collective model of consumption and labor supply estimated on cross sectional data, it is possible to recover the income and wage elasticities of individual demands for each good.

Secondly, the computations above rely on the assumption that labor supply is a continuous variable. This may fail to hold for two reasons. First, in some households one member may elect not to participate; in that case, the person's labor supply is at a corner solution equal to zero. Secondly, the structure of labor markets may put constraints on the number of hours supplied by individuals. For instance, the choice may be only between working part time, full time or not at all; then labor supply should rather be modelled as a discrete variable. Extensions of the previous model to such situations have been studied by Blundell *et al* (2007) and Donni (2007).

Although very convenient, this framework has its limitations. The privateness assumption has been criticized on two grounds. First, while some consumptions are indeed private, others are not. Children expenditures are a typical example of public goods within the household. Blundell, Chiappori and Meghir (2005) analyze a model similar to the previous one but for the consumption good, which is taken to be public. They show that, again, the model is identifiable from the observation of labor supply behavior. They show how their approach can be extended to household production under various specifications.

A second criticism concerns the private nature of individual leisure. It could indeed be argued that leisure is, to some extent, publicly consumed; after all, the utility I derive from my own free time may be higher when my spouse is available as well. The general insight, here, is that a model in which both members' leisure enter each individual utility is still identifiable, provided that some other commodities are exclusive (this is a consequence of the general identifiability results described in Section 2). Fong and Zhang (2001) analyze a framework in which leisure is partly private and partly public; they show that one assignable good is sufficient for identification in the presence of a distribution factor.

Finally, a standard problem with traditional models of labor supply is the implicit assumption that time is divided between market work and leisure - so that any moment not spend working of a wage tends to be assimilated with leisure. This, of course, disregards domestic production, and may result in misleading evaluations. For instance, if a given reform is found to reduce female market labor supply, we may conclude that it increases her leisure, hence her utility, whereas the actual outcome is more domestic work (and ultimately less leisure) for the wife. Donni (2008) shows, however, that the direction of the mistake depends on the properties of the domestic production function. To take an extreme example, consider the case in which the latter is additively separable; that is, when t^s denotes the time spent on domestic production by agent s, then the outcome is:

$$C = f^a(t^a) + f^b\left(t^b\right)$$

Assuming that the domestic good is marketable with price p, the first order conditions require that:

$$\frac{df^s(t^s)}{dt^s} = \frac{w^s}{p}$$

which implies that the time spent on domestic production by s only depends on their wage (and on the price of the domestic good). It follows that any welfare judgment that ignores domestic production is in fact unbiased - that is, a reform that is found to increase the wife's welfare when ignoring domestic production has the same impact even when domestic production is taken into account and conversely. This conclusion, however, does not hold when the productivity of the wife's domestic work depends on the husband's.

5 Empirical evidence.

5.1 Evidence against the unitary model.

As we have seen, there are two strands to testing for the unitary model: the Slutsky conditions and independence of behavior from distribution factors. Regarding the former, Slutsky symmetry is often rejected on household expenditure survey data. Rejections of Slutsky symmetry may be due to many factors other than a failure of the unitary assumption. For example, we might have the wrong functional form or an inappropriate grouping of goods or be wrongly assuming separability from housing and durables or accounting for latent heterogeneity inappropriately and so on. A widely cited piece of evidence that the unitary assumption itself is problematic is from Browning and Chiappori (1998) who model commodity demands using Canadian data. Using a QAIDS formulation, they test for symmetry for three sub-samples: single women, single men and couples with no children. They find that Slutsky symmetry is *not* rejected for single women or single men, while it is (very strongly) for couples. Since most of underlying modelling assumptions are the same across the three strata, this suggests that it is the unitary assumption that is the problem. These findings have been replicated by Kapan (2009) using Turkish data; and Vermeulen (2005) obtains similar results for labor supply.

Although suggestive, the rejection of Slutsky symmetry would not, by itself, warrant abandoning the unitary assumption. Much more convincing are the next set of tests we discuss. The second principle implication of the unitary model is that possible distribution factors do not have any significant impact on the household choice variable being considered. Unlike the test for the Slutsky conditions, such tests can be conducted whether or not we have price variation. Table 1 gives a partial listing of distribution factors that have been considered in the literature. Below we discuss the validity of these factors. The most widely used distribution factor for this is some measure of relative incomes, earnings or wages. Such tests are often called tests of 'income pooling': only household income matters for choice outcomes and not the source of the income.[10] As we have seen, Becker explicitly introduced the RKT to justify income pooling. Tests for the exclusion of other distribution factors constitute a generalization of income pooling.

	Distribution factor
1	Relative income
2	Relative wages
3	Relative unearned income
4	Relative age
5	Relative education
6	Local sex ratio
7	Household income
8	Background family factors
9	Control of land
10	Previous children
11	Reported influence within household
12	Married or cohabiting
13	Divorce laws
14	Alimonies
15	Single parent benefits
16	Gender of a benefit's recipient

Table 1: Distribution factors

[10] Income pooling is a necessary condition for a unitary model but not a sufficient condition. In particular, income pooling can hold locally if we have a noncooperative voluntary contributions game; see section 4 of chapter 3.

Bruce (1989) provides a listing of the research on low income countries documenting tensions within households about the use of household resources. Strauss *et al* (2000) present an exhaustive list of tests for income pooling for low income countries up to their publication date. Table 2 lists some of the studies that have considered non-unitary models.[11] As can be seen from this Table, a wide variety of outcomes and distribution factors have been considered for many different countries. The most widely used distribution factor is relative income (the 'income pooling' test). All of the cited papers find a significant role for the distribution factors that should not affect the outcomes in a unitary model. For instance, an early and influential paper by Thomas (1990), based on Brazilian cross-sectional data, finds that the relative share of non labor income coming from the wife has a very significant impact on the health status of children within the household.

