The Economics of the Family

Chapter 3: Preferences and decision making

Martin Browning
Department of Economics, Oxford University

Pierre-André Chiappori
Department of Economics, Columbia University

Yoram Weiss
Department of Economics, Tel Aviv University

October 2010

1 Preferences

In the last chapter we informally reviewed the gains from marriage in some generality. The existence of potential gains from marriage is not sufficient to motivate marriage and to sustain it. Prospective mates need to form some notion as to whether families realize the potential gains and how they are divided. In this chapter we consider these issues in a very specific context. The context is a two person, woman a and man b household[1] in which the only (static) decision is how much to spend on various market goods that are available at fixed prices, given fixed total household expenditure on all goods. Although very special this context allows us to discuss formally many of the issues that will be used in other contexts in later chapters.

Some commodities are private and some public. Private goods are consumed non-jointly by each partner and public, such as heating, are consumed jointly and non-exclusively by the two partners. In other words, private goods are characterized by an *exclusion restriction* property: the fact that I consume a particular apple de facto excludes anyone else from consuming the same apple; with public goods, on the contrary, no such restriction exists: that I enjoy seeing a beautiful painting on my wall does not preclude my spouse from enjoying it just as much. Several remarks can be made at that point. First, several commodities are sometimes used publicly and sometimes privately; for instance, I can drive my car alone to go to work, or the whole family may take a ride together. Second, the privateness or publicness of a good is quite independent of the type of control existing on that good and who exerts it; typically, parents have control over the (private) consumption of their young children. Finally, and more crucially, there exist subtle interactions between the ('technical') nature of a good and how it enters the member's utilities. The private consumptions of member a certainly enter a's utility; but it also may enter b's - we then call it an *externality*. Conversely, some commodities, although public by nature, may in fact be exclusively consumed by one member; for instance, although both spouses may in principle watch television together without exclusion, one of them may simply dislike TV and never use it. Throughout most of the book, we assume, to keep things simple, that any particular commodity is either purely public or purely private, although many of our results would extend to more general settings.

We introduce some notations that will be used throughout the chapter. There are N public goods, and the market purchase of public good j is denoted Q_j; the N-vector of public goods is given by $\mathbf{Q}$. Similarly, private goods are denoted q_i with the n-vector $\mathbf{q}$. Each private good bought is divided between the two partners so that a receives q_i^a of good i and b receives $q_i^b = q_i - q_i^a$. Hence the vectors a and b receive are $\mathbf{q}^a$ and $\mathbf{q}^b$ respectively, with $\mathbf{q}^a + \mathbf{q}^b = \mathbf{q}$. An *allocation* is a $N + 2n$-vector $\left(\mathbf{Q}, \mathbf{q}^a, \mathbf{q}^b\right)$. The associated market prices are given by the N-vector $\mathbf{P}$ and the n-vector $\mathbf{p}$ for public and private goods respectively.

[1] Children will be introduced at a later point.

We assume that each married person has her or his own preferences over the allocation of family resources. Denote a's utility function by $U^a\left(\mathbf{Q},\mathbf{q}^a,\mathbf{q}^b\right)$ and b's by $U^b\left(\mathbf{Q},\mathbf{q}^a,\mathbf{q}^b\right)$. This general formulation allows that a is concerned directly with b's consumption and also that b's consumption of private goods impacts on a's preferences between her own private goods and the public goods. Any kind of externality is allowed. The presence of the partner's private consumption does not mean necessarily that members are altruistic to each other; for instance, it could simply represent the partner's smoking that bothers the other member by reducing their utility. Throughout the book, we assume, unless stated otherwise, preference orderings are continuous and convex and can be represented by utility functions U^s, $s=a,b$, that are continuously differentiable and strictly concave.

In the subsequent chapters in the first half of this book we shall be discussing the resolution of conflicts that arise between partners. It is important to acknowledge, however, that if marriage is sometimes a battleground, it is can also be a playground. In the context of the family, love or affection might be operating and conflicts are thereby considerably attenuated. If we consider the utility possibility frontier (UPF) available to a couple then love or mutual affection removes the extreme (unequal) points of the UPF. This will generally leave some areas of disagreement but these are smaller than they would be in the absence of love.

Although quite reasonable, the form just described is too general to be used in most contexts - if only because it is difficult to incorporate such preferences into a model in which agents live alone for some part of their life-cycle. Consequently the literature generally takes a special case which is known as *caring*.[2] For this we first assume agents a and b have felicity functions $u^a\left(\mathbf{Q},\mathbf{q}^a\right)$ and $u^b\left(\mathbf{Q},\mathbf{q}^b\right)$ respectively. The most general form has

$$\begin{aligned} U^a\left(\mathbf{Q},\mathbf{q}^a,\mathbf{q}^b\right) &= W^a\left(u^a\left(\mathbf{Q},\mathbf{q}^a\right),u^b\left(\mathbf{Q},\mathbf{q}^b\right)\right) \\ U^b\left(\mathbf{Q},\mathbf{q}^a,\mathbf{q}^b\right) &= W^b\left(u^a\left(\mathbf{Q},\mathbf{q}^a\right),u^b\left(\mathbf{Q},\mathbf{q}^b\right)\right), \end{aligned} \tag{1}$$

where both $W^a(.,.)$ and $W^b(.,.)$ are monotone increasing functions. The weak separability of these 'social' preferences represents an important moral principle; a is indifferent between bundles $\mathbf{q}^b, Q$ that b consumes whenever b is indifferent (similarly for b). In this sense caring is distinguished from paternalism. Caring rules out direct externalities because a's evaluation of her private consumption $\mathbf{q}^a$ does not depend directly on the private goods that b consumes (and *vice versa*). A more commonly used form is the restricted version:

$$\begin{aligned} U^a\left(\mathbf{Q},\mathbf{q}^a,\mathbf{q}^b\right) &= u^a\left(\mathbf{Q},\mathbf{q}^a\right)+\delta^a u^b\left(\mathbf{Q},\mathbf{q}^b\right), \\ U^b\left(\mathbf{Q},\mathbf{q}^a,\mathbf{q}^b\right) &= u^b\left(\mathbf{Q},\mathbf{q}^b\right)+\delta^b u^a\left(\mathbf{Q},\mathbf{q}^a\right). \end{aligned} \tag{2}$$

[2] Sometimes the term *altruistic* is used for preferences taking this form. Pollak (2006) has suggested the term *deferential* since each person defers to the judgment of the other regarding their consumption.

