The Economics of the Family

Chapter 4: 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

July 2010

1 Collective demand functions: a general characterization

1.1 The collective household utility function.

The basic features of the collective model have been described in the previous chapter. We now derive the testable implications for observable demand functions. 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

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

where μ may be a function of $(\mathbf{P}, \mathbf{p}, x, \mathbf{z})$ where z is a vector of distribution factors. We shall always assume that $\mu(.)$ 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 $\mathbf{g} = \left(\mathbf{q}^a + \mathbf{q}^b, \mathbf{Q}\right)$ denoting the quantities consumed by the household and $\mathbf{r} = (\mathbf{p}, \mathbf{P})$ denoting the corresponding price vector. Then the household's behavior is described by the maximization of $u^h(\mathbf{g},\mu)$ under the household budget constraint $\mathbf{r}'\mathbf{g} = x$.

1.2 Structural and observable demand.

The household's program is:

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

which generates *collective demand functions*, $\tilde{\mathbf{g}}(\mathbf{r}, x, \mu)$. It is important to emphasize that this program is *not* equivalent to standard utility maximization (the unitary model) because u^h varies with μ, which in turn depends on prices, income and distribution factors. Yet, for any *fixed* μ, $\tilde{\mathbf{g}}(., \mu)$ is a standard demand function. From standard consumer theory, we therefore know that it satisfies Slutsky symmetry and negativeness[1]. This property is crucial in what follows; it can be exploited in a more formal way.

[1] In all that follows we abbreviate 'negative semi-definite' to 'negative' in the interest of brevity.

Define the generic Slutsky matrix element of $\tilde{\mathbf{g}}(., \mu)$, always holding μ constant, as:

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

and denote its Slutsky matrix by $\Sigma = [\sigma_{ij}]_{i,j}$. We then have that Σ is symmetric and negative. Rearranging (3), we have the standard interpretation of a Slutsky matrix; namely, 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 ($\tilde{g}_j \frac{\partial \tilde{g}_i}{\partial x}$). The intuition is that a marginal increase in the price of any good j 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.

Although the analysis of $\tilde{\mathbf{g}}(\mathbf{r}, x, \mu)$, holding μ constant is conceptually useful, it is crucial to realize that $\tilde{\mathbf{g}}$ *cannot be observed directly*; indeed, such an observation would require changing 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 notational economy, dropped the distribution factors.[2] 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.

[2] 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.

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 two goods. We start with prices and income and the demand at point I. We then change prices so that good 1 is cheaper. The substitution effect is given by the move from I to II and the income effect is II to III. The collective effect associated with the change in μ is represented by the final term in (5) which is shown as the move from III to IV.

1.3 The Slutsky matrix for collective demands.

Using the *observable* functions $\hat{\mathbf{g}}(.)$, we can define the *observable* or *quasi-Slutsky* matrix $S = [s_{ij}]_{i,j}$ by its general term:

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

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} with associated matrix Σ. 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:

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

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 , and an additional matrix R. The latter is the product of a column vector

$(D_\mu \tilde{g})$ and a row vector $(\mathbf{v}')$. Note that such 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 $R.\mathbf{w} = 0$. Finally, this analysis and the homogeneity assumption on $\mu(.)$ yields that the necessary conditions for the collective model demands are the *generalized Slutsky conditions*:

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

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

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

(Browning and Chiappori (1998)). We denote the third property *SNR1.*

One can readily see that these conditions 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, in particular, when either μ is constant (so that $\mathbf{v} = 0$) or when $\tilde{\mathbf{g}}$ does not depend on μ (so that $D_\mu \tilde{g} = 0$). The latter case corresponds to the two partners having the same cardinal preferences; $u^b(\mathbf{g}) = k_0 + k_1 u^a(\mathbf{g})$ with $k_1 > 0$. 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.[3]

The result has a simple, geometric intuition which can be seen in 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 (locally) sufficient for the existence of a collective model. Chiappori and Ekeland (2006) show that any 'smooth'[4] demand function satisfying the three properties above

[3] 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.

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

(homogeneity, adding-up and 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 generate moves along the Pareto frontier (from II to III in Figure **??** in chapter 3). 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:

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

Because the same $\mu\left(.\right)$ function appears in all goods the collective model yields *cross-equation restrictions.* To see this, consider the consequences of a marginal change in distribution factor z_k on the collective demand for commodity i:

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

Comparing the effect of different distribution factors, say z_k and z_l, we find that (assuming $\partial g_i/\partial z_l \neq 0$):

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

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[5]

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

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

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

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 empirically testable (subject to having at least two distribution factors); 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.

1.5 Larger households

The analysis developed above can be extended to larger households. Suppose there are T agents in the household. We continue to assume efficiency so that the collective household utility function is defined as:

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

where the vector $\boldsymbol{\mu} = (\mu_1, ..., \mu_T)$ of Pareto weights is 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, $\tilde{\mathbf{g}}\left(\mathbf{r}, x, \mu_1, ..., \mu_{T-1}\right)$ 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, defined by:

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

Similar computations to the two person case yield:

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

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 generalized Slutsky conditions for a T person household are given by:

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

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:

$$\begin{aligned} E\left(\mathbf{r}, \bar{u}^a, \bar{u}^b\right) &= \min_{\mathbf{q}^a \mathbf{q}^b, \mathbf{Q}} \mathbf{r}' \left(\mathbf{q}^a + \mathbf{q}^b, \mathbf{Q}\right) \\ \text{subject to } u^s\left(\mathbf{q}^a, \mathbf{q}^b, \mathbf{Q}\right) &\geq \bar{u}^s, \; s = a, b. \end{aligned} \tag{23}$$

The collective expenditure function depends on prices and on two utility levels $\left(\bar{u}^a, \bar{u}^b\right)$; it represents the minimum level of expenditures needed at these prices to achieve these utilities. One can then define the *compensated collective demand function*, $\breve{\mathbf{g}}\left(\mathbf{r}, \bar{u}^a, \bar{u}^b\right)$, 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. The basic one is the following. Consider the 'primal' model stated in Chapter 3:

$$\max_{\mathbf{Q},\mathbf{q}^a,\mathbf{q}^b} u^b\left(\mathbf{Q},\mathbf{q}^a,\mathbf{q}^b\right)$$
$$\text{subject to } \mathbf{r}'\left(\mathbf{q}^a+\mathbf{q}^b,\mathbf{Q}\right) \leq x$$
$$\text{and } u^a\left(\mathbf{Q},\mathbf{q}^a,\mathbf{q}^b\right) \geq \bar{u}^a \tag{24}$$

The two programs (23) and (24) are closely related. Indeed, let $\left(\bar{\mathbf{Q}},\bar{\mathbf{q}}^a,\bar{\mathbf{q}}^b\right)$ denote the solution to (24) and let $\bar{u}^b = u^b\left(\bar{\mathbf{Q}},\bar{\mathbf{q}}^a,\bar{\mathbf{q}}^b\right)$. Then:

$$E\left(\mathbf{r},\bar{u}^a,\bar{u}^b\right) = x$$

and $\left(\bar{\mathbf{Q}},\bar{\mathbf{q}}^a,\bar{\mathbf{q}}^b\right)$ solves (23). Conversely, if $\left(\bar{\mathbf{Q}},\bar{\mathbf{q}}^a,\bar{\mathbf{q}}^b\right)$ denotes the solution to (23) for some $\left(\bar{u}^a,\bar{u}^b\right)$, then for $x = E\left(\mathbf{r},\bar{u}^a,\bar{u}^b\right)$ we have that

$$u^b\left(\bar{\mathbf{Q}},\bar{\mathbf{q}}^a,\bar{\mathbf{q}}^b\right) = \bar{u}^b$$

and $\left(\bar{\mathbf{Q}},\bar{\mathbf{q}}^a,\bar{\mathbf{q}}^b\right)$ solves (24) . The intuition is simply that if a particular bundle maximizes b's utility subject to constraints on a's utility and total expenditures - this is program (24) - then one cannot reach the same utilities at a lower total cost than this bundle: if that was possible, the difference in costs could be used to buy extra public commodities and increase both members' utilities, a contradiction. Conversely, if a bundle minimizes total cost for two given utility levels - thereby solving Program (23) - then one cannot increase b's utility without either reducing a's utility or spending more.