This unanimity may be somewhat misleading; our impression is that there is a strong publication bias against not finding an effect. That is, editors may not be interested in papers that confirm a conventional view by finding an insignificant effect. Nonetheless, the evidence seems overwhelming: a principal implication of the unitary model is rejected on a wide set of data sets for a wide range of outcomes.

Many of these rejections may have other explanations than a failure of the unitary assumption. For example, consider a unitary demand model in which the relative (labor or non labor) earnings of the two partners do *not* affect demand behavior directly. Suppose, however, that there is unobserved heterogeneity in tastes between husbands and wives and this heterogeneity is correlated with heterogeneity in earnings. For example, suppose the relative preference for clothing between a husband and wife is correlated with their relative tastes for work. Then we would find that the demand for clothing (conditional on prices, total expenditure and preference factors) will be correlated with relative earnings, with higher earners having relatively more clothing expenditure than their partner. In this case, a finding that relative clothing demands are partially correlated with relative earnings is spurious in the sense that it is due to inadequate control for heterogeneity rather than a failure of the unitary assumption. Attempts to find instruments to wash out this spurious correlation have not been notably successful: it has proven impossible to find observables that are correlated with, say, relative earnings but not with demand heterogeneity.[12] Similarly, Thomas's findings

[11] This listing is by no means exhaustive and tends to focus on results from high income countries.

[12] Luo (2002) estimates a demand system explicitly allowing for uncorrelated heterogeneity and finds that the BC results for Slutsky symmetry hold up

might simply reflect the fact that some women are more willing to invest over the long term than others; such women would be likely to spend more on children, and also to have saved more in the past, hence to receive more non labor income today. Such a mechanism does not rely on a shift in power triggered by the wife's larger relative contribution to total income, but only on unobserved heterogeneity between women; as such, it is fully compatible with a unitary representation.

Several papers provide strong evidence concerning income pooling that can hardly be attributed to heterogeneity biases. Lundberg *et al* (1997) present quasi-experimental evidence based on a reform of the UK child public support system in April 1977. Prior to that time families with children received a child tax allowance and a taxable child allowance. This effectively meant that the child benefits were paid to the higher earner, mostly the father. After April 1977, the old scheme was dropped in favor of a non-taxable child benefit which is paid directly to the mother. This re-allocation of income within the household can reasonably be treated as exogenous to the affected households. Moreover, the child benefit was a sizable transfer (equal to 8% of male earnings for a two child household). Thus we have a large, exogenous 'treatment' which can be used to assess the importance of the distribution of income within the household. The major confounding factor is that the reform was not revenue neutral for all households with children and some saw a substantial rise in net household income. Lundberg *et al* use UK Family Expenditure Survey cross-section data from before and after the change to gauge the effect of the reform on assignable expenditures. They focus attention on the ratio of expenditures on children's clothing and women's clothing, both relative to men's clothing. Their findings are unequivocal: both ratios rose significantly after the reform.[13]

Another strong rejection is provided by Duflo (2003), who analyses a reform of the South African social pension program for elderly that extended the benefits to a large black population who were previously not covered. Due to the eligibility criteria, the coverage is not universal; in some households, in particular one only of the grand parents receives the benefit. Duflo uses a difference in difference approach based on the demographics of the siblings to control for selection in eligibility. She shows that the recipient's gender - a typical distribution factor - is of considerable importance for the

[13] A re-analysis of the Lundberg *et al* episode by Hotchkiss (2005) suggests that it may not be valid. The point at issue is that women in childless couples also appeared to increase their clothing expenditure in the same period. Ward-Batts (2008) convincingly contests this finding: the Hotchkiss timing is not consistent and Ward-Batts uses micro data rather than the grouped data of Lundberg *et al* and Hotchkiss.

impact of the transfers on children's health: a payment to the grandfather has no significant effect, whereas the same amount paid to the grandmother results in a huge improvement in the health status of girls in the family. These contributions and several others (including a subsequent analysis on micro data for all goods by Ward-Batts (2008)) very convincingly suggest that income pooling is indeed strongly rejected on real data.

5.2 Evidence on the collective model.

Although the evidence against the unitary model in specific contexts is not as robust as widely believed, it does add up and most researchers in the field now seem to agree that any reasonable model should account for spouses having different preferences and for the intrahousehold distribution of 'power' to matter for behavior. Evidence against the unitary model does not, however, necessarily constitute evidence for the collective model. Unfortunately it has turned out to be difficult to devise powerful tests for the collective model. This is because such tests must rely either on a test of the quasi-Slutsky condition or the proportionality restriction on distribution factors or on a combination of these conditions; see sections 1.3 and 1.4 of chapter 4. As regards the $SNR1$ restriction (see equation (18)), we need price variation and at least five goods to reject symmetry. This largely restricts our ability to test for $SRN1$ in the labor supply context, although tests based on more specific assumptions - for example, exclusivity of leisure - or on different approaches - typically revealed preferences - are indeed feasible. Tests based on proportionality are in general easier to implement, but they still require at least two unequivocal distribution factors.

Among the few attempts to take $SRN1$ to the data are Browning and Chiappori (1998), Dauphin *et al* (2009) and Kapan (2009). These works share common features: they all estimate a demand system, using a well known and flexible functional form (QUAIDS) that nests both the unitary and the collective settings as specific cases (the former being itself nested within the latter). While the data sets are different (a specific feature of the Turkish data considered by Kapan is the presence of important and largely exogenous variations in relative prices, due to high inflation over the period), they reach similar conclusions. For instance, when testing symmetry and SNR1 on three subsamples - single males, single females, and couples, they all fail to reject the unitary version on singles; but on couples, they strongly reject the unitary version, but not SNR1. In addition, the contributions provide interesting insights on various specific aspects of intrahousehold decision processes. Both Browning and Chiappori and Dauphin *et al* provide additional tests using distribution factors, which tend to support the collective model. Kapan finds that while most Turkish families do not behave as if there was a single decision maker, a notable exception is provided by traditional, rural households, for whom the unitary version is not rejected. Finally, both Kapan and Dauphin *et al* find that older children (above age 16) do play a role in the decision process.