Generally we take the weights δ^a and δ^b to be non-negative parameters such that each person cares for the other but not as much as they care for themselves. For this formulation, $\delta^a = \delta^b = 0$ corresponds to *egotistic* preferences and $1 > \delta^s > 0$ represents altruism.

Some authors use a slightly different representation of altruism, namely

$$\begin{aligned} U^a\left(\mathbf{Q}, \mathbf{q}^a, \mathbf{q}^b\right) &= u^a\left(\mathbf{Q}, \mathbf{q}^b\right) + \tilde{\delta}^a U^b\left(\mathbf{Q}, \mathbf{q}^a, \mathbf{q}^b\right) \\ U^b\left(\mathbf{Q}, \mathbf{q}^a, \mathbf{q}^b\right) &= u^b\left(\mathbf{Q}, \mathbf{q}^b\right) + \tilde{\delta}^b U^a\left(\mathbf{Q}, \mathbf{q}^a, \mathbf{q}^b\right) \end{aligned} \tag{3}$$

The logic here is that a should cares about b's ultimate utility U^b, which includes also b's altruistic feelings towards a. We can then think of (2) as a reduced form obtained by the substitution:

$$U^a\left(\mathbf{Q}, \mathbf{q}^a, \mathbf{q}^b\right) = u^a\left(\mathbf{Q}, \mathbf{q}^b\right) + \tilde{\delta}^a[u^b\left(\mathbf{Q}, \mathbf{q}^b\right) + \tilde{\delta}^b U^a\left(\mathbf{Q}, \mathbf{q}^a, \mathbf{q}^b\right)] \tag{4}$$

If $\tilde{\delta}^a\tilde{\delta}^b \neq 1$ we have:

$$U^a\left(\mathbf{Q}, \mathbf{q}^a, \mathbf{q}^b\right) = \frac{1}{1 - \tilde{\delta}^a\tilde{\delta}^b} u^a\left(\mathbf{Q}, \mathbf{q}^a\right) + \frac{\tilde{\delta}^a}{1 - \tilde{\delta}^a\tilde{\delta}^b} u^b\left(\mathbf{Q}, \mathbf{q}^b\right) \tag{5}$$

Such a reduction yields logical results only if $\tilde{\delta}^a\tilde{\delta}^b < 1$. Clearly, too much caring ($\tilde{\delta}^a\tilde{\delta}^b > 1$) can lead to paradoxical results in which a puts negative weights on both felicity functions. See Bergstrom (1989) and Bernheim and Stark (1988) for further discussion and examples of how excessive altruism can lead to unpalatable outcomes.

In some contexts we wish to impose stronger restrictions on preferences. For example, we shall often consider only one private good. This can be justified if prices are fixed by an appeal to the composite commodity theorem. In that case we can consider the unique private good to be 'money'. A second, particular case that we shall consider in many contexts relies on the assumption of *transferable utility* (TU). This holds if we have egotistic preferences and each felicity function can be (possibly after an increasing transform and a renaming of the private goods) put into a form that is similar to the Gorman polar form:

$$\begin{aligned} u^a\left(\mathbf{Q}, \mathbf{q}^a\right) &= f^a\left(q_2^a, ..., q_n^a, \mathbf{Q}\right) + G\left(\mathbf{Q}\right) q_1^a \\ u^b\left(\mathbf{Q}, \mathbf{q}^b\right) &= f^b\left(q_2^b, ..., q_n^b, \mathbf{Q}\right) + G\left(\mathbf{Q}\right) q_1^b \end{aligned} \tag{6}$$

where $G\left(\mathbf{Q}\right) > 0$ for all $\mathbf{Q}$. Note that the $G\left(.\right)$ function is identical for both members, whereas the $f\left(.\right)$ functions can be individual-specific. In words, the transferable utility assumption implies that, for some well chosen cardinalization of individual preferences, the marginal utility of an additional dollar spent on private consumption of commodity 1 is always the same for both members. Hence utility can be 'transferred' between them (using commodity

1 transfers) at a fixed rate of exchange. Repeatedly in this book, we shall develop examples in which the transferability assumption drastically simplifies the problem to hand.

We shall often need to compare the utility of a given individual in two different marital situations, for instance when married versus when single (or divorced). Various assumptions can be made here. One extreme hypothesis states that marriage may change preferences in an arbitrary way. Then there is simply no relation between pre- and post-marital utility functions - not a very useful property for our purpose. Conversely, we may assume that preferences *over commodities* are not changed by marriage. This by no means implies that the satisfaction derived from any consumption is the same before and after marriage, but simply that the ranking of the various consumption bundles is not affected by the individual's marital status. With egotistic preferences, this will hold if the utility after marriage, u^s, is related to the pre-marital preferences represented by the utility function, $\bar{u}^s(\mathbf{Q}, \mathbf{q}^s)$ by:

$$u^s(\mathbf{Q}, \mathbf{q}^s) = F^s(\bar{u}^s(\mathbf{Q}, \mathbf{q}^s)) \tag{7}$$

where the mapping $F^s(.)$ is strictly increasing. A particularly convenient special case that we shall employ when we consider explicitly the full gains from marriage is:

$$u^s(\mathbf{Q}, \mathbf{q}^s) = F(\bar{u}^s(\mathbf{Q}, \mathbf{q}^s) + \theta^s) \tag{8}$$

Here, θ^s represents non-monetary, marriage-specific aspects of s's idiosyncratic desire to be married. With caring preferences, the same obtains if we normalize the contribution of the spouse's utility to be uniformly zero when the agent is single. This assumption has important consequences on the empirical estimation of the models. If condition (8) is satisfied, then the preferences of married individuals amongst private and public goods are the same when married or single. These preferences can be recovered from data on singles' behavior.

Finally, an intermediate assumption states that single and married individuals have the same basic preferences, but marriage involves a change in the consumption technology, a concept we define in the next subsection.

2 Household production

2.1 The general framework.

Household activities are not limited to private or public consumptions. They are also the source of important production activities that should not be disregarded. In low income countries, a large fraction of GDP consists of agricultural commodities produced at the household (or the village) level. Even in high income economies, a significant fraction of individual available time is spent on household production. This entails immediate tasks (cleaning, cooking, etc.) but also long term investments in health, education and others. In a sense, even such 'commodities' as love, affection or mutual care are 'produced' (and consumed) at the household level. In Becker's (1965) model, the only commodities that are ultimately consumed by individuals are those produced at the household level. Becker views goods that are purchased in the market as inputs in a production system that transforms these purchased goods into final commodities that are actually consumed (and enter individual utilities). These home produced goods can be either public or private for the two partners, denoted by C_j and c_j respectively. The production of commodities also requires time inputs that are provided by the household members in addition to market purchased goods. The technology is described by a production possibility set $\Omega\left(\mathbf{q}, t^a, t^b\right)$ that gives the possible vector of outputs $(\mathbf{c}, \mathbf{C})$ that can be produced given a vector of market purchases $\mathbf{q}$ and the total time spent in household production by each of the two partners, t^a and t^b. This allows that time spent on any activity may produce many goods.