The notion of collective expenditure function - and the duality property just described - is a direct generalization of the standard expenditure function of consumer theory; the only difference is that, now, there are two utility levels that should be reached. Many results follow that generalize standard theorems of consumer theory. In particular if denotes the gradient of $E(.)$ with respect to $\mathbf{r}$ by $\nabla_r E$ (that is, the vector of partial derivatives $\partial E/\partial r_j$), then:

Proposition 1 $\breve{\mathbf{g}}\left(\mathbf{r},u^a,u^b\right) = \nabla_r E\left(\mathbf{r},u^a,u^b\right)$

The result is a consequence of the envelope theorem applied to program (23).

In the case of egotistic preferences of the form $u^s(\mathbf{q}^s, \mathbf{Q})$, we have further results. Define the compensated demand for public goods by $\breve{\mathbf{Q}}\left(\mathbf{p}, \mathbf{P}, u^a, u^b\right)$. Then:

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

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

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

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

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

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*

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

therefore

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

■

2.2 Indirect utilities

We can also define indirect utility functions. Consider first the program

$$\max_{\mathbf{q}^a, \mathbf{q}^b, \mathbf{Q}} \mu u^a(\mathbf{q}^a, \mathbf{Q}) + u^b\left(\mathbf{q}^b, \mathbf{Q}\right)$$

$$\text{subject to } \mathbf{r}'\left(\mathbf{q}^a + \mathbf{q}^b, \mathbf{Q}\right) = x \tag{29}$$

Let $\left(\mathbf{q}^{*a}, \mathbf{q}^{*b}, \mathbf{Q}^*\right)$ denote its solution. Then the function $v^{\ s}$, defined for $s = a, b$ by:

$$v^s\left(\mathbf{r}, x, \mu\right) = u^s\left(\mathbf{q}^{*s}, \mathbf{Q}^*\right)$$

is the direct equivalent, in the collective setting, of the indirect utility concept in standard consumer theory. In particular, v^s only depends on *preferences*, not on the *decision process*; technically, v^s is a function of the Pareto weight μ, and a change in the decision process would result in the same function v^s being applied to a different μ.

A second, and more important definition is obtained by plugging the particular Pareto weight adopted by the household into the previous definition. In this case, the *collective indirect utility of a member* is the level of utility ultimately reached by this member as a function of prices and income and distribution factors. Formally, if the decision process is characterized by a function $\mu\left(\mathbf{r},x,\mathbf{z}\right)$, the collective indirect utility of member s is defined by:

$$V^s\left(\mathbf{r},x,\mathbf{z}\right) = v^s\left(\mathbf{r}, x, \mu\left(\mathbf{r},x,\mathbf{z}\right)\right)$$

The definition of s's collective indirect utility depends not only on s'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.

Also, a key remark, here, is that if one is interested in welfare analysis, then the collective indirect utility is the appropriate concept. Indeed, it preserves the basic interpretation of standard, indirect utilities in consumer theory - namely, it characterizes each agent's final welfare once all aspects of the decision process have been taken into account.

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:

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

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:*

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

In words, consider a household in which, before the reform, total income is x and member s's utility is $u^s = V^s\left(\mathbf{r}, x, \mathbf{z}\right)$. 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 sense that the additional income thus measured could, *if allocated appropriately within the household*, enable both members to reach their pre-reform utility levels. That is, the income $x + \Gamma\left(\mathbf{r}, \mathbf{r}', x, \mathbf{z}\right)$ has the property that the utilities $\left(V^a\left(\mathbf{r}, x, \mathbf{z}\right), V^b\left(\mathbf{r}, x, \mathbf{z}\right)\right)$ belong to the Pareto frontier at prices $\mathbf{r}'$. What is *not* guaranteed, however, is that the point $\left(V^a\left(\mathbf{r}, x, \mathbf{z}\right), V^b\left(\mathbf{r}, x, \mathbf{z}\right)\right)$ 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 $uu = \left(V^a\left(\mathbf{r}, x, \mathbf{z}\right), V^b\left(\mathbf{r}, x, \mathbf{z}\right)\right)$. 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:*

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

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 3. 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.

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 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 there is an equivalent concept which is easier to work with and to think about, if preferences are egotistic and we ignore public goods:

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

It is a very familiar idea in convex economies with independent agents that if there are no externalities, then any efficient outcome can be decentralized by a choice of prices and the (re)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.[6] 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*.[7] Hence the following statement:

[6] This claim should be qualified. One could easily introduce additional, *non-monetary* benefits of marriage (love, sex, companionship etc.).

[7] 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.

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

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

In other words, 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 consumption 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.

In the current setting a sharing rule can be defined for any decision process (one can always consider the outcome and compute the amount spent on private goods for member a). However, Proposition (5) is satisfied (that is , the outcome maximizes a's utility under a's budget constraint) if and only if the process is efficient. Clearly, there exists a close connection (actually, if u^a and u^b are strictly concave, a one-to-one, increasing mapping) between a's share ρ and a's Pareto weight; both reflect a's power in the bargaining phase of the relationship. This implies that the sharing rule depends not only on prices and total expenditures but also on distribution factors.[8]

An advantage of the sharing rule is that, unlike the Pareto weight, it is easy to interpret. 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: any change in, say, a distribution factor that increases ρ makes a

[8] 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.

better off. Of course, this quality comes at a price: the sharing rule interpretation, as presented above, *is valid only when all goods are privately consumed*. We will see in Section 5 to what extent it can be generalized to public goods

Finally, one should keep in mind that the individual demand functions $\tilde{\mathbf{q}}^a(\mathbf{p},\rho)$ and $\tilde{\mathbf{q}}^b(\mathbf{p},x-\rho)$ are 'structural' in the previous sense and cannot be observed, for two reasons. One is that, in general, 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 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 purchases is not observed*. Expenditure surveys invariably collect information about expenditures that are 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 $\rho \simeq 0$ typically generates inefficient allocations since a redistribution of the resulting allocation in favor of a may increase both agents' welfare.

