The Economics of the Family

Chapter 5: The collective model: a formal analysis

Martin Browning
Department of Economics, Oxford University

Pierre-André Chiappori
Department of Economics, Columbia University

Yoram Weiss
Department of Economics, Tel Aviv University

August 2007

1 Collective demand functions: a general characterization

1.1 The collective household utility function.

The basic aspects of the collective model have been described in the previous chapter. As stated earlier, the particular form adopted has testable implications for demand functions. We now describe these implications in detail. We start with the most general version of the model with individual preferences of the form $u^s\left(\mathbf{Q}, \mathbf{q}^a, \mathbf{q}^b\right)$. This allows for any type of consumption externalities between agents. We define the *collective household utility function* by

$$u^f\left(\mathbf{q}, \mathbf{Q}, \mu\right) = \max_{\mathbf{q}^a} \left\{\mu u^a\left(\left(\mathbf{Q}, \mathbf{q}^a, \mathbf{q} - \mathbf{q}^a\right)\right) + u^b\left(\mathbf{Q}, \mathbf{q}^a, \mathbf{q} - \mathbf{q}^a\right)\right\} \tag{1}$$

where μ is a function of $(\mathbf{P}, \mathbf{p}, x, \mathbf{z})$ and z is a vector of distribution factors. We shall always assume that $\mu\left(.\right)$ is zero homogeneous in $(\mathbf{P}, \mathbf{p}, x)$ and any elements of $\mathbf{z}$ that are denominated in monetary terms.

At this level of generality, the distinction between public and private goods is somewhat blurred, and we can leave it aside for the moment. We thus adopt a general notation with $\mathbf{g} = (\mathbf{q}, \mathbf{Q})$ denoting the household demand and $\mathbf{r} = (\mathbf{p}, \mathbf{P})$ denoting the corresponding price vector. Then the household's behavior is described by the maximization of $u^f\left(\mathbf{g},\mu\right)$ under the household budget constraint $\mathbf{r}'\mathbf{g} = x$.

1.2 Structural and observable demand.

The household's program is:

$$\max_{\mathbf{g}} u^f\left(\mathbf{g},\mu\right) \text{ subject to } \mathbf{r}'\mathbf{g} = x \tag{2}$$

which generates *collective demand functions*, $\tilde{\mathbf{g}}\left(\mathbf{r}, x, \mu\right)$. It is important to emphasise that this program is *not* equivalent to standard utility maximization (the unitary model) because u^f varies with μ, which in turn depends on prices, income and distribution factors. For any *fixed* μ, $\tilde{\mathbf{g}}\left(.\right)$ is a standard demand function. From standard consumer theory, we know that it satisfies Slutsky symmetry and negativeness. Define the generic Slutsky matrix element by:

$$\sigma_{ij} = \frac{\partial \tilde{g}_i}{\partial r_j} + \hat{\xi}_j \frac{\partial \tilde{g}_i}{\partial x} \tag{3}$$

and denote the Slutsky matrix by $\Sigma = [\sigma_{ij}]_{i,j}$. We have that Σ is symmetric and negative semi-definite. Rearranging (3) we see that the *Marshallian response* of the demand for good i to changes in the price of good j ($\frac{\partial \tilde{g}_i}{\partial r_j}$) can be decomposed into the difference between a substitution effect (σ_{ij}) and an income effect ($\hat{\xi}_j \frac{\partial \tilde{g}_i}{\partial x}$). The intuition is that a marginal increase in the price of any good i affects, among other things, the real income (the purchasing power) of all agents. The *substitution* term σ_{ij} represents the effect of the infinitesimal variation if it was fully *compensated* in income (that is, accompanied by a variation in income sufficient to exactly offset the loss in purchasing power); for that reason, we often talk of compensated demand. The *income* effect, on the other hand, reflects the fact that the price increase decreases the agent's purchasing power in proportion to the quantity purchased, which in turn influences the demand.

For a fixed Pareto weight the function $\tilde{\mathbf{g}}$ behaves exactly as a standard 'unitary' demand. It is crucial to realize, however, that $\tilde{\mathbf{g}}$, although conceptually relevant, is *not* an operative notion from an empirical perspective. Indeed, $\tilde{\mathbf{g}}$ cannot be observed unless one could change prices and income without modifying μ. Since, in general, μ does depend on $(\mathbf{r}, x)$ this can be, at best, a thought experiment. What we *do* observe is the household demand, in which price and income variations affect both $\tilde{\mathbf{g}}$ and μ. Thus the empirically relevant concept is the demand function defined by:

$$\hat{\mathbf{g}}(\mathbf{r}, x) = \tilde{\mathbf{g}}(\mathbf{r}, x, \mu(\mathbf{r}, x)) \tag{4}$$

where we have, for the moment, dropped the distribution factors.[1] Thus, we make a distinction between the 'structural' demand function, $\tilde{\mathbf{g}}(\mathbf{r}, x, \mu)$, and the observable demand function, $\hat{\mathbf{g}}(\mathbf{r}, x)$. Again, the difference between these collective demand functions and the unitary model (Marshallian) demand functions is the presence of the Pareto weight in the demands.

For the observable demand function we have:

$$\begin{aligned} \frac{\partial \hat{g}_i}{\partial r_j} &= \frac{\partial \tilde{g}_i}{\partial r_j} + \frac{\partial \tilde{g}_i}{\partial \mu}\frac{\partial \mu}{\partial r_j} \\ \frac{\partial \hat{g}_i}{\partial x} &= \frac{\partial \tilde{g}_i}{\partial x} + \frac{\partial \tilde{g}_i}{\partial \mu}\frac{\partial \mu}{\partial x} \end{aligned} \tag{5}$$

Thus we can decompose the price effect into a Marshallian response (the first term on the right hand side) and a collective effect (the second term), which operates through variations of the Pareto weight μ. Figure 1 illustrates for

---

[1] We shall maintain the ˆ notation for an observable function and ˜ for structural throughout the book. Think of the ˆ as denoting a function that could be estimated.

two goods. We start at point I and prices and income $(\mathbf{r}, x)$. We then change prices so that good 1 is cheaper; denote the new prices $(\mathbf{r}', x)$. The substitution effect is given by the move from I to II and the income effect is II to III. The collective effect, the final term in (5), is given by the move from III to IV.

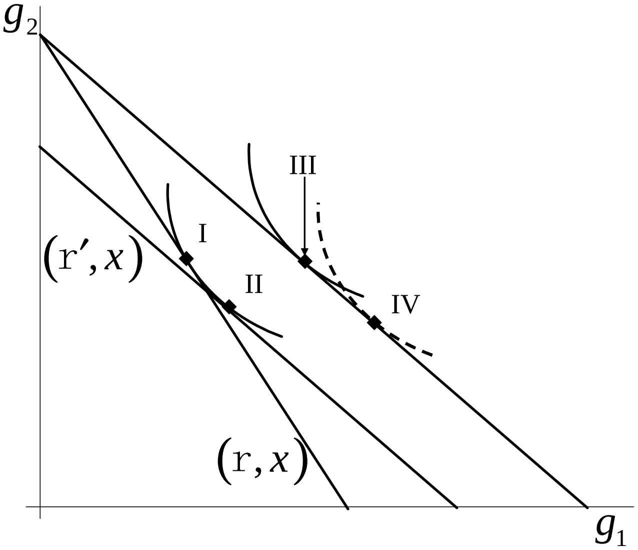

Figure 1: Collective price responses.

1.3 The Slutsky matrix for collective demands.

Using the *observable* functions $\hat{\mathbf{g}}(.)$, we can define the observable Slutsky matrix by its general term:

$$s_{ij} = \frac{\partial \hat{g}_i}{\partial r_j} + \hat{g}_j \frac{\partial \hat{g}_i}{\partial x} \tag{6}$$

with the quasi-Slutsky matrix $S = [s_{ij}]_{i,j}$. From (5) this can be written as:

$$s_{ij} = \left[\frac{\partial \tilde{g}_i}{\partial r_j} + \tilde{g}_j \frac{\partial \tilde{g}_i}{\partial x}\right] + \frac{\partial \tilde{g}_i}{\partial \mu}\left[\frac{\partial \mu}{\partial r_j} + \tilde{g}_j \frac{\partial \mu}{\partial x}\right] \tag{7}$$

From (3), the first term between brackets is the substitution term σ_{ij}. We adopt the following notation:

$$\begin{aligned} D_\mu \tilde{\mathbf{g}} &= \left[\frac{\partial \tilde{g}_i}{\partial \mu}\right]_i \\ \mathbf{v} &= \left[\frac{\partial \mu}{\partial r_j} + \tilde{g}_j \frac{\partial \mu}{\partial x}\right]_j \end{aligned} \tag{8}$$

This gives:

$$S = \Sigma + R = \Sigma + \left(D_\mu \tilde{\mathbf{g}}\right).\mathbf{v}' \tag{9}$$

so that the Slutsky matrix of the observable collective demand $\hat{\mathbf{g}}(\mathbf{r}, x)$ is the sum of a conventional Slutsky matrix Σ, which is symmetric and negative semi-definite and an additional matrix R. The latter is the product of a column vector $(D_\mu \tilde{g})$ and a row vector $(\mathbf{v}')$. An outer product has rank of at most one. Indeed, for any vector $\mathbf{w}$ such that $\mathbf{v}'.\mathbf{w} = 0$ we have that $R.\mathbf{w} = 0$.