2.1.1 Household production function

A special case is when the feasible set can be described by *household production functions* that specify the amount of each commodity that can be produced given the amount of market goods and time assigned to that commodity. We denote the vector of market goods assigned to commodity j by $\mathbf{q}^j$ and the time inputs of a and b devoted to commodity j by t_j^a and t_j^b, respectively. Thus:

$$c_j = f^j\left(\mathbf{q}^j, t_j^a, t_j^b\right) \tag{9}$$

The associated constraints are:

$$\begin{aligned} &\sum_j \mathbf{q}^j = \mathbf{q} \\ &\sum_j t_j^s = t^s, \qquad s = a, b \end{aligned} \tag{10}$$

Each person has preferences defined over household produced goods and the vectors of time use, $U^s\left(\mathbf{C}, \mathbf{c}^a, \mathbf{c}^b, \mathbf{t}^a, \mathbf{t}^b\right)$ for $s = a, b$, where $\mathbf{t}^s$ is the vector of time inputs for j. This framework allows time activities to have two distinct roles. For example, a father who spends time with his child contributes to the development of the child (through $f^j(.)$) and may also enjoy spending time with the child (captured by the presence of t_j^b in $U^b(.)$). Of course, either of these effects could be negative (although not both).

A standard problem with this approach is that the production function, despite its conceptual interest, cannot be estimated independently of the utility function unless the home produced commodities are independently observable; see Pollak and Wachter (1975) and Gronau (2006). Observability of outputs may be acceptable for agricultural production, or even for children's health or education; it is less likely for, say, cleaning, and almost impossible for personal caring.

If only inputs are observed and not outputs we may be able to recover information about the technology if we make auxiliary assumptions such as constant returns to scale and assumptions on preferences. To illustrate this, consider two partners who consume one single public good C and one private good c such that a consumes c^a and b consumes c^b with preferences given by $u^s(C, c^s), s = a, b$. Assume that the private good is purchased in the market and that the public commodity is produced using only the time inputs of the two partners. That is,

$$C = f\left(t^a, t^b\right). \tag{11}$$

Assuming that both partners participate in the labour market at wages w^a and w^b respectively, it can then be shown that for any efficient allocation the partners will minimize the cost of producing the public commodity in terms of the forgone private commodity, yielding

$$\frac{f_1\left(t^a, t^b\right)}{f_2\left(t^a, t^b\right)} = \frac{w^a}{w^b} \tag{12}$$

in any interior solution. Under constant returns to scale, we can write:

$$C = f\left(t^a, t^b\right) = t^b \phi(r) \tag{13}$$

for some function $\phi(r)$ where $r = \frac{t^a}{t^b}$. The condition (12) then reduces to:

$$\frac{\phi'(r)}{\phi(r) - r\phi'(r)} = \frac{w^a}{w^b} \tag{14}$$

The testable implication of this equality is that r only depends on the wage ratio ω; this can be tested on a data set that reports wages and time spent on household production. Defining,

$$h(r) = \frac{\phi'(r)}{\phi(r) - r\phi'(r)} \tag{15}$$

this equation can be re-written as:

$$\frac{\phi'(r)}{\phi(r)} = \frac{1}{r + \frac{1}{h(r)}}$$

Integrating, we have:

$$\phi(r) = K \exp\left(\int_0^r \frac{ds}{s + \frac{1}{h(s)}}\right)$$

where K is an unknown constant. In other words, when outputs are not observable, knowledge of the input supply (as a function of relative wages) allows us to determine the household production function up to a multiplicative scale factor.

It is important to note that the assumptions of constant returns to scale and no joint production (in the sense that t^a and t^b do not appear directly in the utility function) are critical for this particular identification result; see Pollak and Wachter (1975) and Gronau (2006) for further discussion of the role of these assumptions. A further issue that was not challenged in this literature is whether or not the partners are cooperating. The example above shows that in some cases it is sufficient to assume efficiency; other assumptions may also guarantee identification.

2.1.2 Marital technology and indifference scales

Let us briefly come back to the previous discussion on the changes in preferences that may result from marriage. The two extreme assumptions described were either that there are no such changes (in the sense that an individual's preference relationship over consumption bundles was independent of the person's marital status) or that they were arbitrary (that is, there is no relationship between pre- and post- marital utilities). The first assumption is often too restrictive, whereas the second is too general to be useful.

An intermediate approach, proposed by Browning, Chiappori and Lewbel (2003), relies on the notion of production technology. The idea is that marriage leaves ordinal preferences over commodities unchanged, but allows a different (and more productive) technology to be used. Formally, they apply the simple Barten household production technology in which n market goods are transformed into n household commodities in a linear and non-joint way; see the discussion in chapter 2, section **??**. This setting allows us to separate the identification of preferences (which can be done on a subsample of singles) and that of the production function (for which household level data are needed). Not surprisingly, being able to observe identical individuals operating under different technologies (that is, as single or married) considerably facilitates identification. Browning *et al* show that the model can be estimated from the observation of demand functions for individuals and couples.

A crucial outcome of this approach is the computation of each member's *indifference scale,* defined as the minimum fraction of the household's income that this member would need to buy (at market prices) a bundle of privately consumed goods that put her *on the same indifference curve over goods* that she attained as a member of the household. Note that this amount is different (and lower) than what would be needed to purchase, as a single, the same *bundle* the member was consuming when married. Indeed, an obvious effect of the household technology is that the prices implicitly used within the household may differ from market prices; see chapter 2, section 2. It follows that even for a given level of expenditures, the consumption profile of a couple typically differs from that of single individuals.