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 $s = a, b$):

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

Thus $v^s(.)$ 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:

$$V^a(\mathbf{p}, x, \mathbf{z}) = v^a(\mathbf{p}, \rho(\mathbf{p}, x, \mathbf{z})) \tag{38}$$

$$V^b(\mathbf{p}, x, \mathbf{z}) = v^b(\mathbf{p}, x - \rho(\mathbf{p}, x, \mathbf{z})) \tag{39}$$

where V^s is the *collective indirect utility* of member s, according to the definition of the previous section. In particular, the first phase of the decision process (deciding over the sharing rule) can readily be modeled using indirect utilities: whenever some ρ is chosen, a receives $v^a(\mathbf{p}, \rho)$ and b gets $v^b(\mathbf{p}, x - \rho)$. The program would therefore become:

$$\max_{\rho} \mu v^a(\mathbf{p}, \rho) + v^b(\mathbf{p}, x - \rho) \tag{40}$$

More specific processes can also be considered. For instance, Nash bargaining with respective threat points T^a and T^b would solve:

$$\max_{\rho} \left[v^a(\mathbf{p}, \rho) - T^a\right] \left[v^b(\mathbf{p}, x - \rho) - T^b\right] \tag{41}$$

It is important to note that, in general, many different structures (that is , individual preferences and a sharing rule) generate the same collective indirect utilities V^a, V^b. Indeed, for any given pair $\left(V^a, V^b\right)$, let $\left(v^a, v^b, \rho\right)$ be such that (38) and (39) are satisfied, and assume that v^a and v^b are strictly increasing and strictly quasi concave. Pick an arbitrary function $\phi(\mathbf{p})$, and define $\left(v^a_\varepsilon, v^b_\varepsilon, \rho_\varepsilon\right)$ by:

$$\begin{aligned} v^a_\varepsilon(\mathbf{p},r) &= v^a(\mathbf{p}, r-\varepsilon\phi(\mathbf{p})) \\ v^b_\varepsilon(\mathbf{p},r) &= v^b(\mathbf{p}, r+\varepsilon\phi(\mathbf{p})) \\ \rho_\varepsilon(\mathbf{p},x,\mathbf{z}) &= \rho(\mathbf{p},x,\mathbf{z}) + \varepsilon\phi(\mathbf{p}) \end{aligned}$$

then one can readily check that

$$V^a(\mathbf{p},x,\mathbf{z}) = v^a_\varepsilon(\mathbf{p}, \rho_\varepsilon(\mathbf{p},x,\mathbf{z})) \tag{42}$$

$$V^b(\mathbf{p},x,\mathbf{z}) = v^b_\varepsilon(\mathbf{p}, x-\rho_\varepsilon(\mathbf{p},x,\mathbf{z})) \tag{43}$$

In other words, the structures $\left(v^a, v^b, \rho\right)$ and $\left(v^a_\varepsilon, v^b_\varepsilon, \rho_\varepsilon\right)$, although different, generate the same collective indirect utilities. It follows that the welfare conclusions reached by the two structures are always identical. For instance, if a given reform is found to increase his welfare and decrease her's when the evaluation is made using the first structure, using the second instead will lead to the same conclusion. We say that different structures that generate the same collective indirect utilities are *welfare equivalent.*

The notion of welfare equivalence plays an important role, notably in the discussion of identification in Chapter 5. In many situations, welfare equivalent structures are hard to empirically distinguish; in some cases, only the collective indirect utilities can actually be recovered. The key remark is that as far are welfare judgment are concerned, identifying collective indirect utilities is sufficient.

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 member s's leisure (with $0 \leq l^s \leq 1$) 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-labor income. Finally, let z denote a K-vector of distribution factors. The efficiency assumption generates the program:

$$\max_{\{l^a, l^b, q^a, q^b\}} \mu U^a + U^b$$

$$\begin{aligned} \text{subject to } q^a + q^b + w^a l^a + w^b l^b &\leq w^a + w^b + y \\ 0 &\leq l^s \leq 1, \quad s = a, b \end{aligned} \tag{44}$$

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 (44). Then*

$$\frac{\partial \hat{l}^a / \partial z_k}{\partial \hat{l}^a / \partial z_1} = \frac{\partial \hat{l}^b / \partial z_k}{\partial \hat{l}^b / \partial z_1}, \quad \forall k = 2, ..., K. \tag{45}$$

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 $S - \begin{pmatrix} \frac{\partial \hat{l}^a}{\partial z} & \frac{\partial \hat{l}^b}{\partial z} \end{pmatrix}' \mathbf{w}'$ is symmetric and negative.

4.2 Egoistic preferences and private consumption

Much stronger predictions obtain if we add some structure. One way to do that is to assume private consumption and egotistic (or caring) preferences, that is utilities of the form $u^s(l^s, q^s)$. Then there exists a sharing rule ρ, and efficiency is equivalent to the two individual programs:[9]

$$\max_{\{l^a, q^a\}} u^a(l^a, q^a)$$

$$\begin{aligned} &\text{subject to } q^a + w^a l^a \leq w^a + \rho \\ 0 \quad \leq \quad &l^a \leq 1 \end{aligned} \tag{46}$$

and

$$\max_{\{l^b, q^b\}} u^b(l^b, q^b)$$

$$\begin{aligned} &\text{subject to } q^b + w^b l^b \leq w^b + (y - \rho) \\ 0 \quad \leq \quad &l^b \leq 1 \end{aligned} \tag{47}$$

Note that now ρ may be negative or larger than y, since one member may receive all non-labor 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-labor income *allocated to* member a as an outcome of the decision process. This should be carefully distinguished from a's *contribution* to household non-labor income (although the latter may be a distribution factor if it influences the allocation process). That is, if non-labor income comes either from a (denoted y^a, representing, for instance, the 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

[9] 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.

(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 second point is that ρ may be an arbitrary function of wages, non-labor 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, as in Basu's (2006) model of inefficient bargaining described in the previous Chapter.

4.3 Collective labor supply

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

$$\begin{aligned} l^a &= \tilde{l}^a(w^a, \rho) && (48) \\ l^b &= \tilde{l}^b\left(w^b, y-\rho\right) && (49) \end{aligned}$$

where $\tilde{l}^s$ denotes the Marshallian demand for leisure corresponding to u^s. The function $\tilde{l}^s$ 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-labor income and of distribution factors:

$$\begin{aligned} \frac{\partial \hat{l}^a}{\partial w^b} &= \frac{\partial \tilde{l}^a}{\partial \rho}\frac{\partial \rho}{\partial w^b}, \\ \frac{\partial \hat{l}^a}{\partial y} &= \frac{\partial \tilde{l}^a}{\partial \rho}\frac{\partial \rho}{\partial y} \\ \frac{\partial \hat{l}^a}{\partial z_k} &= \frac{\partial \tilde{l}^a}{\partial \rho}\frac{\partial \rho}{\partial z_k} && (50) \end{aligned}$$

The second equation can be rewritten in elasticity terms:

$$\frac{y}{\hat{l}^a}\frac{\partial \hat{l}^a}{\partial y} = \left(\frac{\rho}{\tilde{l}^a}\frac{\partial \tilde{l}^a}{\partial \rho}\right)\left(\frac{y}{\rho}\frac{\partial \rho}{\partial y}\right) \quad (51)$$

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-labor 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-labor 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 (48), we have that:

$$\frac{w^a}{\hat{l}^a}\frac{\partial \hat{l}^a}{\partial w^a} = \frac{w^a}{l^a}\frac{\partial \tilde{l}^a}{\partial w^a} + \left(\frac{\rho}{\tilde{l}^a}\frac{\partial \tilde{l}^a}{\partial \rho}\right)\left(\frac{w^a}{\rho}\frac{\partial \rho}{\partial w^a}\right) \tag{52}$$

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-labor 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 (see the discussion in the next Chapter).