This analysis and the homogeneity assumption on $\mu(.)$ yields that the necessary conditions for the collective model demands are the *generalised Slutsky conditions*:

$$\hat{\mathbf{g}}(\mathbf{r}, x) \text{ is zero homogeneous} \tag{10}$$

$$\mathbf{r}'\hat{\mathbf{g}}(\mathbf{r}, x) \equiv x \tag{11}$$

$$\begin{gathered} S \text{ is the sum of a symmetric,} \\ \text{negative semi-definite matrix and a rank 1 matrix} \end{gathered} \tag{12}$$

(Browning and Chiappori (1998)). We denote the third property *SNR1.* These conditions obviously generalize the conventional Slutsky conditions in the unitary setting. In the particular case of $R = 0$, indeed, we are back to the predictions of the unitary model. This is the case when either μ is constant (so that $\mathbf{v} = 0$) or when $\tilde{\mathbf{g}}$ does not depend on μ (so that $D_\mu \tilde{g} = 0$). In general, however, R is not zero, and the predictions of the model deviate from those of the unitary model; in a sense, matrix R summarizes this deviation. The main result is that *this deviation is only one-dimensional* - which formally translates into the rank of R being at most one. This is a strong result because the *size* of matrix R can be quite large - as many as goods the household buys.[2]

---

[2] In general, R has $(n + N)$ eigenvalues (possibly complex); the rank condition means that all of them, but maybe one, are equal to zero. Equivalently, one can find a basis in which all of the $(n + N)$ columns of R but one are identically zero.

The result has a simple, geometric intuition based on Figure **??** in chapter 3. The move from I to II represents the variation that would obtain if μ was kept constant; as such, it does not violate Slutsky symmetry. The violation comes from the second component,that is, the move from II to III which reflects the impact of changes in μ. This change takes place *along the Pareto frontier*. But this frontier is one dimensional, independently of the number of commodities in the economy. Consequently the matrix R.has at most rank 1.

Finally, it can be shown that these conditions are also sufficient for the existence of a collective model. Indeed, Chiappori and Ekeland (2004) show that any 'smooth'[3] demand function satisfying the three properties above (homogeneity, adding-up, SNR1) can locally be constructed as the collective demand of a well chosen household. This is a very difficult result, that requires complex mathematical tools; it constitutes a generalization of the classical, 'integrability' result in standard consumer theory.

1.4 Distribution factors

We may now reintroduce distribution factors. An interesting feature is that such factors do not change the Pareto frontier, but only the Pareto weight. In geometrical terms, thus, they can only generates moves along the Pareto frontier (from II to III in Figure **??**). This suggests that analyzing the impact of distribution factors may help understanding the nature and the form of such moves. This intuition can be given a formal translation. Equation (4) above can now be rewritten as

$$\hat{\mathbf{g}}\left(\mathbf{r}, x, \mathbf{z}\right) = \tilde{\mathbf{g}}\left(\mathbf{r}, x, \mu\left(\mathbf{r}, x, \mathbf{z}\right)\right) \tag{13}$$

From this, one can readily compute the consequences of a marginal change in distribution factor z_k on the collective demand for commodity i:

$$\frac{\partial \hat{g}_i}{\partial z_k} = \frac{\partial \tilde{g}_i}{\partial \mu}\frac{\partial \mu}{\partial z_k} \tag{14}$$

This equation leads to several interesting conclusions. First, comparing the effect of different distribution factors, say z_k and z_l, we find that (assuming $\partial g_i / \partial z_l \neq 0$):

$$\frac{\partial \hat{g}_i / \partial z_k}{\partial \hat{g}_i / \partial z_l} = \frac{\partial \mu / \partial z_k}{\partial \mu / \partial z_l} \tag{15}$$

---

[3] Technically, the result has been proved for twice continuously differentiable demand functions.

The right hand side term is independent of the good we are considering. Hence we have the *proportionality property* that the ratio of derivatives with respect to two sharing factors is the same for all goods. The result that the impact of z_k and z_l must be proportional across commodities is very important empirically, and can be given various equivalent forms; for instance, we can write that[4]

$$\frac{\partial \hat{g}_i}{\partial z_k} = \frac{\partial \mu / \partial z_k}{\partial \mu / \partial z_l} . \frac{\partial \hat{g}_i}{\partial z_l} \tag{16}$$

If the impact of a change in z_k on household demand for good i is, say, twice as large as that of z_l, then the same must be true for *all* commodities; and we can actually conclude that the impact of z_k on the Pareto weight μ is twice as large as that of z_l. Intuitively, whatever the number of distribution factors, they only operate through their impact on μ; hence their impact is one-dimensional (in a sense, it is as if there was one distribution factor only). This prediction is easy to test empirically, provided that more than one factor is available; possible tests will be discussed in the next chapter.

Another interesting feature of (14) is that it provides additional information about the structure of price and income effects in the collective demand. From (14), we have that:

$$\begin{aligned} \frac{\partial \tilde{g}_i}{\partial \mu} &= \frac{1}{\partial \mu / \partial z_k} \frac{\partial \hat{g}_i}{\partial z_k} \text{ for all } i, k \\ &= \lambda_k \frac{\partial \hat{g}_i}{\partial z_k} \text{ for all } i, k \end{aligned} \tag{17}$$

so that (9) becomes

$$S = \Sigma + R = \Sigma + \lambda_k . \left(D_{z_k} \hat{\mathbf{g}} \right) . \mathbf{v}' \text{ for any } k \tag{18}$$

Thus regarding price and income effects, not only is the deviation from the unitary model (the ‘collective effect’) one-dimensional, but it is closely related to the impact of distribution factors on demand. This is a surprising property, since it establishes links between the impact of purely economic factors - prices and incomes - and that of variables of a different type (say, divorce laws or sex ratios). Again, empirical tests of this property will be discussed in the next chapter.

---

[4] Equivalently, the matrix $D_{\mathbf{z}}\boldsymbol{\xi}$ with general terms $\frac{\partial \xi_i}{\partial z_k}$ is of rank (at most) one.

1.5 Larger households

How do these properties extend to larger households? Given the intuitions described above, the answer is in fact straightforward. Suppose there are T agents in the household. We still assume efficiency so that the collective household utility function is defined as:

$$u^f(\mathbf{q},\mathbf{Q},\mu)=\max_{\mathbf{g}}\left\{\sum_{s=1}^{T}\mu_s u^s\left(\left(\mathbf{Q},\mathbf{q}^1,...,\mathbf{q}^T\right)\right)\right\}$$
$$\text{subject to }\sum_{s=1}^{T}\mathbf{q}^s=\mathbf{q} \tag{19}$$

where the Pareto weights are normalized by $\mu_T=1$. Again, the μ_t are functions of prices, income and distribution factors. The household maximizes this utility under budget constraint. With the same notations as above, we can define a 'structural' demand function, $\mathbf{g}\left(\mathbf{r},x,\mu_1,...,\mu_{T-1}\right)$ as the solution to (19); note that it now depends on $T-1$ Pareto weights. As before, the empirically relevant concept is the observable demand function is defined by:

$$\hat{\mathbf{g}}(\mathbf{r},x,\mathbf{z})=\tilde{\mathbf{g}}\left(\mathbf{r},x,\mu_1(\mathbf{r},x,\mathbf{z}),...,\mu_{T-1}(\mathbf{r},x,\mathbf{z})\right) \tag{20}$$

Similar computations give

$$s_{ij}=\left[\frac{\partial\tilde{g}_i}{\partial r_j}+\tilde{g}_j\frac{\partial\tilde{g}_i}{\partial x}\right]+\sum_{t=1}^{T-1}\frac{\partial\tilde{g}_i}{\partial\mu_t}\left[\frac{\partial\mu_t}{\partial r_j}+\tilde{g}_j\frac{\partial\mu_t}{\partial x}\right] \tag{21}$$

Again, the collective Slutsky matrix is the sum of a symmetric, negative matrix Σ and of a 'deviation' R. However, R is now the sum of $T-1$ terms of the form $\left(D_{\mu_s}\tilde{\mathbf{g}}\right).\mathbf{v}_t'$, in which the vector $\mathbf{v}_t$ is of general term $\left[\frac{\partial\mu_t}{\partial r_j}+\hat{\xi}_j\frac{\partial\mu_t}{\partial x}\right]$; indeed, the deviations now come from the $T-1$ functions μ_t. In particular, its rank is at most $T-1$.

The generalised Slutsky conditions for a T person household are given by:

$$\begin{gathered}\mathbf{g}(\mathbf{r},x,\mathbf{z})\text{ is zero homogeneous}\\ \mathbf{r}'\hat{\mathbf{g}}(\mathbf{r},x,\mathbf{z})\equiv x\\ T\text{ is the sum of a symmetric,}\\ \text{negative semi-definite matrix and a rank }T-1\text{ matrix}\end{gathered} \tag{22}$$

These conditions are sometimes called the SNR($T-1$) conditions. They have a nicely nested structure, in the sense that SNR(k) is a special case of SNR($k+1$). They are more restrictive, the larger the number of goods and the smaller the size of the household. Note, in particular, that when the number of persons in the household is equal to (or larger than) the number of commodities, the SNR($T-1$) conditions are not restrictive at all: any $(n+N) \times (n+N)$ matrix satisfies them (just take $\Sigma = 0$). This is by no means a problem in real life, since the number of commodities available is very large. However, it may be an issue in econometric estimation, which typically use a small number of aggregate 'commodities'.

1.6 Children

Finally, we may briefly come back to the issue of children. We described in the previous chapters two different ways of modelling children: either as a 'public good' that enters parents' utility or as a genuine decision maker. The previous analysis sheds light on the respective implications of these options. In the first case the household has two decision makers, whereas it has three in the second. According to the generalized Slutsky conditions, the demand function should satisfy SNR1 in the first case, but not in the second (it only satisfies SNR2). In words: one can devise a test allowing to find out how many decision makers there are in the household (the precise implementation of the test will be described in the next chapter).