2.2 Children

Modeling children is a complex issue, and one in which even basic methodological choices may be far from innocuous in terms of normative implications. A general approach relies on two basic ideas. One is that, in general, parents care about their children. This could take the form of parent s caring directly about the goods that the child consumes:

$$U^s = U^s\left(\mathbf{Q}, \mathbf{q}^a, \mathbf{q}^b, \mathbf{q}^k, t^a, t^b\right) \tag{16}$$

where t^s are the time inputs of the parents. A widely used special case posits the existence of a child utility function:

$$u^k = u^k\left(t^a, t^b, \mathbf{Q}, \mathbf{q}^k\right) \tag{17}$$

where $\mathbf{q}^k$ denotes the vector of consumption by the child. Then the preferences of adult s can be defined recursively by:

$$U^s\left(\mathbf{Q}, \mathbf{q}^a, \mathbf{q}^b, t^a, t^b\right) + \kappa^s u^k\left(\mathbf{Q}, \mathbf{q}^k, t^a, t^b\right) \tag{18}$$

where κ^s is the weight that parent s gives to the children.[3]

Of course, this approach can be used with any number of children. Depending on the problem to hand, one may either introduce one sub-utility per child or only distinguish between broader 'classes' (for example, boys versus girls, younger children versus older ones, etc.). Timing introduces additional issues, since parents care not only about their children but also their grandchildren. Barro and Becker (1988) have introduced the concept of dynastic utilities, whereby parents actually consider the sum of utility levels of *all* of their descendents, weighted according to some 'distance' factor $\kappa^s < 1$. Then adult s's utility takes the form $u^s + \sum_{t=1}^{\infty} (\kappa^s)^t u^{k(t)}$, with the convention that $u^{k(1)}$ denotes the utility of s's children, $u^{k(2)}$ of his grandchildren, and so on. This model, which relies on restricting (16) to (18), may have strong policy implications. For instance, government subsidies given to children can be completely offset by lower contribution of their parents without any effect on the final outcome. (This type of neutrality is often termed Ricardian Equivalence).

It is important to note that in this context, children matter for the household's choices, but only through the utility their parents derive from their well-being. This is a strong assumption, that can be relaxed in two directions. First, one may, alternatively, consider the child as another decision maker within the household. In this case a couple with one child would be considered as a three person household. Whether a child should be considered as a decision maker or not is a very difficult question, which may depend on a host of factors (age, education, occupation, income, etc.); moreover, its empirical translation introduces subtle differences that are discussed below.

Secondly, throughout this book we stick to a standard practice in economics, and we assume that preferences are given, that is, exogenous and stable. This assumption may be acceptable for adults, but maybe less so for children; after all, many parents invest time and resources into influencing (or 'shaping') their children's preferences regarding work, risk, or 'values' in some general sense. Indeed, a growing literature analyzes the formation of preferences from an economic viewpoint, as a particular 'production' process. These contributions are outside the scope of this book; the interested reader is referred to Becker (1998).

3 The unitary model

Having established the nature of preferences, the notion of household production function and the concept of a distribution factor, we now consider how the partners in the household make decisions on how to spend their time and

[3] A more general formulation would have utilities of the form $u^s\left(\mathbf{Q}, \mathbf{q}^a, \mathbf{q}^b, u^k\right)$.

money. To bring out the main ideas we consider a context in which there are no time allocation decisions, income is given and there is no household production. We denote the market price of public and private goods by $\mathbf{P}$ and $\mathbf{p}$ respectively. In this chapter we generally take the incomes of a and b to be given at levels Y^a and Y^b respectively and we assume that there is no other income into the household. We further assume that household total expenditure, x, is set equal to household income $Y = Y^a + Y^b$, so that there is no borrowing or lending. The household budget constraint for allocations is given by:

$$\mathbf{P}'\mathbf{Q} + \mathbf{p}'\left(\mathbf{q}^a + \mathbf{q}^b\right) = x \tag{19}$$

In general the agents will differ on how to spend household income. There are three broad classes of decision processes: the unitary assumption, non-cooperative processes and cooperative processes.

The most widely used assumption is that choices are made according to a 'unitary' household utility function $\tilde{U}\left(\mathbf{Q}, \mathbf{q}^a, \mathbf{q}^b\right)$. In subsection 5.6 we shall investigate when such an assumption is justified, but for now we simply consider the consequences. One natural assumption, due to Samuelson (1956), is to impose on the household utility function is that it respects the individual preferences in the sense that:

$$\tilde{U}\left(\mathbf{Q}, \mathbf{q}^a, \mathbf{q}^b\right) = W\left(U^a\left(\mathbf{Q}, \mathbf{q}^a, \mathbf{q}^b\right), U^b\left(\mathbf{Q}, \mathbf{q}^a, \mathbf{q}^b\right)\right) \tag{20}$$

where W is a *utility weighting function* which is strictly increasing in the individual utilities. The important feature of this weighting function is that it is fixed and does not vary with prices or income. Given a unitary utility function we define a household utility function over market goods by:

$$U\left(\mathbf{Q}, \mathbf{q}\right) = \max_{\mathbf{q}^a} \tilde{U}\left(\mathbf{Q}, \mathbf{q}^a, \mathbf{q} - \mathbf{q}^a\right) \tag{21}$$

Given this household utility function we can derive market demands in the usual way; namely, it solves the program

$$\max_{(\mathbf{Q},\mathbf{q})} U\left(\mathbf{Q}, \mathbf{q}\right)$$
$$\text{subject to} \quad \mathbf{P}'\mathbf{Q} + \mathbf{p}'\mathbf{q} \le x$$

We assume enough on preferences (continuous differentiability, strict concavity), so that this leads to demands for market goods:

$$\mathbf{Q} = \boldsymbol{\Xi}\left(\mathbf{P}, \mathbf{p}, x\right) \tag{22}$$
$$\mathbf{q} = \boldsymbol{\xi}\left(\mathbf{P}, \mathbf{p}, x\right) \tag{23}$$

The unitary assumption has two important sets of implications. First, market demand functions satisfy the usual *Slutsky conditions;* adding-up, homogeneity, symmetry and negativity of the Slutsky matrix; see, for example, Mas-Colell, Whinston and Green (1995), chapter 3. Second, the demands only depend on prices and total household income, and are independent of the distribution of income; that is, they display *income pooling.* As we shall see below the latter has been the focus of much testing in the empirical literature.