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.[10] 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; see, for example, Mas-Colell, Whinston and Green (1995). 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* $\left(\hat{\mathbf{Q}}, \hat{\mathbf{q}}^a, \hat{\mathbf{q}}^b\right)$ *is efficient. Then there exists a* ρ *and* $2N$ *personal prices* $\mathbf{P}^a = (P_1^a, ..., P_N^a)$ *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* $\left(\hat{\mathbf{q}}^a, \hat{\mathbf{Q}}\right)$ *solves:*

$$\begin{gathered} \max u^a(\mathbf{q}^a, \mathbf{Q}) \\ \textit{subject to } \mathbf{p}'\mathbf{q}^a + (\mathbf{P}^a)'\mathbf{Q} = \rho \end{gathered} \tag{53}$$

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

$$\begin{gathered} \max u^b\left(\mathbf{q}^b, \mathbf{Q}\right) \\ \textit{subject to } \mathbf{p}'\mathbf{q}^b + \left(\mathbf{P}^b\right)'\mathbf{Q} = y - \rho \end{gathered} \tag{54}$$

Note that both the function ρ and the personal prices $\mathbf{P}^a$ and $\mathbf{P}^b$ will in general depend on $(\mathbf{P}, \mathbf{p}, x, \mathbf{z})$.

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

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

and maximizes their 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 been decided upon. Here, individuals need to know not only the 'resources' devoted to them, as described by ρ, but also their 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.[11]

Still, the Lindahl approach generates interesting insights on the outcome of the model. Assuming an interior solution, the first order conditions of (53) give:

$$P_j^a = \frac{\partial u^a / \partial Q_j}{\partial u^a / \partial q_i} p_i \tag{55}$$

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 equalized across private goods. Intuitively P_j^a increases with a's preference for the public good; the intuition of Lindahl prices is precisely that agents with a higher private valuation of the public good should pay more for it. This is required for an efficient allocation of the family income between alternative uses.

Let us now compare the budget constraint the agent is facing in (53) 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, precisely because the other spouse will also contribute to the purchase of the public good.

5.2 The conditional sharing rule.

An alternative approach relies on the notion of the *conditional sharing rule*. Again, let $\left(\hat{\mathbf{Q}}, \hat{\mathbf{q}}^a, \hat{\mathbf{q}}^b\right)$ denote an efficient consumption vector. The total

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

expenditure of a and b on private goods only are $x^a = \mathbf{p}'\hat{\mathbf{q}}^a$ and $x^b = \mathbf{p}'\hat{\mathbf{q}}^b$. This implies that $x^a + x^b = x - \mathbf{P}'\hat{\mathbf{Q}}$. Then:

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

$$\max_{\mathbf{q}} u^s \left(\mathbf{q}, \hat{\mathbf{Q}}\right) \text{ subject to } \mathbf{p}'\mathbf{q} = x^s \tag{56}$$

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 (56) above:

$$\begin{aligned} \tilde{v}^a (\mathbf{p}, x^a; \mathbf{Q}) &= \max_{\mathbf{q}^a} u^a (\mathbf{q}^a, \mathbf{Q}) \\ \text{subject to } \mathbf{p}'\mathbf{q}^a &= x^a \end{aligned} \tag{57}$$

That is, $\tilde{v}^a$ denotes the maximum utility a can ultimately reach given private prices and conditional on the outcomes $(x^a, \mathbf{Q})$ of the first phase decision. We may now consider the first phase, which determines the public consumption, Q, and the disposable income allocated to each spouse, $\left(x^a, x^b\right)$. Efficiency leads to the following program:

$$\begin{gathered} \max_{x^a; x^b; \mathbf{Q}} \left\{\mu \tilde{v}^a (\mathbf{p}, x^a; \mathbf{Q}) + \tilde{v}^b \left(\mathbf{p}, x^b; \mathbf{Q}\right)\right\} \\ \text{subject to} \quad x^a + x^b + \mathbf{P}'\mathbf{Q} = x \end{gathered} \tag{58}$$

The first order conditions give:

$$
\begin{aligned}
\mu \frac{\partial \tilde{v}^a}{\partial x^a} &= \frac{\partial \tilde{v}^b}{\partial x^b} \\
\frac{\partial \tilde{v}^a / \partial Q_j}{\partial \tilde{v}^a / \partial x^a} + \frac{\partial \tilde{v}^b / \partial Q_j}{\partial \tilde{v}^b / \partial x^b} &= P_j, \ j = 1, ..., N
\end{aligned} \tag{59}
$$

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 (56) 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, as argued above. The BLS conditions (the second set of (59)) are necessary and sufficient for efficiency. The choice of a particular allocation on the Pareto frontier is driven by the first condition in (59).

As an application, consider the model of collective labor supply proposed by Donni (2007), who assumes individual preferences of the form:

$$
u^s(1 - h^s, Q),
$$

where Q is a Hicksian good which represents public consumption and h^s is the labor supply of person s. Under this hypothesis, and taking into account the property of homogeneity, labor supplies can be written as:

$$
h^s = h^s \left(\frac{w^s}{\pi_s(y, w^a, w^b, \boldsymbol{z})}, \frac{\rho^s(y, w^a, w^b, \boldsymbol{z})}{\pi_s(y, w^a, w^b, \boldsymbol{z})} \right),
$$

where

$$
\pi_s(y, w^a, w^b, \boldsymbol{z}) = \frac{w^s h^s + \rho^s(y, w^a, w^b, \boldsymbol{z})}{y + w^a h^a + w^b h^b}
$$

denotes member s's Lindahl price for the public good. In this context, Donni shows that the utility functions are identified, up to a positive transformation, from individual labor supplies.

5.3 Application: labor supply, female empowerment and expenditures on public good

While the previous concepts may seem somewhat esoteric, they have important practical applications. For instance, a widely discussed issue in development economics and welfare policy in general is the impact of intrahousehold redistribution on the structure of household consumption, and in particular on household demand for public goods. The notion of 'public goods' should be understood here in a very general sense - any expenditure that benefits both partners. A typical and normatively important example is expenditures on children, at least if we assume that both parents care about the well being of their children. The crucial question, then, is the following: if a given policy 'empowers' women, in the sense that it increases their weight in the household decision process, what will be the impact on household expenditures on children? For instance, by paying a given benefit to the wife instead of the husband, can we expect children health or education to be improved? A large and growing body of empirical evidence suggests that such redistributive effects do exist and can actually be quite large, at least in some countries. As an instance, Duflo (2003), studying elderly benefits in South Africa, concludes that the same transfer has a drastically different impact on the health of female grandchildren depending on whether it is paid to the grandmother or the grandfather.