The validity of proportionality tests, on the other hand, depends crucially on an *a priori* division of demographic and environmental factors between preference factors and distribution factors (additionally, a variable can be both). Typical candidate preference factors include household composition, the age of one of the spouses, the ownership of a car or a house, region of residence etc.. Typical distribution factors are listed in Table 1. A general concern is that the household specific variables could be correlated with constraints or preferences which would invalidate them as distribution factors; societal variables are less susceptible to this problem. Fortunately, as we have shown above (see subsection 2.1) we only need one unequivocal distribution factor to credibly test for proportionality for other candidate distribution factors. To illustrate, suppose we construct an index quantifying the extent to which laws governing divorce favor women, and we take that index as a unequivocal distribution factor. If the index is 'significant' in the choice equations, we can then test for proportionality for other candidate distribution factors. In theory, we could simply take all of the factors that satisfy the proportionality tests as distribution factors and assign other 'significant' variables as preference factors. In practice, this may not be appealing if the factor that fails the proportionality test is unlikely to be a preference factor. For example, if the situation on the marriage market (as measured for instance by the local sex ratio) impacts on demand behavior but fails the proportionality test, we would be very reluctant to designate it a preference factor. Rather, this would cast doubt on our original choice of an unequivocal distribution factor (or the collective model itself!).

There is no evidence against the collective model in the papers listed in Table 1. There is, however, alternative evidence against the efficiency assumption of a different sort. The most convincing evidence of inefficient outcomes is Udry (1996). This is a different style of test than SNR1 and the relevance of distribution factors. Udry uses information on household production.

To sum up: there is considerable evidence against the unitary model and some evidence in favor of the collective model.[14] What is singularly

[14] A notable exception to the latter are the results for efficient risk sharing in low income

lacking in the literature are tests for the collective model against other non-unitary models for high income countries. This is part reflects the lack of non-collective models that can be taken to the data.

5.3 Estimating the collective model

Many of the works mentioned above go beyond testing the collective model; insofar as the predictions are not rejected, they often propose an estimation of the structural components of the model. Although this field is still largely in construction, we may briefly summarize some findings obtained so far.

5.3.1 Demand studies

Many of the papers listed in Table 2 use demand data alone to test for the collective model. Only three of them go beyond testing and impose the collective model restrictions and then estimate the sharing rule and how it depends on distribution factors. The first paper to do this was Browning *et al* (1994). These authors use Canadian Family Expenditure Survey data on men and women's clothing to test for the collective model restrictions and to identify the determinants of the sharing rule. Although they have price data they absorb prices into year/region dummies and treat the data as cross-sectional. Thus the 'no price variation' analysis of section 2 is appropriate. They only consider singles and married couples who are in fulltime employment. The distribution factors they find significant are the difference in ages and the relative earnings of the two partners; they also allow that total expenditure on nondurables and services enters the sharing rule. They address directly the problem that variations in relative earnings may be spuriously correlated with spending on clothing (higher paid jobs might require relatively more expensive clothing) by testing whether singles have clothing demands that depend on earnings. They find that for single men and single women, earnings do not impact on clothing demand once we take account of total expenditure. It is important to note that this does not imply that clothing demand is separable from labor supply (it is not) since they condition on both partners being in fulltime work and effectively test for whether *wages* affect preferences. Given the finding for singles, relative earnings are a reasonable candidate for being a distribution factor for

countries; see, for example, Dercon and Krishnan (2000), Dubois and Ligon (2005), Duflo and Udry (2003), Goldstein (2002), Ligon (2002). These tests, however, are based on specific models that crucially involve specific asusmptions regarding commitment; their discussion is therefore postponed until chapter 6 which deals with dynamic issues.

couples. As discussed in section 2 we cannot generally identify the location of the sharing rule, so Browning *et al* simply set it equal to one half (at the median of total expenditure) if the two partners have the same age and earnings. They find that differences in earnings have a highly significant but small impact on sharing: going from the wife having 25% of total earnings to 75% of total earnings shifts the sharing rule by 2.3 percentage points. Differences in age are similar with significant but small effects: going from being 10 years younger than her husband to being 10 years older raises the wife's share by two percentage points. Conversely, total expenditure (taken as a proxy for lifetime wealth) is less statistically significant but with a large effect: a 60% increase in total expenditure increases the wife's share by 12%. This suggests that wives in high wealth households have a higher share of nondurable expenditure.

Browning and Bonke (2009) use a supplement to the Danish Household Expenditure Surveys for 1999 to 2005. This supplement (designed by the authors) takes the form of respondents recording for every expenditure in a conventional expenditure diary for whom the item was bought: 'mainly for the household', 'for the husband', 'for the wife', 'for the children' and 'outside the household'. This is the first time that such information has been collected in a representative survey in a high income country. Another notable feature of these data is that they contain a richer set of potential distribution factors than most expenditure data sets. For example, questions were asked on the length of the current partnership; the labor force participation of the mothers of the husband and wife when they were 14 and the marital and fertility histories of the two partners. Since all expenditures are allocated in these data, a sharing rule can be constructed for each household. This allows for the identification of the location of the sharing rule as well as its dependence on distribution factors. These authors find that the mean of the sharing rule is very close to one half (at the mean of the data).[15] This equality of the mean total expenditures for the two partners masks that the sharing rule in different households varies widely. For example the first and third quantiles for the wife's share are 0.31 and 0.68 so that close to half of households have one partner receiving twice as much as the other. Some of this variation can be attributed to observable differences in distribution factors but most of it is 'latent' heterogeneity.