4 Non-cooperative models

4.1 No public goods.

If we are not willing to assume a unitary utility function then we must specify a decision process. As always there are very many possibilities here but we shall only explore a small subset of these. We begin with non-cooperative procedures.[4] If household behavior is modeled non-cooperatively, then no binding agreements between members are assumed and the optimal decisions need not be Pareto efficient. However, in some cases efficiency obtains automatically as an outcome of independent decision making. Take the simple situation in which preferences are egotistic, and all commodities are privately consumed. The noncooperative solution boils down to the following two programs:

$$\begin{aligned}
&\max_{\mathbf{q}^a} \left\{ u^a\left(\mathbf{q}^a\right) \text{ subject to } \mathbf{p}'\mathbf{q}^a = Y^a \right\} \\
&\max_{\mathbf{q}^b} \left\{ u^b\left(\mathbf{q}^b\right) \text{ subject to } \mathbf{p}'\mathbf{q}^b = Y^b \right\}
\end{aligned} \tag{24}$$

In words, the noncooperative solution simply implies in that case that each agent chooses his/her level of consumption independently of the other; i.e., they live side by side, but without any economic interaction.[5] Then the consumption of individual s is simply this individual's demands at prices $\mathbf{p}$ and income Y^s. Denote the demand functions for s by $\boldsymbol{\xi}^s\left(\mathbf{p}, Y^s\right)$. Note that the allocation $\left(\boldsymbol{\xi}^a\left(\mathbf{p}, Y^a\right), \boldsymbol{\xi}^b\left(\mathbf{p}, Y^b\right)\right)$ is Pareto efficient: clearly, the utility of, say, a can only be increased by an income transfer from b, which would reduce b's welfare. Generally the associated household demands

$$\boldsymbol{\xi}\left(\mathbf{p}, Y^a, Y^b\right) = \boldsymbol{\xi}^a\left(\mathbf{p}, Y^a\right) + \boldsymbol{\xi}^b\left(\mathbf{p}, Y^b\right) \tag{25}$$

will not satisfy income pooling or the Slutsky conditions. The special case in which income pooling and the Slutsky conditions will hold is if the classic aggregation conditions hold. That is, if the two agents have linear Engel curves with each partner having the same slope for any good:

$$\begin{aligned}
\xi_i^a &= \phi_i^a\left(\mathbf{p}\right) + \varphi\left(\mathbf{p}\right) Y^a \\
\xi_i^b &= \phi_i^b\left(\mathbf{p}\right) + \varphi\left(\mathbf{p}\right) Y^b
\end{aligned} \tag{26}$$

so that the household demand for good i is given by:

$$\begin{aligned}
\xi_i\left(\mathbf{p}, Y^a, Y^b\right) &= \phi_i^a\left(\mathbf{p}\right) + \phi_i^b\left(\mathbf{p}\right) + \varphi\left(\mathbf{p}\right)\left(Y^a + Y^b\right) \\
&= \phi_i^a\left(\mathbf{p}\right) + \phi_i^b\left(\mathbf{p}\right) + \varphi\left(\mathbf{p}\right) Y
\end{aligned} \tag{27}$$

[4] Several authors take the Nash position that any cooperative game should be preceded by a non-cooperative game. Some of the authors cited in this section only develop a noncooperative interaction for this purpose.

[5] Of course, this does not preclude the existence of non economic interactions - love, sex, conversation or others.

4.2 One public good.

Whenever a direct interaction between members is introduced - either because of public consumption, or because one member's consumption has an external effect on the other member's well being - inefficiencies are very likely to appear. To bring out the essential features of the analysis, let us assume that there is only one public good and one private good and that each person has egotistic preferences; see Chen and Wooley (2001) and Browning, Chiappori and Lechene (2006). Given that we have a public good and individual incomes a natural, noncooperative process to consider is a voluntary contributions game in which each person contributes to the purchase of the public good and then uses any money remaining to buy the private good for themselves. That is, the two agents have the problems:

$$\begin{aligned}
&\max_{Q^a,q^a}\left\{u^a\left(Q^a+Q^b,q^a\right)\text{ subject to }PQ^a+pq^a=Y^a\right\}\\
&\max_{Q^b,q^b}\left\{u^b\left(Q^a+Q^b,q^b\right)\text{ subject to }PQ^b+pq^b=Y^b\right\}
\end{aligned}\tag{28}$$

where Q^s denotes agent s's contribution to the public good. Assuming that both goods are normal, this interaction has exactly one Nash equilibrium, which can take one of two forms. In the first form, both agents contribute to the public good. Since this is an interior solution for both we have:

$$\begin{aligned}
\frac{u^a_Q}{u^a_q}\left(\hat{Q},\hat{q}^a\right)&=\frac{P}{p}\\
\frac{u^b_Q}{u^b_q}\left(\hat{Q},\hat{q}^b\right)&=\frac{P}{p}
\end{aligned}\tag{29}$$

If we sum the budget constraints we have:

$$P\hat{Q}+p\left(\hat{q}^a+\hat{q}^b\right)=Y^a+Y^b\tag{30}$$

Thus we have three equations in three unknowns $\left(\hat{Q},\hat{q}^a,\hat{q}^b\right)$ with a solution:

$$\begin{aligned}
\hat{q}^a&=\xi^a\left(P,p,Y^a+Y^b\right)\\
\hat{q}^b&=\xi^b\left(P,p,Y^a+Y^b\right)\\
\hat{Q}&=\Xi\left(P,p,Y^a+Y^b\right)
\end{aligned}\tag{31}$$

We conclude that the household's market demand for both the public good, $\hat{Q}$, and the private good, $\hat{q}=\hat{q}^a+\hat{q}^b$ depends only on total household income Y^a+Y^b and not on how it is distributed. In other words, we have *income pooling* even though we have a non-unitary model. This is an example of the remarkable neutrality result due to Warr (1983) (see also Bergstrom, Blume and Varian (1986) and Bernheim (1986)). This shows that while income pooling is a necessary condition for the unitary model, it is not sufficient.