The collective framework provides a framework for studying these effects. The basic intuition is that while the amount received has a direct impact on the household's budget constraint, the gender of the recipient does not. It can only affect the respective Pareto weights; as such, it is a perfect example of a *distribution factor*. We therefore want to investigate the impact of distribution factors (or equivalently of exogenous changes in the Pareto weights) on household demand. Two questions are of particular interest. First, is it possible to predict, from the knowledge of preferences, which public consumptions will increase when the wife's weight raises? Second, is it always the case that female empowerment also results in more spending on the wife's *private* consumption - or could it be the case that she puts so much emphasis on public consumption that her private consumption actually declines when she has more power?

To investigate these issues, we start with a very simple example. Assume individual preferences are Cobb-Douglas:

$$u^s(\mathbf{q}^s, \mathbf{Q}) = \sum_k \alpha_k^s \log q_k^s + \sum_j \delta_j^s \log Q_j \tag{60}$$

where the coefficients are positive and normalized by $\sum_k \alpha_k^s + \sum_j \delta_j^s = 1$. As above, let μ denote a's Pareto weight. Prices are normalized to 1, so that the budget constraint is simply

$$\sum_k \left(q_k^a + q_k^b\right) + \sum_j Q_j = x$$

Straightforward computations give household demands:

$$
\begin{aligned}
q_k^a &= \frac{\mu \alpha_k^a}{1+\mu} x \\
q_k^b &= \frac{\alpha_k^b}{1+\mu} x \\
Q_j &= \frac{\mu \delta_j^a + \delta_j^b}{1+\mu} x
\end{aligned}
$$

and the following conclusions follow:

1. The private consumptions of a are all increasing in μ
2. The private consumptions of b are all decreasing in μ
3. Since
$$\frac{\partial Q_j}{\partial \mu} = \frac{\delta_j^a - \delta_j^b}{(1+\mu)^2} x$$
household consumption in public commodity j increases if and only if a 'cares more' about that commodity than b does, in the sense that $\delta_j^a > \delta_j^b$.

As above, it is natural to interpret these results in terms of marginal willingness to pay. These are given for any public good j by:

$$MWP_j^s = \delta_j^s \frac{x^s}{Q_j}, \; s = a, b$$

where $x^s = \sum_k q_k^s$ is the conditional sharing rule of member s. Interestingly enough, the condition $\delta_j^a > \delta_j^b$ is *not* equivalent to her MWP being larger than his; rather, it implies that

$$\frac{\partial MWP_j^a}{\partial x^b} > \frac{\partial MWP_j^b}{\partial x^a} \tag{61}$$

In words, the MWP of a must be more *income sensitive* than that of b. Still, it may be the case that $MWP^a < MWP^b$ (particularly if x^b is large with respect to x^a): the absolute magnitude of the respective MWP plays no role in the result.

The interpretation of these findings is quite intuitive. First, one may think of the wife's empowerment (as resulting from an increase in μ) in purely economic terms: she now receives a higher fraction of household resources. With Cobb-Douglas preferences, all commodities are normal, therefore more income always results in more consumption for her; conversely, his share has been reduced and he consumes less. Regarding public goods, however, things are more complex, because a transfer from the husband to the wife typically increases her MWP for each public good but reduces his. The question, here, is whether her increase is sufficient to compensate his reduction - which is exactly what is implied by equation (61). If the condition is satisfied, the impact of the change over total MWP for the public good is positive, and consumption grows; in the opposite situation, it is reduced.

The previous results, natural as they sound, are still dependent on the very specific functional form chosen for utilities. Whether they extend to non-homothetic preferences, for instance, is not clear. In full generality, the comparative statics of the model just described are somewhat complex, if only because, unlike the Cobb-Douglas case, the MWP for a particular commodity depends 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, a case considered by Blundell, Chiappori and Meghir (2005). They show that if preferences are such that both private expenditures and the public good are normal (in the usual sense that an increase in income would raise the corresponding, *individual* demands for these goods), then a marginal improvement 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. Again, it is not the magnitude of the MWP's that matters, but their income sensitivity. Moreover, the private consumptions of the beneficiary member are always increased.

Coming back to the initial motivation, consider the model discussed in Chapter 3 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 be spent either on children or on her private consumption, she would spend a larger fraction of it on children than her husband would.

6 Household production in the collective model

Becker (1965) put forward a generalized approach of consumption and time use in which final consumption is produced within the household by intermediate goods purchased in the market and personal time withdrawn from market work. Although house production is important for singles, it is particular relevant for married (or cohabiting) couples. Household production generates several of the gains from marriage that we mentioned in chapter 3, including increasing returns, specialization and sharing (home produced) public goods. At a global level, household production represents, according to several estimates, up to 20% of total production in developed countries, although it is usually not explicitly taken into account in aggregate measures such as GDP, and much more in developing economies. At the household level, domestic production represents a significant fraction of resources (and especially of time) used and consumed. Finally, at an individual level, utility depends on leisure, which can be defined as time not spent working either at home or on the market (although such a definition raises delicate problems) and also on the consumption of internally produced commodity.

The analysis of household production raises several important issues. One is the choice of the commodities produced at home and their quantity. In many cases, a trade-off exists between home production and market trade. For instance, I can clean my apartment or hire a cleaning person; and in the opposite direction, the vegetables I grow in my garden can be consumed internally by my family or sold on the market.[12] The commodity is then said to be *marketable*. Alternatively, some commodities have to be at least partly internally 'produced'; for instance, a nanny cannot, in many cases, be a perfect substitute for parental care. Another issue is whether and how these decisions depend on the partners' respective 'powers'. Is it the case, for instance, that the allocation of work by each spouses to the domestic production process reflects the bargaining positions of the spouses - or is it exclusively determined by the production technology?

Finally, these issues must be analyzed in an equilibrium context, in which many key factors have drastically evolved over time. In particular, the division of labor within households has changed as married women have dramatically increased their labor force participation. Becker's framework allows one to conceptualize the distinct roles of technological advance in home production and in industrial production in explaining the observed changes in

[12] This issue is particularly important in development economics, since a majority of the population of a developing economy typically work in agriculture, often producing marketable commodities at the household level.

allocation of time. There is extensive research that applies household production approach and tries to sort out the roles of technological advance and changes in norms that have made this revolution possible (Greenwood *et al*, 2005, Fernandez, 2007). Mulligan and Rubinstein (2007) emphasize the role of higher rewards for ability (reflected in the general increase in wage inequality) in drawing married women of high ability into the labor market. See also Albanesi and Olivetti (2009), who emphasize the role of medical progress in child feeding that enabled women to stay out of home.[13]

Another crucial determinant of the time spent on household production is its opportunity cost, which is directly related to the wage the person could receive by working on the market. Over the last decades, a striking phenomenon is the global increase in female education, an evolution that has deeply modified the trade-off between domestic and market work by raising female market wages. Of course, education is not exogenous; it is the outcome of an investment decision based on future (expected) returns, therefore on (among other things) the fraction of time that individuals expect to spend working on the market. In other words, education and current wages affect current decisions regarding household production, but are themselves the outcomes of past expectations about future domestic work. The general equilibrium aspects will be left for the second part of the book; here, we concentrate on a providing a conceptual framework for analyzing the respective impacts of wages, technology and power on domestic production.