Clearly, one has to keep in mind the limits of this exercise. What the theory predicts is that the rank of the R matrix is *at most* $T-1$. Still, it can be less. For instance, if μ_s and $\mu_{s'}$ have a similar impact on household demand (in the sense that $D_{\mu_s}\tilde{\mathbf{g}}$ and $D_{\mu_{s'}}\tilde{\mathbf{g}}$ are colinear) then matrix R will be of rank $T-2$. In other words, if a household demand is found to satisfy SNRk, the conclusion is that there are *at least k* decision makers; there may be more, but there cannot be less. Or, in the case of children: a demand satisfying SNR1 is consistent with children being decision makers; however, if it satisfies SNR2 and not SNR1, then the hypothesis that children are not decision makers is rejected.

2 Duality in the collective model

2.1 The collective expenditure function.

The standard tools of duality theory which have been developed in consumer theory can readily be extended to collective models. They provide useful ways of analyzing welfare issues in the collective setting. We introduce these notions for a two-person household; the extension to larger units is straightforward. The first concept is that of *collective expenditure function*, denoted E, which is defined by:

$$\begin{aligned} E\left(\boldsymbol{\pi}, u^a, u^b\right) &= \min_{\mathbf{g},\mathbf{q}^a\mathbf{q}^b,\mathbf{Q}} \boldsymbol{\pi}'\mathbf{g} \\ \text{subject to } u^i\left(\mathbf{q}^a, \mathbf{q}^b, \mathbf{Q}\right) &\geq u_s, \ s = a, b. \\ \text{and } \mathbf{g} &= \left(\mathbf{q}^a + \mathbf{q}^b, \mathbf{Q}\right) \qquad (23) \end{aligned}$$

The collective expenditure function depends on prices and on two utility levels $\left(u^a, u^b\right)$; it represents the minimum level of expenditures needed at these prices to achieve these utilities. One can then define the *compensated collective demand function*, $\tilde{\mathbf{g}}\left(\mathbf{r}, u^a, u^b\right)$, as a solution to program (23). A key remark is that the definition of household collective expenditure and demand functions depends only on individual preferences and not on the household's decision process.

The properties of the functions just defined are analogous to those of their standard counterpart. In particular:

Proposition 1 *We have:*

$$\tilde{\mathbf{g}}\left(\mathbf{r}, u^a, u^b\right) = \nabla_r E\left(\mathbf{r}, u^a, u^b\right) \qquad (24)$$

where $\nabla_r E$ *denotes the gradient of* E *with respect to* $\mathbf{r}$ *(that is, the vector of partial derivatives* $\partial E / \partial r_j$*).*

The result is a consequence of the envelope theorem applied to program (23). In the case of egotistic preferences of the form $u^s\left(\mathbf{q}^s, \mathbf{Q}\right)$, we have further results. Define the compensated demand for public goods by $\tilde{\mathbf{Q}}\left(\mathbf{p}, \mathbf{P}, u^a, u^b\right)$. The we have:

Proposition 2 *If* u^s *only depends on* $\left(\mathbf{q}^s, \mathbf{Q}\right)$, $s = a, b$*, then:*

$$\begin{aligned} E\left(\mathbf{p}, \mathbf{P}, u^a, u^b\right) &\leq e^a\left(\mathbf{p}, \mathbf{P}, u^a\right) + e^b\left(\mathbf{p}, \mathbf{P}, u^b\right) \\ E\left(\mathbf{p}, \mathbf{P}, u^a, u^b\right) &\geq e^a\left(\mathbf{p}, \mathbf{P}, u^a\right) + e^b\left(\mathbf{p}, \mathbf{P}, u^b\right) - \mathbf{P}'\tilde{\mathbf{Q}}\left(\mathbf{p}, \mathbf{P}, u^a, u^b\right) \qquad (25) \end{aligned}$$

where $e^s\left(\mathbf{p}, \mathbf{P}, u^s\right)$ *denotes the (individual) expenditure function of member* s*.*

Proof. *The last inequality stems from the definition of individual expenditure functions, since*

$$e^s\left(\mathbf{p},\mathbf{P},u^s\right)\leq\mathbf{p}'\mathbf{q}^s\left(\mathbf{p},\mathbf{P},u^a,u^b\right)+\mathbf{P}'\mathbf{Q}\left(\mathbf{p},\mathbf{P},u^a,u^b\right)\tag{26}$$

For the first inequality, let $\left(\bar{\mathbf{q}}^s,\bar{\mathbf{Q}}^s\right)$ *denote the* individual *compensated demand of* s *(corresponding to prices* $\mathbf{p},\mathbf{P}$ *and utility* u^s*). If* $\bar{\mathbf{Q}}^a=\bar{\mathbf{Q}}^b$ *the conclusion follows. If not, say* $\bar{\mathbf{Q}}^a>\bar{\mathbf{Q}}^b$*, then*

$$\begin{aligned}u^a\left(\bar{\mathbf{q}}^a,\bar{\mathbf{Q}}^a\right)&=u^a\\u^b\left(\bar{\mathbf{q}}^b,\bar{\mathbf{Q}}^a\right)&>u^b\end{aligned}\tag{27}$$

therefore

$$\begin{aligned}E\left(\mathbf{p},\mathbf{P},u^a,u^b\right)&\leq\mathbf{p}'\left(\bar{\mathbf{q}}^a+\bar{\mathbf{q}}^b\right)+\mathbf{P}'\bar{\mathbf{Q}}^a\\&\leq\mathbf{p}'\left(\bar{\mathbf{q}}^a+\bar{\mathbf{q}}^b\right)+\mathbf{P}'\bar{\mathbf{Q}}^a+\mathbf{P}'\bar{\mathbf{Q}}^b\\&=e^a\left(\mathbf{p},\mathbf{P},u^a\right)+e^b\left(\mathbf{p},\mathbf{P},u^b\right)\end{aligned}\tag{28}$$

■

2.2 Indirect utilities

We can also define indirect utility functions. A first concept of indirect utility is:

$$\begin{aligned}\tilde{v}^f\left(\mathbf{r},x,\mu\right)&=\max_{\mathbf{g}}\ \mu u^a\left(\mathbf{g}\right)+u^b\left(\mathbf{g}\right)\\\text{subject to }\mathbf{r}'\mathbf{g}&=x\end{aligned}\tag{29}$$

This version is the direct equivalent, in the collective setting, of the indirect utility concept in standard consumer theory. Again, $\tilde{v}$ only depends on preferences, not on the decision process; different processes result in different μ's, hence to different values for $\tilde{v}^f\left(.\right)$. A second definition is obtained by plugging the particular Pareto weight adopted by the household into the previous definition. In this case the household's *collective indirect utility* is defined as:

$$\hat{v}^f\left(\mathbf{r},x,\mathbf{z}\right)=\tilde{v}^f\left(\mathbf{r},x,\mu\left(\mathbf{r},x,\mathbf{z}\right)\right)$$

The function $\tilde{v}^s\left(.\right)$ characterizes the utility level reached by the household, at prices $\mathbf{r}$ and for income x, as a function of the weight μ of the partner

a (b's weight being normalized to unity). On the other hand, the collective indirect utility of the household, $\hat{v}^f(.)$, takes into account the Pareto weight resulting from the actual decision process.

Finally, one may define the *collective indirect utility of a member* as the level of utility ultimately reached by this member as a function of prices and income and distribution factors. If $\hat{\mathbf{g}}(\mathbf{r},x,\mathbf{z})$ denotes a solution to the initial program (29), then:

$$\begin{aligned}\hat{v}^a(\mathbf{r},x,\mathbf{z}) &= u^a(\hat{\mathbf{g}}(\mathbf{r},x,\mathbf{z})) \\ \hat{v}^b(\mathbf{r},x,\mathbf{z}) &= u^b(\hat{\mathbf{g}}(\mathbf{r},x,\mathbf{z}))\end{aligned}$$

It is important to note that the definition of a's indirect utility depends not only on a's preferences, but also on the whole decision process. In other words, collective indirect utilities are specific to a particular match between agents and a particular decision rule (summarized by the function μ). This is in sharp contrast with the unitary case, where there exists a one-to-one correspondence between direct and indirect utility at the individual level. Still, the basic interpretation remains: indirect utilities provide information on each agent's final welfare once all aspects of the decision process have been taken into account. In particular, they play a key role in any normative analysis.

2.3 Welfare

An important application of consumer theory relates to welfare issues, such as the cost-benefit evaluation of economic reforms. A standard tool is the notion of *compensating variation.* Consider a reform that changes the price vector from $\mathbf{r}$ to $\mathbf{r}'$. For an agent with initial income x, the compensating variation (CV) is defined as the change in income that would be needed to exactly compensate the agent. That is, the income that would allow her to remain on the same indifference curve. For a single person this is defined by:

$$CV = e\left(\mathbf{r}', v(\mathbf{r},x)\right) - x$$

where e and v respectively denote the agent's expenditure and indirect utility functions. This concept can directly be extended to a collective setting. This leads to the following definition:

Definition 3 *The* potentially compensating variation *is the function* $\Gamma(.)$ *such that:*

$$\Gamma_1\left(\mathbf{r},\mathbf{r}',x,\mathbf{z}\right) = E\left(\mathbf{r}', \hat{v}^a(\mathbf{r},x,\mathbf{z}), \hat{v}^b(\mathbf{r},x,\mathbf{z})\right) - x$$

In words, consider a household in which, before the reform, total income is x and member s's utility is $u^s = \hat{v}^s(\mathbf{r}, x, \mathbf{z})$. The potentially compensating variation measures the change in income that has to be given to the household for the previous utility levels to be affordable at the new prices $\mathbf{r}'$. Natural as this extension may seem, it nevertheless raises problems that are specific to a multi-person setting. The variation is *potentially* compensating in the following sense: the additional income thus measured could, *if allocated appropriately within the household*, enable members to reach their pre-reform utility levels. That is, the income $x + \Gamma(\mathbf{r}, \mathbf{r}', x, \mathbf{z})$ has the property that the utilities $\left(\hat{v}^a(\mathbf{r}, x, \mathbf{z}), \hat{v}^b(\mathbf{r}, x, \mathbf{z})\right)$ belong to the Pareto frontier at prices $\mathbf{r}'$. What is *not* guaranteed, however, is that the point $\left(\hat{v}^a(\mathbf{r}, x, \mathbf{z}), \hat{v}^b(\mathbf{r}, x, \mathbf{z})\right)$ will still be chosen on the new frontier. In other words, the compensation is such that the welfare level of each members could be maintained despite the reform. Whether the household will choose to do so is a different story.