It is important to note that income pooling here is a local property and holds only as long as both partners contribute to the public good. The other case we have to consider is the one in which only one person contributes. If this is person a, the first order condition in (29) holds for her. Person b spends all of his income on the private good, so that:

$$\frac{u_Q^b}{u_q^b}\left(\hat{Q}, \frac{Y^b}{p}\right) \leq \frac{P}{p} \tag{32}$$

with a strict inequality if the agent is not on the margin of contributing to the public good. In this case a redistribution of income from a to b will generally change the market demand since b will increase his demand for the private good and a generally will not change her demands to exactly offset these revisions. Thus we have market demands:

$$\begin{aligned} \hat{q} &= \hat{q}^a + \frac{Y^b}{p} = \xi^a\left(P, p, Y^a\right) + \frac{Y^b}{p} \\ \hat{Q} &= \hat{Q}^a = \Xi\left(P, p, Y^a, Y^b\right) \end{aligned} \tag{33}$$

In both cases, the allocations the non-cooperative procedure leads to an inefficient outcome (except for the cases in which one or the other has all of the income); this is the standard under-provision for the voluntary contributions public goods game. To see that for the case of an interior solution, add the two first order conditions (29), yielding

$$\frac{u_Q^a}{u_q^a}\left(\hat{Q}, \hat{q}^a\right) + \frac{u_Q^b}{u_q^b}\left(\hat{Q}, \hat{q}^b\right) = \frac{1}{2}\frac{P}{p},$$

while Samuelson's (1954) condition for an efficient allocation of public goods requires that

$$\frac{u_Q^a}{u_q^a}\left(\hat{Q}, \hat{q}^a\right) + \frac{u_Q^b}{u_q^b}\left(\hat{Q}, \hat{q}^b\right) = \frac{P}{p}. \tag{34}$$

That is, the sum of the willingness to pay for the public good of the two partners, should equal to the opportunity cost of the public good in terms of the private good. In this regard, there is an under provision of the public good. [6]

We now present an example to illustrate some of the points made here. Normalize prices to unity, $P = p = 1$, and take preferences represented by $u^a = q^a Q^\alpha$ and $u^b = q^b Q$. The parameter α governs how much a likes the public good; if $\alpha > 1$ then she values it more than b if they have the same

[6] Results on dynamic contributions games suggest that inefficiencies cam be eliminated if players contribute sequentially and cannot reduce previous contributions; see, for example, Matthews (2006).

private consumption. We set $Y^a = \rho$ and $Y^b = (1-\rho)$ so that household income is unity and ρ is a's share of household income. It is straight forward to show that the decisions of the agents are given by:

$$\hat{Q}^a = \min\left(\max\left(0, \rho - \frac{1}{1+2\alpha}\right), \frac{\alpha\rho}{1+\alpha}\right)$$
$$\hat{Q}^b = \min\left(\max\left(0, (1-\rho) - \frac{\alpha}{1+2\alpha}\right), \frac{(1-\rho)}{2}\right)$$

(with the demands for private goods being given by the budget constraints). It is easiest to see the implications of this if we graph $\hat{Q} = \hat{Q}^a + \hat{Q}^b$ against a's share of income, ρ. In figure 1 we do this for two values of α, 0.8 and 1.2. There are a number of notable features to figure 1. First, if b has all the income ($\rho = 0$) then the level of public goods provision corresponds to his preferred level; here a value of one half. If we now redistribute some income from b to a we see that the level of the public good falls whether or not a has a higher valuation for the public good ($\alpha \lesseqgtr 1$). This is because a uses all of her income for her own private good and b reduces spending on both the public good and his private good. As we continue shifting income form b to a the level of the public good falls until at some point a starts to contribute. The level at which a starts to contribute is lower the higher is the value of her liking for the public good (compare the curves for $\alpha = 0.8$ and $\alpha = 1.2$). Once both partners are contributing to the public good, further small transfers from b to a leave all allocations unchanged as a increases her contributions and b reduces his in an exactly offsetting way (this is the local income pooling result). At some point b's income falls to the point at which he stops contributing. This level of income is lower the higher is the level of the provision of the public good. After this, transfers of income from b to a cause a to increase her contribution to the public good until she has all of the income ($\rho = 1$). To illustrate that the level of provision of the public god is inefficiently low for any value of $\rho \in (0,1)$, consider the case $\alpha = 1$ and $\rho = 0.5$. The equilibrium choices are $\hat{Q} = 1/3$ and $q^a = q^b = 1/3$. This gives each a utility level of 1/9. If we instead impose that each contributes 0.25 to the public good and spends 0.25 on their own private good then each has a utility level of 1/8, which is a Pareto improvement on the equilibrium outcome.

4.3 Altruism and the Rotten Kid Theorem.

An important extension to this analysis is to move beyond the egoistic assumption and to allow for altruism. In this case, if one person has most (or all) of the income and cares for the other then they may make a transfer of private goods to the poorer partner as well as being the sole contributor to the public good. This adds flat segments at the extremes of figure 1, as shown in figure 2. In this figure the demand for the public good if a's income share is between ρ_1 and ρ_4 is of the same form as for the egoistic case. If, however, we have an extreme distribution of income then the figure changes.

For example, if b has most of the income ($\rho < \rho_1$) and cares for a then he will transfer some private good to her and will be the only contributor to the public good (since ρ_2 is the distribution at which a starts to contribute). In this region we have three important implications. First, there is local income pooling and small re-distributions of income within the household would not change the allocations $\left(\hat{Q}, \hat{q}^a, \hat{q}^b\right)$. Second, the outcome is efficient since b is an effective dictator; any other feasible allocation will make b worse off. Third, the household demands for private goods $\left(\hat{q}^a(P, p, Y) + \hat{q}^b(P, p, Y)\right)$ and public goods $(\hat{Q}(P, p, Y))$ will satisfy the Slutsky conditions. Note, however, that the range of this unitary-like behavior and efficiency will depend on the degree of altruism; as drawn, b cares more for a than a cares for b (the interval $[0, \rho_1]$ is shorter than the interval $[\rho_4, 1]$).

In chapter 8 of his revised Treatise of 1991 Becker refers to the unitary style implications (efficiency, income pooling and the Slutsky conditions) as the Rotten Kid Theorem (RKT); see also Becker (1974). If one person has enough income relative to the other and cares for them then they internalize all decisions and the household behaves as though it is one. A corollary is that each household member will be motivated to maximize total household income. For example, if we have $\rho < \rho_1$ and a can take some action that lowers her income but increases b's income by more, she will choose to do it, safe in the knowledge that b will increase the transfer to her sufficiently to make her better off. This presentation makes it clear that the scope of the RKT (in this version) is limited; it only applies locally and requires an extreme distribution of household income and altruism. In subsection 5.10 below we present a more general version of the RKT that is closer in spirit to Becker's original formulation in Becker (1974). This version widens the scope at the cost of imposing restrictions on preferences.