Some of the significant distribution factors in Browning and Bonke (2009)

[15] This equality of total assigned expenditures is not reflected in the expenditures on individual goods. For example, the individual allocations show that, in mean, wives spend more on clothing but less on alcohol and tobacco than their husbands.

are familiar from earlier studies; for example, if the wife has a higher share of gross income then she has a higher share of total expenditure. On the other hand, these authors do not find a significant role for the difference in age nor for total expenditure. Of more interest (because they have never been used in this context before) are the family and individual background variables. The two highly significant variables here are whether the husband's mother was in full-time employment when he was 14 and whether the partners have children from before the partnership. A husband having grown up in a household in which his mother was in full-time employment increases his share of expenditure. This is consistent with the theory model in which such men make desirable husbands (perhaps because they contribute more in housework) and hence do better in any match than an otherwise similar male who does not have this background. The other finding is less easy to rationalize. If either the husband or the wife has a previous child then the wife's share is lower. Thus a women who has had a previous child and is married to a man who has also had a previous child receives a share of total expenditure that is about nine percentage points lower than an otherwise comparable women in which neither partner has children from before the marriage. This is a very large effect which defies easy rationalization.

Browning, Chiappori and Lewbel (2009) also present identification results and estimates of the location of the sharing rule. These are based on making the strong assumptions that the preferences of singles and married people are the same and that only the household technology changes at marriage. This allows them to identify the location of the sharing rule as well as its dependence on distribution factors. Differences between the demands of singles and couples are picked up by a Barten style technology (see section 2 of chapter 2). For example, 'transport' is largely a public good whereas 'food at home' is largely private. The data used is the same as in Browning *et al* (1994) with the important difference that explicit account is taken of price variations across time and over regions. The distribution factors are very similar to those used in Browning *et al* (1994): the wife's share in total gross income, the difference in age between husband and wife, a home-ownership dummy and household total expenditure. The point estimate for the sharing rule (at the mean of the distribution factors) is 0.65; this is much higher than found in any other study. Mechanically it arises since the budget shares of couples are more similar to those of single women than to the budget shares of single men; this suggests that some relaxing of the unchanging preferences assumption is called for in future work. Having the allocations of total expenditure to each partner allows us to calculate budget shares for husbands and wives; see Table 3. Wives have higher budget shares for clothing, personal services and recreation whereas husbands have higher budget shares for food inside and outside the home, alcohol and tobacco and transport. Where comparisons can be made, this is similar to the Danish data discussed in the previous paragraph.

The results presented here on the location and determinants of the sharing rule do not sit together comfortably. This partly reflects the fact that potential distribution factors differ widely across different data sets and the excluded distribution factors are correlated with the included ones. For example, only one study can take account of the impact of previous children but this is correlated with the difference in age between the partners. More fundamentally, there is no coherent theory of the sharing rule. Without such a theory a 'kitchen sink' approach is adopted in which whatever variables are available in a particular data set are included as distribution factors (if they are not obviously preference or constraint factors) with limited explicit concern for biases due to endogeneity (a particular worry for income shares), omitted distribution factors or correlated latent heterogeneity. Equally worrying is the widespread assumption that private assignable goods are separable from public goods (see Donni (2009)). It is clear that much remains to be done and that 'much' probably requires better data than we have had available until now.

5.3.2 Labor supply

The first empirical estimations of a collective model of labor supply are due to Fortin and Lacroix (1997) and Chiappori, Fortin and Lacroix (2002). Using data from the 1988 PSID, the latter analyze the total number of hours worked each year by single males, single females and couples, concentrating exclusively on couples without children in which both spouses work. They consider two distribution factors, namely the state of the market for marriage, as summarized by the sex ratio computed by age and race at the state level, and the legislation governing divorce, summarized by an aggregate index with the convention that a larger value indicates laws that are more favorable to women. Their main findings can be summarized as follows:

- The distribution factors have a significant impact on both labor supplies. The signs are as predicted by the theory; that is, a higher sex ratio (denoting a smaller percentage of women on the marriage market), as well as divorce laws more favorable to women, reduce the wife's labor supply and increase the husband's, suggesting a transfer of resources to the wife. Interestingly, these effects are not present for singles; divorce laws do not impact singles' labor supplies in a significant way, whereas the sex ratio has no effect on the labor supply of single males and *increases* the labor supply of single women. Finally, the authors do not reject the prediction from the collective model that the impacts of the two factors on the two labor supplies should proportional.

- The corresponding transfers can be evaluated, since the sharing rule is identified up to an additive constant. A one percentage point increase in the sex ratio (representing roughly one standard deviation from the mean) is found to result in an annual transfer to the wife of more than \$2,000, or about 5% of the average household income. Likewise, a one point increase in the Divorce Laws Index (which varies from 1 to 4, with a mean at 2.8) induces husbands to transfer and additional \$4,300 to their wives. Both estimates are statistically significant at conventional levels.

- In addition, one can recover the impact of wages and non labor incomes on the sharing rule. For instance, a one dollar increase in the wife's wage rate (which is equivalent to an annual increase of about \$1,750 in her labor income, at the mean of hours worked by women) translates into more income being transferred to her husband. At sample mean, the transfer amounts to more than \$1,500, although this effect is not precisely estimated. Also, a one dollar increase in the husband's wage rate (equivalent to an annual increase of \$2,240 in his labor income) translates into \$600 being transferred to his wife, although again this effect is imprecisely estimated. Finally, a one dollar increase in household nonlabor income will increase the wife's nonlabor income by 70 cents; that non labor income goes mostly to the wife on average is actually a common finding of most empirical studies based on the collective framework.