The idea is illustrated in Figure 2. The potentially compensating variation is such that the new frontier (the dashed frontier) goes through $uu = \left(\hat{v}^a(\mathbf{r}, x, \mathbf{z}), \hat{v}^b(\mathbf{r}, x, \mathbf{z})\right)$. However, the reform changes both the frontier and the Pareto weights. While the initial allocation uu is still affordable (it belongs to the new frontier), the household may instead choose the allocation uu'. It follows that although both members could have been exactly compensated, in practice one partner will strictly gain from the reform (a in Figure 2), whilst the other will strictly lose. This is despite the fact that, as drawn, the Pareto weight for a has actually gone down.

This suggests an alternative definition of the compensation, which is the following:

Definition 4 *The* actually compensating variation *is the function* Γ_2 *such that:*

$$\Gamma_2\left(\mathbf{r}, \mathbf{r}', x, \mathbf{z}\right) = \min_{x'} \left\{ \left(x' - x\right) \ \textit{subject to } \hat{v}^s\left(\mathbf{r}', x', \mathbf{z}\right) \geq \hat{v}^s(\mathbf{r}, x, \mathbf{z}),\ s = a, b \right\} \tag{30}$$

Thus $\Gamma_2(\mathbf{r}, \mathbf{r}', x, \mathbf{z})$ is the minimum amount to be paid to the household for each agent to be actually compensated for the reform, taking into account the intrahousehold allocation of additional income. This is illustrated in figure **??**. The actually compensating change moves the Pareto frontier out until b is no worse off. On the new frontier uu'' is the chosen allocation. Note, still, that while b is then exactly compensated for the reform, a gains strictly; the initial point uu lies strictly within the new frontier.

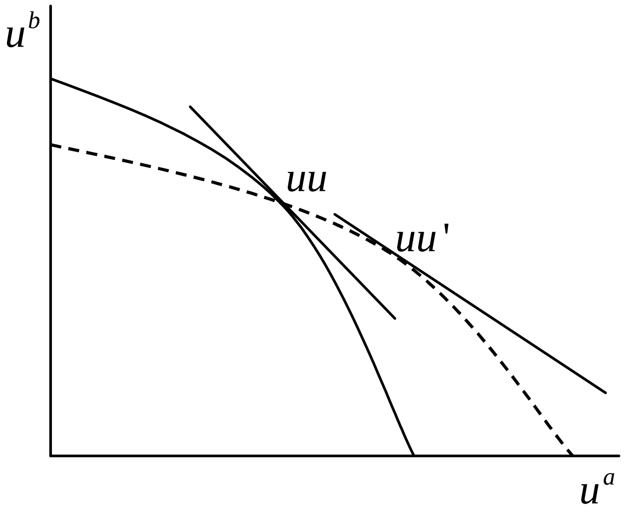

Figure 2: A potentially compensating variation.

Clearly, both concepts raise specific difficulties. The concept of potential compensation disregards actual decision processes, and ignores intrahousehold inequality. In a fully compensated household, the reform may worsen the situation of one of the members. This may have a social cost, at least if we accept that the actual intrahousehold decision process need not always be optimal from a normative, social viewpoint. On the other hand, the notion of actual compensation may lead to costly compensations, resulting in a bias in favor of the *status quo*. Moreover, it *de facto* rewards (marginal) unfairness, since the amount paid to the household has to be larger when most of the additional transfers goes to the dominant member. These issues are still largely open. We may simply make two remarks. First, these issues are inherent to any context in which the social planner cannot fully control intragroup redistribution; they are by no means specific to the collective approach, or for that matter to cooperative models. The obvious conclusion is that welfare economics can hardly do without a precise analysis of intrafamily decision processes.

Secondly, the notion of distribution factors suggests an additional direction for public intervention. Some of these factors can indeed be controlled

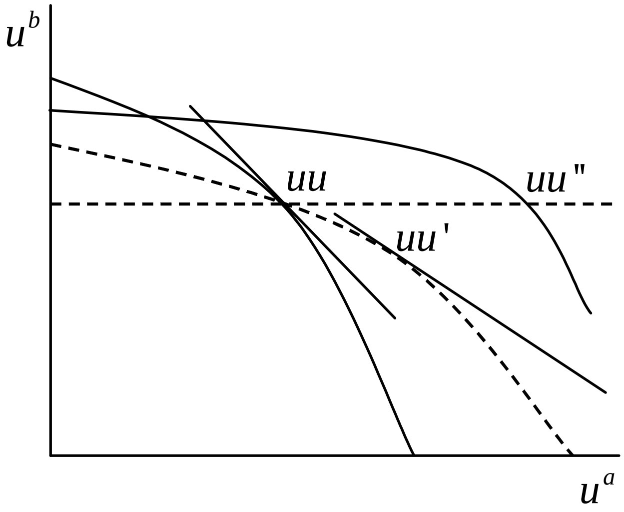

Figure 3: An actually compensating variation.

by the planner. For example, a benefit can be paid to the husband or to the wife, in cash or in kind. The benefit should then be designed taking into account the planner's ability to influence the decision process; technically, the maximization in (30) should be over x' *and* $\mathbf{z}$. For instance, several authors have suggested that a benefit aimed at improving the welfare of children should be paid to the mother, because such a shift may increase her weight in the decision process. Again, we may conclude that a theoretical and empirical analysis of intrahousehold allocation is a key step in any policy design.

3 The case of purely private consumptions

3.1 The sharing rule.

Although the Pareto weight captures very clearly our intuitive idea about power it turns out that if preferences are egotistic and we ignore public goods:

$$\begin{aligned} u^a\left(\mathbf{q}^a,\mathbf{q}^b\right) &= u^a\left(\mathbf{q}^a\right) \\ u^b\left(\mathbf{q}^a,\mathbf{q}^b\right) &= u^b\left(\mathbf{q}^b\right) \end{aligned} \tag{31}$$

then there is an equivalent concept which is easier to work with and to think about. It is a very familiar idea in economies with independent agents that if there are no externalities then any efficient outcome can be decentralised by a choice of prices and the distribution of income. This is the Second Fundamental theorem of Welfare Economics. In collective models we can exploit a similar idea. The efficiency assumption has a very simple and natural translation. With preferences of this kind, the economic interactions within the household are minimal: neither externalities, nor public goods are involved - agents essentially live side by side and consume independently.[5] Efficiency simply means that *for each agent*, the consumption bundle is optimal, in the sense that no other bundle could provide more utility at the same cost. In other words, take any particular (re)distribution of total income between members, and assume each member chooses his/her preferred consumption bundle subject to the constraint that the corresponding expenditures cannot exceed his/her share of total income. Then the resulting consumption will be Pareto efficient. Conversely, when preferences are quasi concave, any Pareto efficient allocation can be obtained in this way.

Suppose a household faces prices $\mathbf{p}$ and has decided on a level of total expenditure x. Let the resulting allocation be denoted $\left(\hat{\mathbf{q}}^a,\hat{\mathbf{q}}^b\right)$ so that $\mathbf{p}'\left(\hat{\mathbf{q}}^a+\hat{\mathbf{q}}^b\right)=x$. The decentralisation procedure is simple: each person is given a share of total expenditure and allowed to spend it on their own private goods, using their own private sub-utility function $u^s\left(\mathbf{q}^s\right)$. In what follows, let x^s denote s's total expenditures; then $x^a=\mathbf{p}'\hat{\mathbf{q}}^a$, $x^b=\mathbf{p}'\hat{\mathbf{q}}^b$, and $x^a+x^b=x$. Traditionally, a's part of total expenditures x^a is denoted ρ (so that $x^b=x-\rho$), and called the *sharing rule*.[6] Hence the following statement:

Proposition 5 *Define $\rho=\mathbf{p}'\hat{\mathbf{q}}^a$ so that $x-\rho=\mathbf{p}'\hat{\mathbf{q}}^b$. We have:*

- $\hat{\mathbf{q}}^a$ *solves*

$$\max u^a\left(\mathbf{q}^a\right) \text{ subject to } \mathbf{p}'\mathbf{q}^a=\rho \tag{32}$$

---

[5] This claim should be qualified. One could easily introduce additional, *non monetary* benefits of marriage (love, sex, companionship etc.). The crucial features is that individual preferences over consumption should be separable from these benefits.

[6] The terminology is not completely tied down with some authors referring to the fraction of expenditures going to A (that is, x^A/x) as the sharing rule.

- $\hat{\mathbf{q}}^b$ *solves*

$$\max u^b \left(\mathbf{q}^b\right) \text{ subject to } \mathbf{p}'\mathbf{q}^b = x-\rho \tag{33}$$

Conversely, for any ρ*, if* $\hat{\mathbf{q}}^a$ *and* $\hat{\mathbf{q}}^b$ *solve (32) and (33) then the allocation* $\left(\hat{\mathbf{q}}^a, \hat{\mathbf{q}}^b\right)$ *is Pareto efficient.*

The demands functions $\tilde{\mathbf{q}}^a$ and $\tilde{\mathbf{q}}^b$, *as functions of* $(\mathbf{p},\rho)$ and $(\mathbf{p},x-\rho)$, are conventional conditional demand functions and have all of the usual (Slutsky) properties.