6.1 The basic model

We have already discussed home production in section 2 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 time being, we ignore time inputs and let $\mathbf{q}$ denote the purchases of market goods that are used in home production. Assuming for the moment that household commodities are not marketable, the Pareto program thus becomes

$$\max \mu U^a\left(\mathbf{C}, \mathbf{c}^a, \mathbf{c}^b\right) + U^b\left(\mathbf{C}, \mathbf{c}^a, \mathbf{c}^b\right) \tag{62}$$

[13] Another application is De Vries (1994, 2008) who applied this framework to identify an "industrious revolution", characterized by an increased production of marketable goods within households, which "preceded and prepared the way for the Industrial Revolution".

subject to

$$\begin{aligned}\mathbf{F}\left(\mathbf{C}, \mathbf{c}^a + \mathbf{c}^b, \mathbf{q}\right) &= 0 \\ \mathbf{p}'\mathbf{q} &= x \end{aligned} \tag{63}$$

where $\mathbf{F}$ is the production function. As above, what is observed is the the household's demand function $\mathbf{q} = \hat{\mathbf{q}}(\mathbf{p}, x, \mathbf{z})$. Note that the model implicitly assumes that *all* commodities are input for household production. This is without loss of generality: if commodity i is directly consumed, the corresponding row of the production equation simple reads $c_i^a + c_i^b = q_i$ for private consumption, or $C_i = q_i$ if the consumption if public.

When compared with the household production model in the unitary framework, (62) 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.

6.2 Domestic production and time use

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.[14] 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-labor income y and a distribution factor z.[15] For simplicity, we ignore the tax system and assume that budget sets are linear;[16] similarly, we exclude joint production.[17] Finally, we assume that

[14] The model can easily be generalized by adding other inputs to the production process; the main conclusions below would not change.

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

[16] For a comprehensive analysis of taxation with household production, the reader is referred to Apps and Rees (2009).

[17] See Pollak and Wachter (1975), and Apps and Rees (2009) for a general presentation.

preferences are 'egoistic', so that s's are represented by $U^s(q^s, c^s, l^s)$, where l^s denotes leisure and total time is normalized to unity so that

$$l^s + t^s + h^s = 1 \text{ for } s = a, b \tag{64}$$

When the domestic good is not marketable, the previous model therefore becomes:

$$\max \mu U^a\left(q^a, c^a, l^a\right) + U^b\left(q^b, c^b, l^b\right) \tag{65}$$

subject to

$$\begin{aligned} c^a + c^b &= F\left(t^a, t^b\right) & (66) \\ q^a + q^b &= y + w_a h^a + w_b h^b & (67) \end{aligned}$$

and the time constraint (64).[18] Conversely, if the commodity is marketable - that is , if good c can be bought and sold on a market, we let c^M denote the quantity sold (or bought if negative) on the market and p its market price, which the household takes 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 $(c^a + c^b)$ and sell the difference, if $c^M < 0$ the household produces only a fraction of the amount it consumes and purchases the rest. The production equation is now:

$$c^a + c^b + c^M = F\left(t^a, t^b\right)$$

and the budget constraint at the household level becomes:

$$q^a + q^b = w^a h^a + w^b h^b + y + pc^M. \tag{68}$$

In our analysis of household production models, we shall first consider the benchmark situation in which both spouses are working outside the family, and their working time is flexible enough to allow for marginal variations. Then the opportunity cost of a person's time is determined by the person's wage, which is taken as given for the family decision process. We later consider 'corner' solutions, in which one spouse works exclusively at home.

[18] Note that utility depends only on consumption and leisure and that, by assumption, time spent at work either at home or in the market do not enter utility directly.

6.2.1 Marketable production

Cost minimization Let us first assume that good c is marketable. In this context, efficiency has an immediate implication, namely profit maximization. Specifically, t^a and t^b must solve:

$$\max_{\left(t^a,t^b\right)} pF\left(t^a,t^b\right) - w^a t^a - w^b t^b \tag{69}$$

implying the first order conditions:

$$\frac{\partial F}{\partial t^s}\left(t^a,t^b\right) = \frac{w_s}{p}, \quad s = a, b \tag{70}$$

The economic interpretation of these equations is clear. The opportunity cost of an additional unit of time spent on domestic production is the person's wage. If this is not equated to the marginal productivity of domestic labor, efficiency is violated. For instance, if this marginal productivity is smaller than the wage, then the person should spend less time working at home and more working for a wage, keeping total leisure constant. Intrahousehold production would decline, but household income would increase by more than the amount needed to purchase the missing production on the relevant market. To put it differently, the condition reflects cost minimization; if it not satisfied, then the household could achieve the same level of leisure and domestic consumption while saving money that could be used to purchase more of the consumption goods - clearly an inefficient outcome.

The same argument can be presented in a more formal way. 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 second welfare theorem, 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. Then market equilibrium requires profit maximization, which does not depend on individual preferences. This is the well-known separation principle, according to which the production side is fully determined by profit maximization, irrespective of individual preferences.

Choosing domestic work The first order conditions of the profit maximization program give

$$\frac{\partial F}{\partial t^s}\left(t^a, t^b\right) = \frac{w^s}{p}, \quad s = a, b \tag{71}$$

If F is strictly concave (that is , if the domestic technology exhibits decreasing returns to scale), these relations can be inverted to give:

$$t^s = f^s\left(\frac{w^a}{p}, \frac{w^b}{p}\right), \quad s = a, b \tag{72}$$

Knowing the $f^s(.)$ functions is strictly equivalent to knowing F. The relationships (72) can in principle be econometrically estimated, leading to a complete characterization of the production side. It is important to note that, in this logic, *the time spent by each spouse on domestic production is totally determined by 'technological' consideration*: it depends only on wages and on the household production function F, but neither on preferences nor on 'power' (as measured by Pareto weights). The model predicts, for instance, that when a change in a distribution factor redistributes power in favor of the wife (say, a benefit that used to be paid to the husband is now paid to the wife), the result will be a different consumption pattern (as discussed above, the household now consumes more of the commodities preferred by the wife), but the times spent on domestic production by the husband and the wife remain unchanged. On the contrary, an exogenous increase in female wage reduces her domestic labor; the impact on his domestic work then depends on the domestic production technology (that is , are male and female work complement or substitutes?).

It should be stressed that the marketability assumption is demanding. Strictly speaking, it requires that households can freely buy *or sell* the domestic good. Selling the domestic good is natural in some contexts (for example, agricultural production in developing countries), but less so in others (many people clean their own house but would not think of selling their cleaning services to a third party). If domestic goods can only be purchased but not sold, our analysis still applies whenever wages and technology are such that they always consume more than what they produce - that is , the household as a positive net demand of the domestic good. However, some households may reach a corner solution, in which the market purchase of domestic goods is nil, and the normalized marginal productivity of a person's domestic work exceeds the person's wage. In practice, this is equivalent to the domestic good not being marketable, a case we consider below.