- Finally, wage elasticities can be computed in two ways. A direct estimation gives a positive, significant elasticity for women, close to 0.2, while men's wage elasticities are very small and not statistically significant. The structural model also allows us to estimate the 'true' own-wage elasticities of individual labor supplies, taking into account the impact of wages on the sharing of nonlabor income. Both women's and men's elasticities are significant but smaller than those reported previously - reflecting the fact that a marginal increase in either spouse's wage rate reduces their share of the nonlabor income, which in turn increases their labor supply through an income effect. Indeed, both men's and women's labor supply elasticities with respect to nonlabor income are negative and significant.

Recent empirical developments involving cooperative models of labor supply include Donni (2003), which generalizes the standard approach to corner solutions and non-linear budget constraints, and Blundell, Chiappori, Magnac and Meghir (2007), who consider a model in which female labor supply is continuous whereas male labor supply is discrete; they show that the sharing rule can equally be recovered in this case. Moreau and Donni (2002) also introduce distribution factors, applied to French data, and take into account the non-linearity of taxation. Other empirical analyses include Bloemen (2009), Clark, Couprie and Sofer (2004) and Vermeulen (2005) on Dutch, British and Belgian data respectively.

In a series of recently published papers, several authors apply the collective model to welfare issues, including the impact of changes in the tax/benefit system, in different European countries. The basic methodology, as described in Vermeulen et al (2006), presents interesting features. One is its scope: the approach addresses standard problems of welfare analysis of labor supply, such as non linear taxation, non convex budget sets and discrete participation decisions, within a collective framework. In addition, individual preferences are more general than in the standard collective model of labor supply (Chiappori 1988, 1992) in the sense that they allow for interactions between individual leisures (that is, the marginal utility of a spouse's leisure is a function of the other spouse's labor supply). Since individual leisures are treated as public goods, the standard identification results do not apply. The identification strategy relies on a different assumption - namely, that the 'direct' trade-off between individual leisure and consumption (disregarding the impact of the spouse's leisure) is identical for singles and married individuals, and can therefore be directly estimated from the labor supply of singles; of course the additional, 'external' effect of one spouse's leisure on the other's utility can only be estimated from the sample of married couples. This approach allows to calibrate a collective model that can then be used for welfare analysis. Myck *et al* (2006) uses this framework to analyze the impact of a recent welfare reform in the UK, namely the introduction of the Working Families' Tax Credit (WFTC). In particular, they consider two hypothetical versions of the reform: one in which the recipient remains the main carer (as for the previous Family Credit), and another in which the benefit is paid to the main earner. The model allows to predict the impact of each version on the spouses' respective Pareto weights, and the corresponding labor supply responses; they conclude that, indeed, the two versions have different impacts on individual labor supplies and ultimately welfares. Similar studies have been undertaken in various countries, including Belgium, France, Germany, Italy and Spain; the findings are summarized in Myck *et al* (2006). Finally, Beninger *et al* (2006) provide a systematic comparison of the evaluations of tax policy reforms made within the unitary or the collective approaches respectively. They show, in particular, that the unitary version tends to overestimate male (and underestimate female) labor supply responses vis a vis the collective counterpart. Moreover, for a significant fraction of households, a tax reform that appears to be Pareto improving in the collective setting is found to reduce household utility in the unitary version - a possibility that had already be mentioned by the theoretical literature but had not received an empirical confirmation so far.

Another interesting analysis is provided by Lise and Seitz (2009), who study consumption inequality in the UK from 1968 to 2001. The main findings of the paper is that ignoring consumption inequality within the household produces misleading estimates of inequality. Using a rich version of the collective model that allows for public consumption and caring preferences, they reach two important conclusions. First, the standard analysis of inequality, based on adult equivalence scales and the implicit assumption of equal sharing of consumption within the household, underestimates the level of cross sectional consumption inequality in 1968 by 50%; the reason being that large differences in the earnings of husbands and wives translate into large intrahousehold inequality in consumption. Second, the large and well known rise in between household inequality during the 80's was largely offset by a drastic reduction in intrahousehold inequality, due to changes in female labor supply. As a result, inequality *between individuals*, once (properly) computed by taking into account changes in intrahousehold allocation, turns out to be practically the same in 2000 as in 1970 - a conclusion that sharply contrasts with standard studies. Other works on intrahousehold inequality include Kalugina, Radchenko and Sofer (2009a, b) and Lacroix and Radchenko (2009).

Natural experiments can provide a rich source of applications for the collective approach to labor supply. Kapan (2009) studies the impact of a change in UK divorce laws in 2000, whereby the allocation of wealth, initially based on a principle of separate ownership of assets, shifted to 'the yardstick of equal division'. A change of this kind is a typical distribution factor; however, because of its discrete nature, the analysis cannot rely on the same technique as Chiappori, Fortin and Lacroix (2002). Kapan shows how the estimation strategy can be adapted to take advantage of discrete distribution factors. He finds that, indeed, the shift resulted in an additional transfer to women, at least when their wealth was smaller than their husband's; in turn, this reallocation had a significant impact of labor supplies and individual welfares.

Finally, models involving domestic productions have been empirically analyzed in a number of contributions. For example, Apps and Rees (1996) and Rapoport, Sofer and Solaz (2004) estimate the canonical model with Australian, French and Dutch data, respectively, whereas Couprie (2007) and van Klaveren, van Praag and Maassen van den Brink (2008) consider models where the domestic good is public and present empirical results on various data sets.

References

[1] Anderson, Siwan and Jean-Marie Baland, 'The Economics of Roscas and Intrahousehold Resource Allocation', **The Quarterly Journal of Economics,** 117 (2002), 963-995.

[2] Apps, Patricia F. and Ray Rees, 'Taxation and the Household', **Journal of Public Economics,** 35 (1988), 355-369.