When all commodities are privately consumed the decision process can be decomposed into two phases: a *sharing* phase in which agents determine the sharing rule and a *consumption* phase, in which agents allocate their share between the various commodities available. In this context, efficiency only relates to the second phase: whatever the sharing rule, the resulting allocation will be efficient provided that agents maximize their utility during the second phase. On the other hand, the collective part of the process (whether it entails bargaining, formal rules or others) takes part in the first stage.

Clearly, there exists a close connection (actually, a one-to-one increasing mapping) between a's share ρ and a's Pareto weight μ; both reflect a's power in the bargaining phase of the relationship. In particular, the sharing rule depends not only on prices and total expenditures but also on distribution factors.[7] An advantage of the sharing rule is that it is easy to interpret, unlike the Pareto weight. In particular, it is independent of the cardinal representation of individual utilities. For this reason, it is often more convenient to use the sharing rule as an indicator of the agent a's 'weight' in the decision process. Of course, this quality comes at a price: the sharing rule interpretation, as presented above, is valid only when all goods are privately consumed - although we will see an extension to public goods in the next section.

Finally, one should keep in mind that the functions $\tilde{\mathbf{q}}^a(\mathbf{p},\rho)$ and $\tilde{\mathbf{q}}^b(\mathbf{p},x-\rho)$, although 'structural' in the previous sense, cannot be observed, for two reasons. One is that, in general (and as in the general case before), one cannot change prices without changing the sharing rule as well; what can be observed, *at best*, are the functions $\hat{\mathbf{q}}^a(\mathbf{p},x,\mathbf{z})$ and $\hat{\mathbf{q}}^b(\mathbf{p},x,\mathbf{z})$, which are related

---

[7] The sharing rule depends on prices and income even if the Pareto weight is independent of the latter. Thus even in a unitary model with egotistic preferences we have a sharing rule and it depends on prices and total expenditure. However, the sharing rule cannot depend on distribution factors unless the Pareto weight does.

to the previous ones by the relationships:

$$\begin{aligned} \hat{\mathbf{q}}^a(\mathbf{p}, x, \mathbf{z}) &= \tilde{\mathbf{q}}^a(\mathbf{p}, \rho(\mathbf{p}, x, \mathbf{z})) \\ \hat{\mathbf{q}}^b(\mathbf{p}, x, \mathbf{z}) &= \tilde{\mathbf{q}}^b(\mathbf{p}, x - \rho(\mathbf{p}, x, \mathbf{z})) \end{aligned} \tag{34}$$

However, even these functions are in general unknown, because most of the time *the intrahousehold allocation of resources is not observed*. All expenditure surveys collect information about consumption that is aggregated at the household level; but who consumes what remains largely unknown, except, maybe, for some specific commodities (for example, expenditure surveys typically distinguish between male and female clothing). In general what we observe is the household demand which is equal to the sum of the individual demands:

$$\begin{aligned} \hat{\mathbf{q}}(\mathbf{p}, x, \mathbf{z}) &= \hat{\mathbf{q}}^a(\mathbf{p}, x, \mathbf{z}) + \hat{\mathbf{q}}^b(\mathbf{p}, x, \mathbf{z}) \\ &= \tilde{\mathbf{q}}^a(\mathbf{p}, \rho(\mathbf{p}, x, \mathbf{z})) + \tilde{\mathbf{q}}^b(\mathbf{p}, x - \rho(\mathbf{p}, x, \mathbf{z})) \end{aligned} \tag{35}$$

As we shall see below, one can often use this relationship to derive the properties of collective demand functions.

3.2 Caring preferences

Let us now consider the case of preferences of the 'caring' type, namely

$$\begin{aligned} U^a\left(\mathbf{q}^a, \mathbf{q}^b\right) &= u^a(\mathbf{q}^a) + \delta^a u^b\left(\mathbf{q}^b\right) \\ U^b\left(\mathbf{q}^a, \mathbf{q}^b\right) &= u^b\left(\mathbf{q}^b\right) + \delta^b u^a(\mathbf{q}^a) \end{aligned} \tag{36}$$

Here, the Welfare Theorems do not directly apply, since caring involves an externality component. Two points should however be remembered. First, any allocation that is Pareto efficient for caring preferences is also Pareto efficient for the egotistic preferences u^a and u^b. This implies that the first part of Proposition 5 still applies: whenever an allocation is efficient, it can be decentralized through a sharing rule. The converse, however, no longer holds in general. We know that some allocations may be efficient for egotistic preferences, but not so for caring ones. It follows that *only a subset* of possible sharing rules generate efficient allocations for caring preferences. For instance, a sharing rule such as $\rho = 0$ typically generates inefficient allocations, because a re-distribution of the resulting allocation in favor of a will in general increase both agents' welfare (if $\delta^b > 0$).

3.3 Indirect utilities

In the private good case, there exists a simple link between the collective indirect utilities defined above and the standard, individual indirect utilities. Denote the indirect utility corresponding to u^s (for $s = a, b$):

$$\begin{aligned} v^s\left(\mathbf{p}, x^s\right) &= \max u^s\left(\mathbf{q}\right) \\ \text{subject to } \mathbf{p}.\mathbf{q} &= x^s \end{aligned} \tag{37}$$

Thus $v^s\left(.\right)$ denotes the (maximum) utility level reached by s when facing prices $\mathbf{p}$ and consuming a total amount x^s. This is the standard, unitary concept, which makes no reference to the intrahousehold decision process. Now, in the case of private goods, the decision process is fully summarized by the sharing rule. It follows that:

$$\begin{aligned} V^a\left(\mathbf{p}, x, \mathbf{z}\right) &= v^a\left(\mathbf{p}, \rho\left(\mathbf{p}, x, \mathbf{z}\right)\right) \\ V^b\left(\mathbf{p}, x, \mathbf{z}\right) &= v^b\left(\mathbf{p}, x - \rho\left(\mathbf{p}, x, \mathbf{z}\right)\right) \end{aligned} \tag{38}$$

where V^s is the *collective indirect utility* of member s, according to the definition of the previous section. In particular, the bargaining phase of the decision process (deciding over the sharing rule) can readily be modeled using indirect utilities: whenever some ρ is chosen, a receives $v^a\left(\mathbf{p}, \rho\right)$ and b gets $v^b\left(\mathbf{p}, x - \rho\right)$. For instance, Nash bargaining with respective threat points T^a and T^b would solve:

$$\max_{\rho}\left[v^a\left(\mathbf{p}, \rho\right) - T^a\right]\left[v^b\left(\mathbf{p}, x - \rho\right) - T^b\right] \tag{39}$$

4 Application: labor supply with private consumption

4.1 The general setting

An example that has been widely analyzed in the literature concerns labor supply. In the most stripped down model without household production, labor supply is modelled as a trade off between leisure and consumption: people derive utility from leisure, but also from the consumption purchased with labor income. In a couple, however, an additional issue is the division of labor and of labor income: who works how much, and how is the resulting income distributed between members? As we now see, the collective approach provides a simple but powerful way of analyzing these questions.

Let l^s denote respectively member s's leisure (with $0 \leq l^s \leq 1$) and and q^s the consumption of a private Hicksian composite good whose price is set to unity. We start from the most general version of the model, in which member s's welfare can depend on his or her spouse's consumption and labor supply in a very general way, including for instance altruism, public consumption of leisure, positive or negative externalities, etc.. In this general framework, member s's preferences are represented by a utility function $U^s\left(l^a, q^a, l^b, q^b\right)$. Let w^a, w^b, y denote respectively real wage rates and household non-labour income. Finally, let z denote a K-vector of distribution factors. The efficiency assumption generates the program**:**

$$\begin{aligned} & \max_{\{l^a,l^b,q^a,q^b\}} \mu U^a + U^b \\ \text{subject to } q^a + q^b + w^a l^a + w^b l^b & \leq w^a + w^b + y \\ 0 & \leq l^s \leq 1, \quad s = a, b \end{aligned} \tag{40}$$

where μ is a function of $\left(w^a, w^b, y, \mathbf{z}\right)$, assumed continuously differentiable in its arguments.

In practically all empirical applications we observe only $q = q^a + q^b$. Consequently our statement of implications will involve only derivatives of q, l^a and l^b. In this general setting and assuming interior solutions, the collective model generates one set of testable restrictions, given by the following result:

Proposition 6 *Let* $\hat{l}^s(w^a, w^b, y, z)$, $s = a, b$ *be solutions to program (40). Then*

$$\frac{\partial \hat{l}^a / \partial z_k}{\partial \hat{l}^a / \partial z_1} = \frac{\partial \hat{l}^b / \partial z_k}{\partial \hat{l}^b / \partial z_1}, \quad \forall k = 2, ..., K. \tag{41}$$

This result is by no means surprising, since it is just a restatement of the proportionality conditions (15). The conditions are not sufficient, even in this general case, because of the SNR1 condition (12). Namely, one can readily check that the Slutsky matrix (dropping the equation for q because of adding up) takes the following form:

$$S = \begin{pmatrix} \frac{\partial \hat{l}^a}{\partial w^a} - \left(1 - \hat{l}^a\right) \frac{\partial \hat{l}^a}{\partial y} & \frac{\partial \hat{l}^a}{\partial w^b} - \left(1 - \hat{l}^b\right) \frac{\partial \hat{l}^a}{\partial y} \\ \frac{\partial \hat{l}^b}{\partial w^a} - \left(1 - \hat{l}^a\right) \frac{\partial \hat{l}^b}{\partial y} & \frac{\partial \hat{l}^b}{\partial w^b} - \left(1 - \hat{l}^b\right) \frac{\partial \hat{l}^b}{\partial y} \end{pmatrix}$$