4.4 Many public goods.

When we turn to the more realistic case with more than one public good, the important features we saw above persist but some new ones emerge. The main points can be seen in a model with no altruism, N public goods, a single private good and prices normalized to unity. The voluntary contributions model has:

$$\max_{\mathbf{Q}^a, q^a} \left\{ u^a \left(\mathbf{Q}^a + \mathbf{Q}^b, q^a \right) \text{ subject to } \mathbf{e}'\mathbf{Q}^a + q^a = Y^a \right\} \tag{35}$$

$$\max_{\mathbf{Q}^b, q^b} \left\{ u^b \left(\mathbf{Q}^a + \mathbf{Q}^b, q^b \right) \text{ subject to } \mathbf{e}'\mathbf{Q}^b + q^b = Y^b \right\} \tag{36}$$

where $\mathbf{e}$ is an N-vector of ones. Let $\left(\hat{Q}_1^s, ...\hat{Q}_N^s\right)$ for $s = a, b$ be a Nash equilibrium.[7] We say that person s *contributes* to good j if $\hat{Q}_j^s > 0$. Let m^a (respectively, m^b) be the number of goods to which a (respectively, b)

[7] We assume enough to ensure the existence of at least one Nash equilibrium. We do not impose uniqueness.

contributes. Browning, Chiappori and Lechene (2006) show that if all public goods are bought ($\hat{Q}_j^s > 0$ for at least one s) then either $m^a + m^b = N$ or $m^a + m^b = N + 1$ (generally). This striking result shows that there is at most one public good to which both partners contribute.[8] To see the result informally, suppose that both partners simultaneously contribute to two public goods, i and j. Then both set the marginal rates of substitution between the two goods to unity (the relative prices) and hence equalize the mrs's:

$$\frac{u_i^a}{u_j^a} = \frac{u_i^b}{u_j^b} \tag{37}$$

Without restrictions on preferences and incomes, this is unlikely to hold. Moreover, if it does hold, if we make an infinitesimal change in Y^a or Y^b the property (37) will generally not hold. If there is some overlap in contributions ($m^a + m^b = N + 1$) then we have the local income pooling result, just as in the one public good case when both contribute. As before, the outcomes are generally inefficient since each person under-contributes to their own set of public goods. The result that each partner has a set of public goods which are his or her 'domain' suggests a gender division of allocation within the household. Note, however, that the goods that each takes as their domain is determined endogenously by preferences *and* the division of income within the household. As we move from b having all the income to a having all the income (holding total income constant) the number of goods that she contributes to will generally rise and the number of goods to which he contributes will generally fall.

We illustrate with an example with egoistic preferences from Browning *et al* (2006) for the case of two public goods, G and H. Let the two partners have preferences represented by the pair of Cobb-Douglas utility functions

$$u^a(q^a, G, H) = \ln q^a + \frac{5}{3}\ln G + \frac{8}{9}H$$
$$u^B(q^B, G, H) = \ln q^B + \frac{15}{32}G + \frac{1}{2}\ln H$$

The relative weights on the two public goods are $\frac{45}{24}$ and $\frac{15}{16}$ for a and b respectively; that is, a likes good G relative to good H, more than b. Figure 3 shows the purchases of public goods against a's share of income. When a has a low share of income (region I on the x-axis) she does not contribute to either public good ($m^a = 0$ and $m^b = 2$). In this region, increases in a's income share lead her to spend more on the private good and lead b to spend less on both public goods. If her income is increased to region II then she starts contributing to one of the public goods (good G in this case) and he continues contributing to both ($m^a = 1$ and $m^b = 2$); this is a region of income pooling. As we continue to take income from b and give it to a we move to region

[8]This result is generic in the sense that it is possible to find 'knife-edge' configurations of preferences and incomes for which the two partners contribute to more than one common public good.

III where she contributes to one public good and he contributes to the other ($m^a = 1$ and $m^b = 1$). This is again a region in which the distribution of income matters (locally). Regions IV and V are analogous to regions II and I, with a and b exchanged. One feature to note about this model is the point at which flat segments begin is the same for the two goods (and household expenditures on the private good).

Lundberg and Pollak (1993) introduce a model that is similar to the many public goods version above, except that they restrict contributions to exogenously given sets of public goods for each partner, which they term *separate spheres*. They have two public goods and assume that each partner has a public good to which they alone can contribute; this is their 'sphere' of responsibility or expertise. These spheres are determined by social norms; this is the principal difference from the model developed in the last subsection in which the 'sphere' of influence depends on preferences and the distribution of income within the household and is hence idiosyncratic to each household.

5 Cooperative models: the collective approach

The main problem with non-cooperative procedures is that they typically lead to inefficient outcomes. In a household context this is a somewhat unpalatable conclusion. If each partner knows the preferences of the other and can observe their consumption behavior (the assumption of symmetric information) and the two interact on a regular basis then we would expect that they would find ways to exploit any possibilities for Pareto improvements. This does not preclude the existence of power issues; as we shall see, the notion of 'power' plays a crucial role (and has a very natural interpretation) in cooperative models. The cooperative approach does recognize that the allocation of resources within the household may (and generally) will depend on the members' respective 'weights'; it simply posits that however resources are allocated, none are left on the table.

There are various ways of modeling cooperative behavior. In what follows, we mainly explore the implications of the only restriction common to all cooperative models, namely that household decisions are Pareto efficient, in the usual sense that no other feasible choice would have been preferred by *all* household members. This approach was originally suggested by Chiappori (1988, 1992) and Apps and Rees (1988). Following Chiappori, we refer to such models as *collective* and refer to households that always have Pareto efficient outcomes as *collective households*. More specific representations, based on bargaining theory, are briefly discussed at the end at this section. In the remainder of this chapter we briefly introduce the collective model. Chapters 4 and 5 expand on this discussion.

The collective approach relies on two fundamental assumptions. First, there exists a decision process in the household and it is stable. Second, this process leads to Pareto efficient outcomes. We discuss these aspects successively.

5.1 Decision processes

A fundamental assumption in unitary demand theory is that individual preferences are *stable*, in the sense of not changing capriciously from moment to moment. This is not a logical requirement; in principle, the world could be such that people are intrinsically inconsistent, and a person's preferences today are unconnected with those of yesterday. Clearly, in such a world, very little could be say about individual behavior: a minimum level of stability is necessary if we wish to make predictions based on our models.