[3] Aronsson, Thomas, Daunfeldt, Sven-Olov and Magnus Wikstrom, 'Estimating Intrahousehold Allocation in a Collective Model with Household Production', **Journal of Population Economics,** 14 (2001), 569-584.

[4] Attanasio, Orazio and Valérie Lechene, 'Tests of Income Pooling in Household Decisions', **Review of Economic Dynamics,** 5 (2002), 720-748.

[5] Bargain, Olivier, Beblo, Miriam, Beninger, Denis, Blundell, Richard, Carrasco, Raquel, Chiuri, Maria-Concetta, Laisney, Francois, Lechene, Valérie, Moreau, Nicolas, Myck, Michal, Ruiz-Castillo, Javier and Frederic Vermeulen, 'Does the Representation of Household Behavior Matter for Welfare Analysis of Tax-benefit Policies? An Introduction', **Review of Economics of the Household,** 4 (2006), 99-111.

[6] Barmby, Tim and Nina Smith, 'Household Labour Supply in Britain and Denmark: Some Interpretations Using a Model of Pareto Optimal Behaviour', **Applied Economics,** 33 (2001), 1109-1116.

[7] Bayudan, Connie G., 'Wives' Time Allocation and Intrahousehold Power: Evidence from the Philippines', **Applied Economics,** 38 (2006), 789-804.

[8] Beninger, Denis, Bargain, Olivier, Beblo, Miriam, Blundell, Richard, Carrasco, Raquel, Chiuri, Maria-Concetta, Laisney, Francois, Lechene, Valérie, Longobardi, Ernesto, Moreau, Nicolas, Myck, Michal, Ruiz-Castillo, Javier and Frederic Vermeulen, 'Evaluating the Move to a Linear Tax System in Germany and Other European Countries', **Review of Economics of the Household,** 4 (2006), 159-180.

[9] Bloemen, Hans G.,'An Empirical Model of Collective Household Labour Supply with Non-Participation', **Economic Journal,** 120 (2010), 183-214.

[10] Blundell, Richard, Chiappori, Pierre-André and Costas Meghir, 'Collective Labor Supply with Children', **Journal of Political Economy,** 113 (2005), 1277-1306.

[11] Blundell, Richard, Chiappori, Pierre-André, Magnac, Thierry and Costas Meghir, 'Collective Labour Supply: Heterogeneity and Non-Participation', **Review of Economic Studies,** 74 (2007), 417-445.

[12] Bourguignon, Francois, Browning, Martin and Pierre-André Chiappori, 'Efficient Intra-Household Allocations and Distribution Factors: Implications and Identification', **Review of Economic Studies,** 76 (2009), 503-528.

[13] Bourguignon, Francois, Browning, Martin, Chiappori, Pierre-André and Valérie Lechene, 'Intra Household Allocation of Consumption: A Model and Some Evidence from French Data', **Annales d'Economie et de Statistique,** 29 (1993), 137-156.

[14] Browning, Martin, 'Saving and the Intra-Household Distribution of Income: An Empirical Investigation', **Ricerche Economiche,** 49 (1995), 277-292.

[15] Browning, Martin and Jens Bonke, 'Allocation within the Household: Direct Survey Evidence', University of Oxford, Working Paper No. 429, (2009).

[16] Browning, Martin, Bourguignon, Francois and Valérie Lechene, 'Income and Outcomes: A Structural Model of Intrahousehold Allocation', **Journal of Political Economy,** 102 (1994), 1067-1096.

[17] Browning, Martin and Pierre-André Chiappori, 'Efficient Intra-Household Allocations: A General Characterization and Empirical Tests', **Econometrica,** 66 (1998), 1241-1278.

[18] Browning, Martin and Metter Gørtz, 'Spending Time and Money within the Household', University of Oxford, Working Paper No. 288, (2006).

[19] Bruce, Judith, 'Homes divided', **World Development,** 17 (1989), 979-991.

[20] Burda, Michael C., Hamermesh, Daniel S. and Philippe Weil, 'The Distribution of Total Work in the EU and US', IZA Discussion Paper No. 2270, (2006).

[21] Cherchye, Laurens, De Rock, Bram, Sabbe, Jeroen and Frederic Vermeulen, 'Nonparametric Tests of Collectively Rational Consump-

tion Behavior: An Integer Programming Procedure', **Journal of Econometrics,** 147 (2008), 258-265.

[22] Cherchye, Laurens, De Rock, Bram and Frederic Vermeulen, 'The Collective Model of Household Consumption: A Nonparametric Characterization', **Econometrica,** 75 (2007), 553-574.

[23] Chiappori, Pierre-André, 'Collective Labor Supply and Welfare', **Journal of Political Economy,** 100 (1992), 437-467.

[24] Chiappori, Pierre-André, 'Rational Household Labor Supply', **Econometrica,** 56 (1988), 63-90.

[25] Chiappori, Pierre-André, 'Collective Labour Supply with Many Consumption Goods', Columbia University, (2008).

[26] Chiappori, Pierre-André and I. Ekeland, 'The Micro Economics of Group Behavior: General Characterization', **Journal of Economic Theory,** 130 (2006), 1-26.

[27] Chiappori, Pierre-André and I. Ekeland, 'The Microeconomics of Efficient Group Behavior: Identification', **Econometrica,** 77 (2009), 763-799.

[28] Chiappori, Pierre-André, Fortin, Bernard and Guy Lacroix, 'Marriage Market, Divorce Legislation, and Household Labor Supply', **Journal of Political Economy,** 110 (2002), 37-72.

[29] Clark, Andrew, Couprie, Helene and Catherine Sofer, 'La modélisation collective de l'offre de travail. Mise en perspective et application aux données britanniques', **Revue économique,** 55 (2004), 767-789.

[30] Couprie, Helene, 'Time Allocation within the Family: Welfare Implications of Life in a Couple', **Economic Journal,** 117 (2007), 287-305.