As above, S must be the sum of a symmetric negative matrix and a matrix of rank one. With three commodities, the symmetry requirement is not restrictive: any 2×2 matrix can be written as the sum of a symmetric matrix and a matrix of rank one. Negativeness, however, has a bite; in practice, it requires that there exists at least one vector $\mathbf{w}$ such that $\mathbf{w}'\mathbf{S}\mathbf{w} < 0$. With distribution factors, the necessary and sufficient condition is actually slightly stronger, For $K = 1$, there must exist a vector $\mathbf{w}$ such that $S - \begin{pmatrix} \frac{\partial \hat{l}^1}{\partial z} & \frac{\partial \hat{l}^2}{\partial z} \end{pmatrix}' \mathbf{w}'$ is symmetric and negative. semi-definite

4.2 Egoistic preferences and private consumption

However, much stronger predictions obtain if we add some structure. One way to do that is to assume private consumption and egotistic (or caring) preferences, i.e. utilities of the form $u^s(l^s, q^s)$. Then there exists a sharing rule ρ, and efficiency is equivalent the two individual programs:[8]

$$\begin{aligned} & \max_{\{l^a,q^a\}} u^a(l^a, q^a) \\ \text{subject to } q^a + w^a l^a & \leq w^a + \rho \\ 0 & \leq l^a \leq 1 \end{aligned} \tag{42}$$

and

$$\begin{aligned} & \max_{\{l^b,q^b\}} u^b(l^b, q^b) \\ \text{subject to } q^b + w^b l^b & \leq w^b + (y - \rho) \\ 0 & \leq l^b \leq 1 \end{aligned} \tag{43}$$

Note that now ρ may be negative or larger than y, since one member may receive all non-labour income plus part of the spouse's labor income. Two remarks can be made at this point. First, ρ is the part of total non-labour income *allocated to* member a as an outcome of the decision process. This should be carefully distinguished from a's *contribution* to household non-labour income (although the latter may be a distribution factor if it influences the allocation process). That is, if non-labour income comes either from a (denoted y^a, representing for instance return on a's capital) or from b (denoted y^b, representing, say, a benefit paid exclusively to b), so that $y = y^a + y^b$, then a's part of total expenditures, denoted ρ, may depend (among other things) on y^a or on the ratio y^a/y - just as it may depend on any relevant distribution factor. But it is *not* equal to y^a in general. The

---

[8] In what follows, we shall assume for simplicity that only one distribution factor is available; if not, the argument is similar but additional, proportionality conditions must be introduced.

secondly point is that ρ may be an arbitrary function of wages, non-labour income and distribution factors. However, our assumptions imply that ρ *cannot* depend on the agents' total labor income, $w^s(1-l^s)$. Indeed, efficiency precludes a person's allocation to depend on an *endogenous* variable such as the labor supply of this person. The intuition is that such a link would act as a subsidy that would distort the price of leisure faced by the agents; see Basu (2006).

4.3 Collective labor supply

In turn, these programs shed light on various aspects of household labor supply. First, we have that

$$\begin{aligned} l^a &= \tilde{l}^a\left(w^a,\rho\right) && (44)\\ l^b &= \tilde{l}^b\left(w^b,y-\rho\right) && (45)\end{aligned}$$

where Λ^s denotes the Marshallian demand for leisure corresponding to u^s. The function Λ^s is structural (in the sense that it depends on preferences), but only l^s is observed. The first implication of this model is that the spouse's wage matters for an individual's demand for leisure, but only through its impact on the sharing rule; that is,, through an income effect. The same is true of non-labour income and of distribution factors:

$$\begin{aligned} \frac{\partial\hat{l}^a}{\partial w^b} &= \frac{\partial\tilde{l}^a}{\partial\rho}\frac{\partial\rho}{\partial w^b},\\ \frac{\partial\hat{l}^a}{\partial y} &= \frac{\partial\tilde{l}^a}{\partial\rho}\frac{\partial\rho}{\partial y}\\ \frac{\partial\hat{l}^a}{\partial z_k} &= \frac{\partial\tilde{l}^a}{\partial\rho}\frac{\partial\rho}{\partial z_k} && (46)\end{aligned}$$

The second equation can be rewritten in elasticity terms:

$$\frac{y}{\hat{l}^a}\frac{\partial\hat{l}^a}{\partial y}=\left(\frac{\rho}{\tilde{l}^a}\frac{\partial\tilde{l}^a}{\partial\rho}\right)\left(\frac{y}{\rho}\frac{\partial\rho}{\partial y}\right) \qquad (47)$$

Thus the income elasticity of a's *observed* demand for leisure is the product of two terms. The first is the *structural* income elasticity which characterizes a's preferences - what would be observed if a's fraction of total non-labour income could be independently monitored. The second term is the income elasticity of ρ, reflecting the change (in percentage) of a's allocation resulting from a given percentage change in household non-labour income. Hence if a member's allocation is elastic, then the elasticity of this person's demands for leisure, *as computed as the household level*, will exceed (in absolute value) the 'true' value (as observed for instance on singles, assuming that preferences are not changed by marriage). Conversely, if the allocation is inelastic (< 1), then her income elasticity will be found to be smaller than the 'true' value.

The same argument applies to own wage elasticities. From (44), we have that:

$$\frac{w^a}{\hat{l}^a}\frac{\partial \hat{l}^a}{\partial w^a} = \frac{w^a}{l^a}\frac{\partial \tilde{l}^a}{\partial w^a} + \left(\frac{\rho}{\tilde{l}^a}\frac{\partial \tilde{l}^a}{\partial \rho}\right)\left(\frac{w^a}{\rho}\frac{\partial \rho}{\partial w^a}\right) \tag{48}$$

Thus the own wage elasticity observed at the household level is the sum of two terms. The first is the 'structural' elasticity, corresponding to the agent's preferences; the second is the product of the person's structural income elasticity by the wage elasticity of the sharing rule. To discuss the sign of the latter, consider the consequences for intrahousehold allocation of an increase in a's wage. If leisure is a normal good, then the observed own wage elasticity (the left hand side) is smaller than the structural value (the first expression on the right hand side) if and only if ρ is increasing in w^a. This will be case if the wage increase dramatically improves a's bargaining position, so that a is able to keep all the direct gains *and* to extract in addition a larger fraction of household non-labour income. Most of the time, we expect the opposite; that is, part of a's gain in labor income is transferred to b, so that ρ is decreasing in w^a. Then the observed own wage elasticity (the left hand side) will be larger than the structural value.

The impact of distribution factors is in principle much easier to assess, because they leave the budget set unchanged and can only shift the distribution of power. Assuming that leisure is normal we have that if a change in a distribution factor favors member a, then a's share of household resources will increase which will reduce labor supply through a standard income effect.

This simple mechanism has been repeatedly tested, using distribution factors such as sex ratios and 'natural experiments' such as the legalization of divorce (in Ireland) or abortion (in the United States). Interestingly enough, all existing studies tend to confirm the theory. The effects are found to be significant and of the predicted sign; moreover, they are specific to married people and are typically not significant when singles are considered.

5 Public goods

5.1 Lindahl prices

We now consider a more general version of the model with egotistic preferences in which we allow for public goods. Hence individual utilities are of the form $u^s(\mathbf{q}^s, \mathbf{Q})$, $s = a, b$. While the general form of the Pareto program remains unchanged, its decentralization is trickier, because the welfare theorems do not apply in an economy with public goods.[9] One solution, which generalizes the previous intuitions, is to use *individual* (or 'Lindahl') prices. It relies on an old idea in public economics, namely that decisions regarding public commodities can be decentralized using agent-specific prices. In a sense, this is part of the standard duality between private and public consumptions. When a good is private, all agents face the same price and choose different quantities; with public goods, they all consume the same quantity but would be willing to pay different marginal prices for it.

A precise statement is the following:

Proposition 7 *For any* $(\mathbf{P}, \mathbf{p}, x, \mathbf{z})$*, assume that the consumption vector* $\left(\mathbf{Q}, \mathbf{q}^a, \mathbf{q}^b\right)$ *is efficient. Then there exists a* ρ *and* $2N$ *personal prices* $\mathbf{P}^a = \left(P_1^a, ..., P_N^a\right)$ *and* $\mathbf{P}^b = \left(P_1^b, ..., P_N^b\right)$*, with* $P_j^a + P_j^b = P_j$*,* $j = 1, ..., N$ *, such that* $(\mathbf{q}^a, \mathbf{Q})$ *solves*

$$\begin{gathered}\max u^a\left(\mathbf{q}^a, \mathbf{Q}\right) \\ \textit{subject to } \mathbf{p}'\mathbf{q}^a + \left(\mathbf{P}^a\right)'\mathbf{Q} = \rho\end{gathered} \tag{49}$$

and $\left(\mathbf{q}^b, \mathbf{Q}\right)$ *solves*

$$\begin{gathered}\max u^b\left(\mathbf{q}^b, \mathbf{Q}\right) \\ \textit{subject to } \mathbf{p}'\mathbf{q}^b + \left(\mathbf{P}^b\right)'\mathbf{Q} = y - \rho\end{gathered} \tag{50}$$

Note that both the function ρ and the personal prices $\mathbf{P}^a$ and $\mathbf{P}^b$ will in general depend on $(\mathbf{P}, \mathbf{p}, x, \mathbf{z})$. Also, it can readily be proved that efficiency obtains as soon as personal prices add up to market prices, as stated in the Proposition.