Finally, the model above assumes that all forms of labor are equally costly - that is , that the subjective disutility of one hour of labor is the same, whether it is spent working in a factory or taking care of children. This assumption, however, can readily be relaxed. One may posit, for instance, that for some activities (say domestic work), one hour of work 'costs' to spouse s only a fraction α^s of one hour of leisure (intuitively, the remaining fraction $(1-\alpha^s)$ *is* leisure). Under this extension, the time constraint (64) should be replaced with:

$$l^s + \alpha^s t^s + h^s = 1 \text{ for } s = a, b \tag{73}$$

and the first order conditions become:

$$\frac{\partial F}{\partial t^s}\left(t^a, t^b\right) = \alpha^s \frac{w_s}{p}, \quad s = a, b \tag{74}$$

In words, the opportunity cost of domestic work should be adjusted for the associated amenities. Note, however, that the same logic applies; that is, the time spent by each spouse on domestic production is fully determined by wages, technology and the individual preferences captured here by the amenity parameter α^s. However, they do *not* depend on the power of the spouses as measured by μ.

The demand side The separability principle implies that the demand side is totally divorced from production. Indeed, the household's total 'potential' income is

$$Y = w^a\left(1-t^a\right) + w^b\left(1-t^b\right) + y + pc \tag{75}$$

This potential income has to be split between the members and spent on individual leisures and consumptions of the two goods. Since all commodities are private, efficiency is equivalent to the existence of a sharing rule. As above, thus, there exists two functions $\rho^a\left(w^a, w^b, y, p\right)$ and $\rho^b\left(w^a, w^b, y, p\right)$, with $\rho^a + \rho^b = Y$, such that each member s solves:

$$\max U^s\left(q^s, c^s, l^s\right)$$

under the member-specific budget constraint

$$q^s + pc^s + w^s l^s = \rho^s$$

At this stage, we are back to the standard collective model of labor supply.

6.2.2 Non-marketable production

The other polar case obtains when no market for the domestic good exists (then $c^M = 0$). Then we are back to maximizing (65) under the constraints (66), (67) and (64). 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 model - that is , it is determined by the maximization program.

A particularly interesting case obtains when the domestic production function exhibits constant returns to scale (CRS). Then:

$$F\left(t^a, t^b\right) = t^b \Phi\left(\frac{t^a}{t^b}\right) \tag{76}$$

for some function Φ. The first order conditions (FOC) imply that:

$$\frac{\partial F/\partial t^a}{\partial F/\partial t^b} = \frac{w^a}{w^b},$$

which give in this case:

$$\Phi\left(\frac{t^a}{t^b}\right) - \frac{t^a}{t^b}\Phi'\left(\frac{t^a}{t^b}\right) = \frac{w^a}{w^b},$$

This relationship, which is a direct consequence of the efficiency assumption, pins down the ratio $\frac{t^a}{t^b}$ to be some function $\phi\left(\frac{w^a}{p}\right)$. In other words, it is now the case that *the ratio of male to female domestic work depends only on wages and household production technology* - a natural consequence of cost minimization. On the other hand, the *scale* of production - that is , the quantity eventually produced - is indeterminate from the production perspective; it depends on preferences and the decision process. We conclude that preferences and power determine the total quantity of household goods produced; however, *conditional on that quantity*, the particular combination of male and female time is determined by respective wages and the production technology, and does not depend on preferences or power.[19]

[19] Pollak and Wachter (1975) discuss the roles of constant returns to scale and joint production. They show that with joint production (i.e., activities that generate more than one final good), it is generally impossible to separate household technology from preferences, even under a constant return to scale technology.

The (household-specific) price of the domestic good can readily be recovered. Indeed, an interior solution under constant returns require zero profits, therefore it must be the case that:

$$p = \frac{w^a \left(t^a/t^b\right) + w^b}{\Phi \left(t^a/t^b\right)}$$

Again, this price depends only on wages and on the technology. It is household specific in the sense that two households with different wages will price the household good differently, even if they have access to the same domestic technology. However, for given wages and domestic technology, it depends neither on preferences nor on respective powers. Finally, the separation result still holds. That is, each member's decision can be modeled as if they were maximizing their own utility under the member specific budget constraint defined by a sharing rule; this mechanism determines all the components of consumption, including The only difference with the marketable case is that p is no longer a market price; instead, it is determined by the wages and the technological constraints.

6.2.3 Power and domestic work

While the previous conclusions are not really surprising, at least from a general equilibrium perspective - they basically illustrate standard results in welfare economics - their implications can be somewhat unexpected. Consider, for instance, a change in Pareto weights that benefits women - say, through the impact of a distribution factor - while wages and incomes are unaffected. As discussed in the previous subsection, a first consequence is that the structure of consumption will change; intuitively, the household will now consume more of the commodities that the wife 'likes more'. If, as it is often argued, women generally care more about the goods that are domestically produced (child care being a primary example), the total consumption of these commodities should increase. If the commodity is marketable and initially (partly) purchased on the market, the result will be higher market purchases of these goods, with no impact on domestic labor by the partners. In all other cases, domestic labor will increase, and the distribution of the additional effort between spouses is completely driven by the technology. For instance, under a standard, Cobb-Douglas production function, inputs are complements; at constant prices (here wages), more production requires increasing both inputs. We conclude that more power to the wife may actually imply more domestic work for *both* spouses. Note, however, that because a transfer of income to the wife does not affect her time input into home production, the income effect will induce her to reallocate the remaining time so that her market work should decline and her leisure increase. This conclusion should be contrasted with the impact of an increase in the wife's market wage, which *always* affects her domestic labor supply. When the commodity is marketable, her domestic work is always reduced. In the alternative situation, her domestic work decreases *with respect to her husband's*, but the absolute impact also depends on the structure of consumption - especially if her Pareto weight is boosted by her higher wage.

6.2.4 Extensions

Public goods In the previous analysis, the internally produced commodity was privately consumed. What if, instead, the commodity is public within the household - as it is the case for childcare, for instance? Interestingly, not much is changed, because the separation principle still applies. If the commodity, although public within the household, is marketable, then its production is driven by profit maximization; the only change is on the demand side, where the decision process can no longer be decentralized using a sharing rule. Even in the non marketable case, the logic of cost minimization prevails. In particular, under constant returns to scale, it still the case that the level of production is determined by preferences and the decision process, while for any given level the time allocation of domestic work between spouses stems from technological considerations.[20]

Specialization Another special (but empirically relevant) case obtains when one of the spouses - say b - does not enter the labor market, and specializes instead in home production. This happens when, for the chosen allocation of time and consumption, b's potential wage, w^b, is smaller than both b's marginal productivity in household production and b's marginal rate of substitution between leisure and consumption. In words: the marginal hour can indifferently be spent in leisure or household production, and both uses dominate market work.[21]

The situation, here, is more complex, because the opportunity cost of labor for b is no longer exogenously given; instead, it is now endogenous to

[20] The reader is referred to Blundell, Chiappori and Meghir (2005) for a more detailed investigation.

[21] Technically, this result is true at the marginal level only in the absence of non convexities. In the presence of fixed costs of work or constraints on the number of hours worked, the same constraint must be redefined at a more global level.

the program. Still, if we keep the assumption of constant return to scale domestic technology, some of the previous conclusions remain valid. Indeed, in the marketable case, efficiency in a's allocation of time still requires that:

$$\frac{\partial F}{\partial t^a}\left(t^a, t^b\right) = \frac{w^a}{p}$$

while the CRS condition (76) implies that

$$\frac{\partial F}{\partial t^a}\left(t^a, t^b\right) = \Phi'\left(\frac{t^a}{t^b}\right)$$

It follows that, again, the ratio t^a/t^b is pinned down by technological constraints - namely, it must be such that

$$\Phi'\left(\frac{t^a}{t^b}\right) = \frac{w^a}{p}$$

In words: the volume of domestic production is now determined by preferences, but the distribution (between spouses) of effort needed to produce that amount is fixed by the technology.