[31] d'Aspremont, Claude and Rodolphe Dos Santos Ferreira, 'Household behavior and individual autonomy', Working Paper of BETA No. 2009-17, UDS, Strasbourg, (2009).

[32] Dauphin, Anyck, El Lahga, AbdelRahmen, Fortin, Bernard and Guy Lacroix, 'Are Children Decision-Makers Within the Household?', IZA Discussion Paper No. 3728, (2008).

[33] Del Boca, Daniela and Christopher J. Flinn, 'Household Time Allocation and Modes of Behavior: A Theory of Sorts', IZA Discussion Paper No. 1821, Institute for the Study of Labor (IZA), (2005).

[34] Dercon, Stefan and Pramila Krishnan, 'In Sickness and in Health: Risk Sharing within Households in Rural Ethiopia', **Journal of Political Economy,** 108 (2000), 688-727.

[35] Donni, Olivier, 'Collective Household Labor Supply: Nonparticipation and Income Taxation', **Journal of Public Economics,** 87 (2003), 1179-1198.

[36] Donni, Olivier, 'Collective Consumption and Welfare', **Canadian Journal of Economics,** 39 (2006), 124-144.

[37] Donni, Olivier, 'Les modèles non coopératifs d'offre de travail: théorie et évidence', **L'Actualité économique**, 82 (2006), 181-206.

[38] Donni, Olivier, 'Collective Female Labour Supply: Theory and Application', **Economic Journal,** 117 (2007), 94-119.

[39] Donni, Olivier, 'Choix collectif du temps travaillé, consommation publique, et équilibre de Lindahl', **Revue économique**, 58 (2007), 535-543.

[40] Donni, Olivier, 'Labor Supply, Home Production, and Welfare Comparisons', **Journal of Public Economics,** 92 (2008), 1720-1737.

[41] Donni, Olivier, 'A Simple Approach to Investigate Intrahousehold Allocation of Private and Public Goods', **The Review of Economics and Statistics,** 91 (2009), 617-628.

[42] Donni, Olivier and Nicolas Moreau, 'Collective Labor Supply: A Single-Equation Model and Some Evidence from French Data', **Journal of Human Resources,** 42 (2007), 214-246.

[43] Dubois, Pierre and Ethan Ligon, 'Incentives and Nutrition for Rotten Kids: Intra household Food Allocation in the Philippines', UC Berkeley, (2005).

[44] Duflo, Esther, 'Grandmothers and Granddaughters: Old-Age Pensions and Intrahousehold Allocation in South Africa', **World Bank Economic Review**, 17 (2003), 1-25.

[45] Duflo, Esther and Christopher Udry, 'Intrahousehold Resource Allocation in Cote d'Ivoire: Social Norms, Separate Accounts and Consumption Choices', NBER Working Paper No. 10498, (2004).

[46] Ermisch, John and Chiara Pronzato, 'Intra-Household Allocation of Resources: Inferences from Non-resident Fathers' Child Support Payments', **Economic Journal,** 118 (2008), 347-362.

[47] Fong, Yuk-fai and Junsen Zhang, 'The Identification of Unobservable Independent and Spousal Leisure', **The Journal of Political Economy,** 109 (2001), 191-202.

[48] Fortin, Bernard and Guy Lacroix, 'A Test of the Unitary and Collective Models of Household Labour Supply', **Economic Journal,** 107 (1997), 933-955.

[49] Gimenez-Nadal, Jose Ignacio and Almudena Sevilla-Sanz, 'The Time-crunch Paradox', Department of Economics, University of Oxford, WP 483, (2010).

[50] Goldstein, Markus P., 'Intrahousehold Efficiency and Individual Insurance in Ghana', LSE STICERD Research Paper No. DEDPS38, (2004).

[51] Haddad, Lawrence and John Hoddinott, 'Women's Income and Boy-Girl Anthropometric Status in the Cote d'Ivoire', **World Development,** 22 (1994), 543-553.

[52] Hoddinott, John and Lawrence Haddad, 'Does Female Income Share Influence Household Expenditures? Evidence from Cote d'Ivoire', **Oxford Bulletin of Economics and Statistics,** 57 (1995), 77-96.

[53] Hotchkiss, Julie L., 'Do Husbands and Wives Pool Their Resources?: Further Evidence', **Journal of Human Resources,** 40 (2005), 519-531.

[54] Kalugina, Ekaterina, Radtchenko, Natalia and Catherine Sofer, 'How Do Spouses Share Their Full Income? Identification Of The Sharing Rule Using Self-Reported Income', **Review of Income and Wealth,** 55 (2009), 360-391.

[55] Kalugina, Ekaterina, Sofer, Catherine and Natalia Radtchenko, 'Intra-household Inequality in Transitional Russia', **Review of Economics of the Household,** 7 (2009), 447-471.

[56] Kapan, T., 'Essays in Household behavior', PhD Dissertation, Columbia University, (2009).

[57] Klaveren, Chris, Praag, Bernard and Henriette Maassen van den Brink, 'A Public Good Version of the Collective Household Model: an Empirical Approach with an Application to British Household Data', **Review of Economics of the Household,** 6 (2008), 169-191.

[58] Lacroix, Guy and Natalia Radtchenko, 'The Changing Intra-household Resource Allocation in Russia', **Journal of Population Economics,** DOI No. 10.1007/s00148-009-0275-2, (2009).

[59] Lechene, Valérie and Ian Preston, 'Demand properties in household Nash equilibrium', IFS Working Paper No. W07/01, Institute for Fiscal Studies, (2007).

[60] Ligon, Ethan, 'Dynamic Bargaining in Households (with application to Bangladesh)', UC Berkeley, Working Paper Series No. 972, (2002).