These programs correspond to a decentralization of the efficient allocation in the sense that each agent is faced with her own budget constraint,

---

[9] Private contributions to the public goods are ruled out, since they generate inefficient outcomes (see Chapter 3).

and maximizes her utility accordingly. There is however a clear difference with the private good case, in which all relevant information was readily available to each agent as soon as the sharing rule has be decided upon. Here, individuals need to know not only the 'resources' devoted to them, as described by ρ, but also the set of personal prices. Computing the personal prices is a difficult task, that is basically equivalent to solving for the efficient allocation; hence the 'decentralization' only obtains in a specific sense.[10]

Still, the Lindahl approach generates interesting insights on the outcome of the model. The first order conditions of (49) give:

$$P_j^a = \frac{\partial u^a / \partial Q_j}{\partial u^a / \partial q_i} p_i$$

The right hand side of this equation is often called a's *marginal willingness to pay* (or MWP) for commodity j; indeed, it is the maximum amount a would be willing to pay to acquire an additional unit of public good j, if the amount was to be withdrawn from a's consumption of private good i. Note that this amount does *not* depend on the private good at stake since the marginal utility of any private good divided by its price is equalised across private goods. Intuitively, thus, P_j^a increases with a's preference for the public good; the intuition of Lindahl prices is precisely that agents pay more, the greater their private valuation of the commodity.

Let us now compare the budget constraint the agent is facing in (49) with what the same agent would face if she was a single: $\mathbf{p}'\mathbf{q}^a + \mathbf{P}'\mathbf{Q} = y^a$, where y^a denotes a's income as single. An obvious difference is that the amount of resources has changed - from y^a to ρ; this is similar to the private goods case. However, another difference, which is specific to the public good case, is that *the relative prices of the public commodities have been changed*, from P_j/p_i to P_j^a/p_i. Since $P_j^a + P_j^b = P_j$ and $P_j^b > 0$, we have that $P_j^a < P_j$. Intuitively, the publicness of good j makes it less expensive relatively to any private good, preciselybecausetheotherspousewillalsocontributetothepurchaseofthepublicgood.

5.2 The conditional sharing rule.

An alternative approach relies on the notion of the *conditional sharing rule.* Again, let $\left(\mathbf{Q}, \mathbf{q}^a, \mathbf{q}^b\right)$ denote an efficient consumption vector, and let x^s denote s's (b's) total expenditures on private goods only; then $x^a = \mathbf{p}'\mathbf{q}^a$, $x^b = \mathbf{p}'\mathbf{q}^b$ and $x^a + x^b = x - \mathbf{P}'\mathbf{Q}$. Then:

---

[10] The literature on planning has developed several procedures through which information exchanges may lead to the determination of Lindahl prices.

Proposition 8 *For $s = a, b$, $\mathbf{q}^s$ solves:*

$$\max_{\mathbf{q}} u^s(\mathbf{q}, \mathbf{Q}) \text{ subject to } \mathbf{p}'\mathbf{q} = x^s \tag{51}$$

Note that, in the two programs above for $s = a, b$, individuals maximize over private consumptions taking public consumptions as given. The value x^a is called the *conditional sharing rule* precisely because its definition is conditional to the level of public expenditures. The proof is clear: if a could, through a different choice of her private consumption bundle, reach a higher utility level while spending the same amount, then the initial allocation had to be inefficient, a contradiction.

Again, the decision process can be interpreted as operating in two phases, although the precise definition of the phases differs from the private good case. Specifically, during the first phase agents determine *both* the level of public expenditures *and* the conditional sharing rule; then comes the *consumption* phase, when agents allocate their conditional share between the various private commodities available. It is important to note that in sharp contrast with the private good case, the existence of a conditional sharing rule is necessary for efficiency, but by no means sufficient. The reason for that is that, in general, efficiency introduces a strong relationship between the level of public expenditures and the conditional sharing rule. Broadly speaking, for any given level of public expenditures, most (actually, almost all) sharing rules would be incompatible with efficiency.

Before analyzing in more detail the first phase, it is useful to define a's indirect conditional utility $\tilde{v}^a$ as the value of program (51) above:

$$\begin{aligned} \tilde{v}^a(\mathbf{p}, x^a; \mathbf{Q}) &= \max_{\mathbf{q}^a} u^a(\mathbf{q}^a, \mathbf{Q}) \\ \text{subject to } \mathbf{p}'\mathbf{q}^a &= x^a \end{aligned} \tag{52}$$

In words, $\tilde{v}$ denotes the maximum utility a can ultimately reach given private prices and conditional on the outcomes $(x^a, \mathbf{Q})$ of the first period decision. We may now consider the first period. Efficiency leads to the following program:

$$\begin{gathered} \max_{x^a; x^b; \mathbf{Q}} \left\{ \mu \tilde{v}^a(\mathbf{p}, x^a; \mathbf{Q}) + \tilde{v}^b\left(\mathbf{p}, x^b; \mathbf{Q}\right) \right\} \\ \text{subject to} \quad x^a + x^b + \mathbf{P}'\mathbf{Q} = x \end{gathered} \tag{53}$$

In this formulation there is effectively a single private good: for given prices $\mathbf{p}$, the second phase only depends on the allocation $\left(x^a, x^b\right)$ and the choice of public consumption $\mathbf{Q}$. The first order conditions give:

$$\begin{aligned} \mu \frac{\partial \tilde{v}^a}{\partial x^a} &= \frac{\partial \tilde{v}^b}{\partial x^b} \\ \frac{\partial \tilde{v}^a / \partial Q_j}{\partial \tilde{v}^a / \partial x^a} + \frac{\partial \tilde{v}^b / \partial Q_j}{\partial \tilde{v}^b / \partial x^b} &= P_j, \ j = 1, ..., N \end{aligned} \tag{54}$$

The second set of conditions are often called the Bowen-Lindahl-Samuelson (BLS) conditions. The ratio $\frac{\partial \tilde{v}^a / \partial Q_j}{\partial \tilde{v}^a / \partial x^a}$ is exactly a's willingness to pay for public good j. To see this, note that the first order conditions of (51) imply that $\frac{\partial u^a}{\partial q_i^a} = \lambda^a p_i$, where λ^a is the Lagrange multiplier of a's budget constraint; and the envelope theorem applied to the definition of $\tilde{v}^a$ gives that $\frac{\partial \tilde{v}^a}{\partial x^a} = \lambda^a$, hence $\frac{\partial \tilde{v}^a}{\partial x^a} = \frac{1}{p_i} \frac{\partial u^a}{\partial q_i^a}$.Thus the conditions simply state that MWP's (or private prices) must add up to the market price of the public good, a result derived above. The BLS conditions (the second set of (54)) are necessary and sufficient for efficiency. The choice of a particular allocation on the Pareto frontier is driven by the first condition in (54).

The comparative statics of the model just described are somewhat complex, if only because the MWP for a particular commodity depend in an *a priori* arbitrary way on the quantities of the other public goods. However, a clearer picture obtains when there is only one public good. The efficiency conditions become, in that case:

$$\begin{aligned} \mu \frac{\partial \tilde{v}^a}{\partial x^a} &= \frac{\partial \tilde{v}^b}{\partial x^b} \\ MWP^a + MWP^b &= P \end{aligned} \tag{55}$$

where MWP^a denotes a's willingness to pay for the public good.

An important issue, here, is the impact of respective 'powers' (or the Pareto weight) on household decision making. Assume that the weight μ of agent a is increased. What will be the consequences on the distribution of private consumption between members and the household's demand for the public good? As it turns out, the model allows us to answer these questions (see Blundell, Chiappori and Meghir (2005)). Differentiating the first order conditions, we have:

$$\frac{\partial x^a}{\partial \mu} = \frac{1}{D} \left[\frac{\partial MWP^a}{\partial Q} + \frac{\partial MWP^b}{\partial Q} - \frac{\partial MWP^b}{\partial x^b} \right] \left(\frac{\partial \tilde{v}^a}{\partial x^a} + \frac{\partial \tilde{v}^b}{\partial x^b} \right) \tag{56}$$

$$\frac{\partial Q}{\partial \mu} = -\frac{1}{D} \left[\frac{\partial MWP^a}{\partial x^a} - \frac{\partial MWP^b}{\partial x^b} \right] \left(\frac{\partial \tilde{v}^a}{\partial x^a} + \frac{\partial \tilde{v}^b}{\partial x^b} \right) \tag{57}$$

where D is the determinant

$$D = \begin{vmatrix} \left(\mu \frac{\partial^2 \tilde{v}^a}{\partial x^a \partial Q} - (1-\mu)\frac{\partial^2 \tilde{v}^b}{\partial x^b \partial Q} - P\mu \frac{\partial^2 \tilde{v}^a}{(\partial x^a)^2}\right) & -\left((1-\mu)\frac{\partial^2 \tilde{v}^b}{\partial x^b \partial x^b} + \mu \frac{\partial^2 \tilde{v}^a}{(\partial x^a)^2}\right) \\ \frac{\partial MWP^a_j}{\partial Q} + \frac{\partial MWP^b_j}{\partial Q} - \frac{\partial MWP^a_j}{\partial x^a} P & \frac{\partial MWP^b_j}{\partial x^b} - \frac{\partial MWP^a_j}{\partial x^a} \end{vmatrix} \tag{58}$$

Assume, now, that preferences are such that the ‘goods’ x^a, x^b and Q are normal, in the following, usual sense: if an agent was purchasing these commodities as a single, the quantity purchased would increase with the agent's income (or total expenditures). Technically, this means that agent i's marginal willingness to pay, MWP^i, is increasing in x^i and decreasing in Q. Therefore, in the right hand side of (56), the expression in the square brackets is negative, while $\frac{\partial \tilde{v}^a}{\partial x^a}$ and $\frac{\partial \tilde{v}^b}{\partial x^b}$ are both positive. The sign of $\frac{\partial x^a}{\partial \mu}$ is thus the opposite of that of the determinant D.

The expression in (58) is quite complex, and there is little hope to derive the sign of D from it. But the problem can easily be circumvented. The argument goes as follows. Assume that the difference $DMWP \equiv \frac{\partial MWP^a}{\partial x^a} - \frac{\partial MWP^b}{\partial x^b}$ is, say, positive. Then $\frac{\partial Q}{\partial \mu}$ has the same sign as $\frac{\partial x^a}{\partial \mu}$ (i.e., minus the sign of D). But this sign cannot be negative, for that would mean that an increase in a's Pareto weight reduces a's public and private consumptions, hence a's utility. Therefore $\frac{\partial Q}{\partial \mu} \geq 0$ and $\frac{\partial x^a}{\partial \mu} \geq 0$.