Finally, in the case of specialization into the production of a non marketable good, both the price of the domestic good and b's opportunity cost of labor are endogenous. Then all aspects of household production are potentially affected by the distribution of power within the couple.

6.3 Empirical issues

To what extent can the previous analysis generate testable restrictions? Note first that, as discussed in section 2 of 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 (62). The same is true for the various versions of the SNR conditions, with and without distribution factors, which rely on the same ideas.

Moreover, if time use data are available, then the previous models generate several, testable restrictions regarding the impact of wages, income and power on domestic production. If we consider the benchmark case of CRS technology, the basic prediction is that the *proportion* of total domestic time spent by each member only depends on wages and the technology. Therefore 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 the ratio t^a/t^b. On the other hand, changes in wages do affect the ratio; as expected, a (proportionally) higher female wage reduces the ratio of her domestic work to his.

Regarding identification, note first that if the internally produced commodity is marketable (as will often be the case for, say, agricultural production in developing countries), then conditions (72) above can in principle be econometrically estimated, leading to a complete characterization of the production side. In the opposite case, however, 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.

Finally, 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] Albanesi, Stefania and Claudia Olivetti, 'Home Production, Market Production and the Gender Wage Gap: Incentives and Expectations', **Review of Economic Dynamics**, 12 (2009), 80–107.

[2] Apps, Patricia F. and Ray Rees, 'Collective Labor Supply and Household Production', **Journal of Political Economy,** 105 (1997), 178-90.

[3] Apps, Patricia F. and Ray Rees, **Public Economics and the Household**, (Cambridge, Cambridge University Press, 2009).

[4] Basu, Kaushik, 'Gender and Say: A Model of Household Behaviour With Endogenously Determined Balance of Power', **Economic Journal**, 116 (2006), 558-580.

[5] Becker, Gary S., 'A Theory of the Allocation of Time', **The Economic Journal**, 75 (1965), 493-517.

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

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

[8] Chiappori, Pierre-André, 'Introducing Household Production in Collective Models of Labor Supply', **Journal of Political Economy,** 105 (1997), 191-209.

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

[10] De Vries, Jan, 'The Industrial Revolution and the Industrious Revolution', **The Journal of Economic History,** 54 (1994), 249-270.

[11] De Vries, Jan, **The Industrious Revolution: Consumer Behavior and the Household Economy, 1650 to the Present,** (Cambridge: Cambridge University Press, 2008).

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

[13] Fernandez, Raquel, 'Alfred Marshall Lecture Women, Work, and Culture', **Journal of the European Economic Association,** 5 (2007), 305-332.

[14] Greenwood, Jeremy, Seshadri, Ananth and Mehmet Yorukoglu, 'Engines of Liberation', **Review of Economic Studies,** 72 (2005), 109-133.

[15] Mas-Colell, Andreu, Whinston, Michael D. and Jerry R. Green, **Microeconomic Theory,** (New York: Oxford University Press, 1995).

[16] Mulligan, Casey B. and Yona Rubinstein, 'Selection, Investment, and Women's Relative Wages Over Time', **The Quarterly Journal of Economics,** 123 (2008), 1061-1110.

[17] Pollak, Robert A. and Michael L. Wachter, 'The Relevance of the Household Production Function and Its Implications for the Allocation of Time', **Journal of Political Economy,** 83 (1975), 255-277.

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

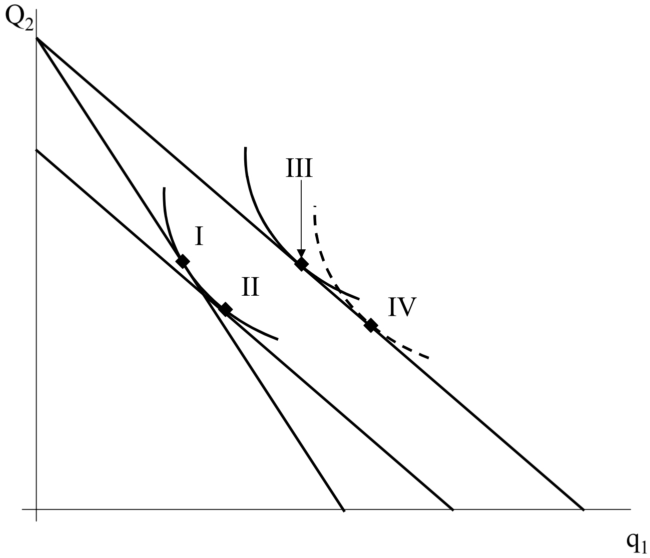

Figure 1: Collective price responses

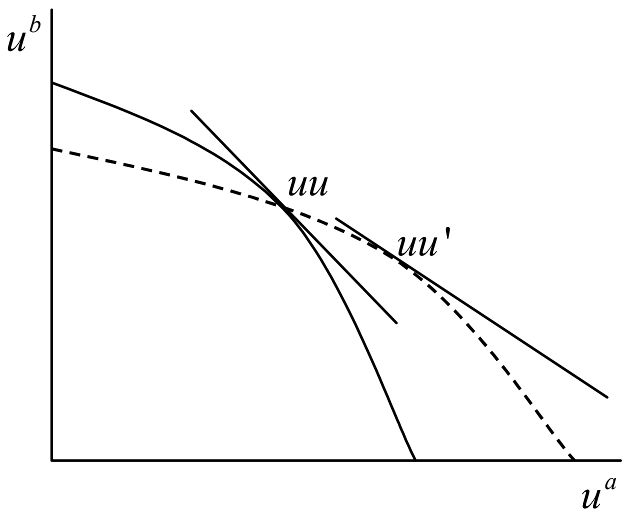

Figure 2: A potentially compensating variation.

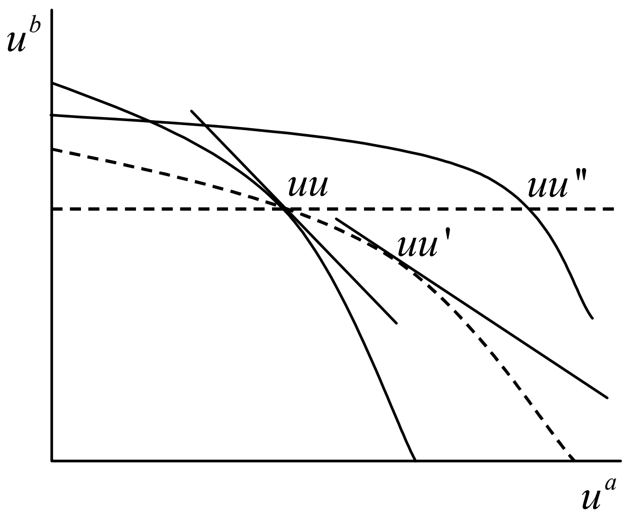

Figure 3: An actually compensating variation.