[61] Lise, Jeremy and Shannon Seitz, 'Consumption Inequality and Intra-Household Allocations', **Labor and Demography,** EconWPA, No. 0504001, (2005).

[62] Lundberg, Shelly J., Pollak, Robert A. and Terence J. Wales, 'Do Husbands and Wives Pool Their Resources? Evidence from the United Kingdom Child Benefit', **The Journal of Human Resources,** 32, 463-480 (1997).

[63] Luo, Guo Ying, 'Collective Decision-Making and Heterogeneity in Tastes', **Journal of Business & Economic Statistics,** 20 (2002), 213-226.

[64] Myck, Michal, Bargain, Olivier, Beblo, Miriam, Beninger, Denis, Blundell, Richard, Carrasco, Raquel, Chiuri, Maria-Concetta, Laisney, Francois, Lechene, Valérie, Longobardi, Ernesto and Nicolas Moreau, 'The Working Familiess' Tax Credit and Some European Tax Reforms in a Collective Setting', **Review of Economics of the Household,** 4 (2006), 129-158.

[65] Phipps, Shelley A. and Peter S. Burton, 'What's Mine Is Yours? The Influence of Male and Female Incomes on Patterns of Household Expenditure', **Economica,** 65 (1998), 599-613.

[66] Rapoport, Benoit, Sofer, Catherine and Anne Solaz, 'Household Production in a Collective Model: Some New Results ', **Journal of Population Economics,** DOI No. 10.1007/s00148-010-0308-x, (2009).

[67] Schultz, T. Paul, 'Testing the Neoclassical Model of Family Labor Supply and Fertility', **The Journal of Human Resources,** 25 (1990), 599-634.

[68] Strauss, John, Mwabu, G. and Kathleen Beegle, 'Intrahousehold Allocations: a Review of Theories and Empirical Evidence', **Journal of African Economies,** 9 (2000), 83-143.

[69] Thomas, Duncan, 'Intra-Household Resource Allocation: An Inferential Approach', **The Journal of Human Resources,** 25 (1990), 635-664.

[70] Udry, Christopher, ‘Gender, Agricultural Production, and the Theory of the Household’, **The Journal of Political Economy,** 104 (1996), 1010-1046.

[71] Vermeulen, Frederic, ‘And the Winner is... An Empirical Evaluation of Unitary and Collective Labour Supply Models’, **Empirical Economics,** 30 (2005), 711-734.

[72] Vermeulen, Frederic, ‘A Collective Model for Female Labour Supply with Non-Participation and Taxation’, **Journal of Population Economics,** 19 (2006), 99-118.

[73] Vermeulen, Frederic, ‘Collective Household Models: Principles and Main Results’, **Journal of Economic Surveys,** 16 (2002), 533-564.

[74] Vermeulen, Frederic, Bargain, Olivier, Beblo, Miriam, Beninger, Denis, Blundell, Richard, Carrasco, Raquel, Chiuri, Maria-Concetta, Laisney Francois, Lechene, Valérie, Moreau, Nicolas and Michal Myck, ‘Collective Models of Labor Supply with Nonconvex Budget Sets and Nonparticipation: A Calibration Approach’, **Review of Economics of the Household,** 4 (2006), 113-127.

[75] Vermeulen, Frederic and Watteyne, André, ‘Quand un et un ne font plus deux’, **L’Actualité Economique,** 82 (2006), 155-180.

[76] Ward-Batts, Jennifer, ‘Out of the Wallet and into the Purse: Using Micro Data to Test Income Pooling’, **Journal of Human Resources,** 43 (2008), 325-351.

Reference	Outcome	Country	Df's (Table 1)
Anderson & Baland (2002)	Participation in a rosca	Kenya	1
Aronsson *et al* (2001)	Leisure demand	Sweden	2,3,4,5,6
Attanasio & Lechene (2002)	Commodity demands;	Mexico	1
	influence on various decisions		
Barmby & Smith (2001)	labor supplies	Denmark, UK	2
Bayudan (2006)	Female labor supply	Philippines	2, 11
Bourguignon *et al* (1993)	Commodity demands	France	1
Browning (1995)	Saving	Canada	1
Browning & Bonke (2006)	Commodity demands	Denmark	1,7,8,10
Browning & Gørtz (2006)	Commodity demands, leisures	Denmark	2,4,7
Browning *et al* (1994)	Demand for clothing	Canada	1,4,7
Browning & Chiappori (1998)	Commodity demands	Canada	1,4
Chiappori *et al* (2002)	labor supplies	US	6
Couprie (2007)	labor supply and leisure	UK	2,3
Donni (2007)	labor supplies, demands	France	1,7
Duflo (2003)	Child health	South Africa	1
Ermisch & Pronzato (2006)	Child support payments	UK	1
Fortin & Lacroix (1997)	Joint labor supply	Canada	1,2
Haddad & Hoddinott (1994)	Child health	Cote D'Ivoire	1
Hoddinott & Haddad (1995)	Food, alcohol and tobacco	Cote D'Ivoire	1
Lundberg, *et al* (1997)	Clothing demands	UK	3
Oreffice (2008)	Labor supply	US	1
Phipps & Burton (1998)	Commodity demands	Canada	1
Schultz (1990)	labor supplies and fertility	Thailand	3
Thomas (1990)	Child health	Brazil	3
Udry (1996)	Farm production	Burkina Faso	9
Vermeulen (2005)	labor supplies	Netherlands	3,4,12
Ward-Batts (2008)	Household demands	UK	3

Table 2: Empirical collective studies

Model	Budget shares ($\times$100)	
	Wife	Husband
Food at home	13.9	20.7
Restaurants	9.9	12.8
Clothing	16.3	7.1
Alcohol and tobacco	6.3	11.8
Transport	22.0	27.6
Personal services	15.2	12.1
Recreation	16.4	7.8

Table 3: Budget shares for husbands and wives