If, on the other hand, the difference $DMWP$ is negative, the above analysis is valid for b; we conclude by the same token that an increase in μ, which *reduces* b's (relative) weight, decreases both Q and x^b (then x^a must increase).

We have thus found that a marginal change in a member's Pareto weight increases the household's expenditures on the public good if and only if the marginal willingness to pay of this member is *more sensitive to changes in his/her share* than that of the other member. Interestingly, it is not the *magnitude* of the MWP that matters, but its income elasticity.

It is useful to present an important application of this result. Consider the models discussed in Chapter ??? in which children's well being is modeled as a public good that enters the parents' utility. Assume that some policy measure may increase the relative weight of the wife within the household. It is often argued that children should benefit from such a change; the (somewhat hazy) intuition being that ‘mothers care more about children than do fathers’. What is the exact meaning of such a statement, and what exactly does it assume about preferences? The answer is given by the previous result. She ‘cares more’ means, in this context, that her MWP for children is more income-sensitive: should she receive an additional dollar to spend either on children or on her private consumption, she would spend a larger fraction of it on children than her husband would. As an illustration, assume individual preferences are Cobb-Douglas:

$$U^s(x^s, Q) = \log x^s + \delta^s \log Q \tag{59}$$

Straightforward computations show that the partial $\frac{\partial Q}{\partial \mu}$ is positive if and only if $\delta^a > \delta^b$ (which is indeed equivalent to $\frac{\partial MWP^a}{\partial x^a} > \frac{\partial MWP^b}{\partial x^b}$); note, however, that even with $\delta^a > \delta^b$ it may be the case that $MWP^a = \delta^a \frac{x^a}{Q} < MWP^b = \delta^b \frac{x^b}{Q}$ (particularly if x^b is large with respect to x^a).

6 Household production in collective models

6.1 General considerations.

In the presence of household production, purchased commodities are not directly consumed, but are used as inputs for some intrahousehold production process, the output of which are then consumed. We have already discussed home production in section **??** in chapter 3; here we focus on the novel aspects that arise in a collective model. Let $\mathbf{c}^s$ denote the vector of private consumption of the home produced commodity by s and let $\mathbf{C}$ denote public home produced goods. For the moment we shall ignore time inputs and let $\mathbf{q}$ denote the purchases of market goods that are used in home production. The Pareto program thus becomes

$$\begin{aligned} & \max \ \mu U^a\left(\mathbf{C}, \mathbf{c}^a, \mathbf{c}^b\right) + U^b\left(\mathbf{C}, \mathbf{c}^a, \mathbf{c}^b\right) \\ \text{subject to } \mathbf{F}\left(\mathbf{C}, \mathbf{c}^a + \mathbf{c}^b, \mathbf{q}\right) &= 0 \\ \mathbf{p}'\mathbf{q} &= x \end{aligned} \tag{60}$$

where $\mathbf{F}$ is the production function. As above, what is observed is the the household's demand function $\mathbf{q} = \hat{\mathbf{q}}\left(\mathbf{p}, x, \mathbf{z}\right)$.

When compared with the household production model in the unitary framework, (60) exhibits some original features. For instance, the outcome of the intrahousehold production process can be consumed either privately or publicly; the two situations will lead to different conclusions, in particular in terms of identification. On the other hand, two main issues - whether the goods produced within the household are marketable or not, and whether the output is observable - remain largely similar between the collective and the unitary frameworks.

As discussed in section **??** in chapter 3, when the outcome is observable, efficiency can directly be tested empirically. Indeed, a straightforward implication of efficiency is cost minimization: whatever the value of the output, it cannot be the case that the same value of output could be produced with a cheaper input combination. Udry (1996) provides a test of this sort on data from Burkina-Faso. Also, it is in general possible to directly estimate the production function; then one can refer to the standard, collective setting, using the methods presented above. Usually, however, the output of the intrahousehold production process is not observable. Still, some of the techniques described for models without home production can be extended to the case of production. For instance, distribution factor proportionality should still hold in that case; the basic intuition (distribution factors matter only through the one-dimensional Pareto weight μ) remains perfectly valid in Program (60). The same is true for the various versions of the SNR conditions, with and without distribution factors, which rely on the same ideas.

6.2 Housework.

Of particular interest are the various versions of the collective model with production involving labor supply. For simplicity, we present one version of the model, initially analyzed by Apps and Rees (1997) and Chiappori (1997), in which the two partners supply labor and consume two *private* consumption goods, one (denoted q and taken as numeraire) purchased on a market and the other (denoted c) produced domestically, according to some concave function $F\left(t^a, t^b\right)$, where t^s is member s's household work. Market and domestic labor supplies for person s, h^s and t^s, are assumed observed as functions of wages w^a, w^b, non-labour income y and a distribution factor $\mathbf{z}$.[11] For simplicity, we forget about the tax system and assume that budget sets are linear. Finally, we assume that preferences are 'egotistic', so that s's are represented by $u^s(q^s, c, l^s)$, where total time is normalized to unity so that

$$l^s + t^s + h^s = 1 \text{ for } s = a, b \tag{61}$$

6.3 Marketable production

Let us first assume that good c can be bought and sold on the market; let c^M denote the quantity sold (or bought if negative) on the market and p its

---

[11] With several distribution factors, the approach is similar. In addition, one can perform the proportionality tests described above.

market price; the household takes this price as given. Then total production of the good is $c = c^a + c^b + c^M$; if $c^M > 0$ then the household produces more than it consumes $(c^a + c^b)$. The budget constraint at the household level is given by :

$$q^a + q^b = w^a l^a + w^b l^b + y + pc^M. \tag{62}$$

Consider the household as a small economy, defined by preferences u^a and u^b and by *two* production constraints - namely, the production of the household good (here $c = F\left(t^a, t^b\right)$) *and* the budget constraint. By the welfare theorems, any Pareto efficient allocation can be decentralized as a market equilibrium. On the production side, the second constraint (the budget constraint) implies that the *intrahousehold* prices of the consumption goods q and c and the leisures l^a and l^b are proportional to $\left(1, p, w^a, w^b\right)$; we can normalize the proportionality factor to be one, and keep $\left(1, p, w^a, w^b\right)$ as intrahousehold prices as well. Moreover, t^a and t^b must solve:

$$\max_{\left(t^a, t^b\right)} pF\left(t^a, t^b\right) - w^a t^a - w^b t^b \tag{63}$$

This is the well-known separability principle, according to which the production side is fully determined by profit maximization, irrespective of individual preferences. This provides testable restrictions: any variable that does not affect the production side of the household (but only, say, preferences or the decision process) should not be relevant for the determination of t^a and t^b. Moreover, first order conditions give

$$\frac{\partial F}{\partial t^s}\left(t^a, t^b\right) = \frac{w_s}{p}, \quad s = a, b \tag{64}$$

If F is strictly concave, these relations can be inverted to give:

$$t^s = f^s\left(\frac{w^a}{p}, \frac{w^b}{p}\right), \quad s = a, b \tag{65}$$

Knowing the $f^s(.)$ functions is strictly equivalent to knowing F. The relationships (65) can in principle be econometrically estimated, leading to a complete characterization of the production side.

Considering, now, the demand side, one sees that the household's total 'potential' income is

$$w^a\left(T - t^a\right) + w^b\left(T - t^b\right) + y + pc \tag{66}$$

which has to be split between the members and spent on individual leisures and consumptions of the two goods. In particular, for any efficient allocation, respective welfare levels are fully determined by the sharing rule. Hence we are back to the case studied in Section 4, up to one difference - namely, there are two consumption goods instead of one.

6.4 Non-marketable production

The other polar case obtains when no market for the domestic good exists (then $c^M = 0$). One can still define a price p for the domestic good, equal to the marginal rate of substitution between the domestic and the market goods for each of the members (the MRS are equalized across members as a consequence of the efficiency assumption). The difference, however, is that p is now *endogenous* to the household preferences and decision process. As above, this implies that the separability property no longer applies; the price p has to be estimated as well. As discussed by Chiappori (1997), identifiability does not obtain in general; however, it can still be achieved under additional assumptions.

However, a much stronger result obtains when the produced good is publicly consumed. Blundell, Chiappori and Meghir (2005) consider a model which is formally similar to the previous one, except that the second commodity is public and its production requires labor and some specific input, Q. Technically, individual utilities take the form $u^s(q^s, C, l^s)$, and the production constraint is $C = F\left(Q, t^a, t^b\right)$. A natural (but not exclusive) interpretation of C is in terms of children's welfare, which enters both utilities and is 'produced' from parental time and children expenditures Q. Blundell, Chiappori and Meghir show that strong testable restrictions are generated. Moreover, the structure (that is, utilities and the Pareto weights) are identifiable from labor supplies (both domestic and on the market) and children's expenditures, provided that one distribution factor (at least) is available.

References

[1] Apps, P. F. and R. Rees (1997). "Collective labor supply and household production." *Journal of Political Economy* 105(1): 178-190.

[2] Basu, K. (2006). "Gender and say: A model of household behaviour with endogenously determined balance of power." *Economic Journal* 116(511): 558-580.

[3] Blundell, R., P. A. Chiappori, et al. (2005). "Collective labor supply with children." *Journal of Political Economy* 113(6): 1277-1306.

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

[5] Chiappori, P. A. (1997). "Introducing household production in collective models of labor supply." *Journal of Political Economy* 105(1): 191-209.

[6] Chiappori, P. A. and I. Ekeland (2004). chapter 5.

[7] Chiappori, P. A. and I. Ekeland (2005). Chapter 6.

[8] Chiappori, P. A. and I. Ekeland (2006). Chapter 6.

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