The same requirement applies to any model aimed at describing the behavior of a group. The notion of stability, in that case, must be given a broader interpretation: it relates not only to preferences, but also to the decision process. Again, the world could be such that a given household, faced with the same environment in different time periods, adopts each time a different decision process leading to different outcomes. Again, in such a world not much could be predicted about household behavior. We rule out such situations by assuming the existence of a stable decision process. Formally, we define the fundamentals of the model as the preferences of the members and the domestic technologies they can use. A *decision process* is a mapping that associates, to given fundamentals and given vectors of prices, incomes and distribution factors, a probability distribution over the set of consumption bundles. Our first basic assumption is thus the following:

Axiom 1 (Stability) *Each household is characterized by a unique decision process.*

In words: there is a stable relationship between the fundamentals of the model, the economic environment and the chosen outcomes. Note that, in full generality, this relationship needs not be deterministic. It may be the case, for instance, that in some circumstances the process will lead to explicit randomization.[9] What the stability axiom requires in that case is that the randomization be stable, in the sense that, keeping the fundamentals fixed, a household faced with the same economic environment will always randomize in exactly the same way (that is, using the same probability). Nevertheless, in what follows we essentially consider deterministic decision processes.

An important remark, however, is that while *household members* can observe all the factors influencing the decision process, the *econometrician* analyzing their behavior may not have such luck. If some determinants of the process are not observed, we are in a classical situation of unobserved heterogeneity. Then the model may (and typically will), for empirical purposes, involve probability distributions (of unobserved heterogeneity), even when the decision process in each household is purely deterministic. The corresponding randomness should not be considered as intrinsic; it simply reflects the econometrician's inability to observe all the relevant variables.

[9] For instance, a basic conclusion of second best theory is that in the presence of non convexities, randomization may be needed to achieve Pareto efficiency. See Chiappori (2009) for an application to collective labor supply.

Clearly, the stability axiom is not specific to the collective approach; *any* model of group behavior must assume some kind of stability to be able to make predictions. Most of the time, the stability is implicit in the formulation of the model. For instance, in the unitary framework, a unique utility is maximized under a budget constraint; the outcome is the solution to the maximization problem. Similarly, in noncooperative models based on Nash equilibria, the decision process selects, for given fundamentals and a given environment, the fixed point(s) of the best response mapping. In the collective approach, one way to justify the stability axiom could be to assume that the household uses a well specified bargaining protocol, which can be cooperative (Nash, Kalai-Smorodinsky) or noncooperative (e.g., Rubinstein's 'shrinking cake' model). In all cases, the concept under consideration exactly pins down the particular outcome of the negotiations as a function of prices, incomes and distribution factors, which can in that context be interpreted as factors influencing the bargaining game (for example, the status quo points). But, of course, assuming bargaining is by no means necessary for stability.

5.2 Assuming efficiency

The second key assumption of the collective approach is that the outcomes of the decision process are always efficient, in the (usual) sense that no alternative decision would have been preferred by all members. The efficiency assumption is standard in many economic contexts, and has often been applied to household behavior. For instance, the axiomatic concepts of bargaining used in cooperative game theory typically assume efficiency, and noncooperative models of bargaining under symmetric information typically generate Pareto efficient outcomes. Among the alternative approaches that have been proposed in the literature, many, from simple dictatorship (possibly by a 'benevolent patriarch', as in Becker, 1974) to the existence of some household welfare function (as in Samuelson, 1956) assume or imply Pareto efficiency. In the same line, the 'equilibrium' approaches of Becker (1991) and Grossbard-Schechtman (1993), in which household members trade at existing market prices, typically generate efficient outcomes.

Natural as it seems for economists, the efficiency assumption nevertheless needs careful justification. Within a static context, this assumption amounts to the requirement that married partners will find a way to take advantage of opportunities that make both of them better off. Because of proximity and durability of the relation, both partners are aware of the preferences and actions of each other. They can act cooperatively by reaching some binding agreement. Enforcement of such agreements can be achieved through mutual care and trust, by social norms and by formal legal contracts. Alternatively, the agreement can be supported by repeated interactions, including the possibility of punishment. A large literature in game theory, based on several 'folk theorems', suggests that in such situations, efficiency should prevail.[10]

There are, however, two situations (at least) in which the efficiency as-

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

sumption may fail to apply. One is when existing social norms impose patterns of behavior that may conflict with efficiency. One example for apparent inefficiency is when, due to the traditional norms (as is still the case in some Muslim societies) the wife is expected to stay at home and the husband to work in the market. Although such a division of labor may have been efficient in the past, it certainly conflicts with efficiency in modern societies in which women are often more educated than their husband. Another example is the finding by Udry (1996) that households in Burkina-Faso fail to allocate inputs efficiently among various crops because of the socially imposed division of labor between genders, which implies that some crops can only be grown by a particular gender.[11] Secondly, some decisions are taken only once (or a few times), which implies that the repeated game argument does not apply; see Lundberg and Pollak (2003). Deciding whether to engage in a long training program or to move to another city, are not frequent decisions. The kind of 'equilibrium punishments' that are needed to implement efficient outcomes in repeated games may then be unavailable. The main theoretical insight here is that for medium or long-term decisions, efficiency requires commitment; conversely, any limitation of the members' ability to commit may result in inefficient outcomes. As we know, commitment within a household is only partial; for instance, it is impossible to commit not to divorce, although marriage contracts can be used to make divorce costly for one of the members (or both). For that reason, we will treat the issue of efficiency in a different manner depending upon whether we deal with a dynamic or static context. In most of the book the setting is static and efficiency is assumed. However, in chapters 6, 7, 8 and 12, which discuss saving and investment, we allow departures from efficiency.

In the remainder of this Chapter, we investigate the properties of models based on the collective approach. Before entering the technical analysis, however, one point should be stressed. The great advantage of the collective model is that we do not have to specify the (stable) mechanism that households use, but only assume that such a mechanism exists. In other words, the collective strategy does not make assumptions about decision processes (beyond efficiency); it rather lets the data tell us what these processes are. This naturally leads to two crucial questions. One is whether the efficiency assumption is testable; that is, whether empirically falsifiable predictions can be derived from it. The second question relates to identification: when (and under which assumptions) is it possible to recover the decision process from observable behavior? Answering the questions is a strong motivation for what follows.

[11] A program of research in economics tries to explain existing institutions (including social norms) as an efficient response to a particular context; the argument being that competition will tend to promote the 'best performing' institutions. However, when technology changes deviations from efficiency can arise, because social norms may change slowly.

5.3 Distribution factors

The generality of the collective approach comes at a cost. While the efficiency assumption restricts the set of possible allocations, we are still left with 'many' of them - a continuum, in fact. If we want to make more specific predictions on household behavior, additional information - more precisely, additional sources of variations - may be useful. Among the various factors that can influence household behavior, many have a 'direct' impact on either preferences or the budget constraint. A change in prices, wages and non labor income are likely to affect demands and labor supplies, simply because they modify the set of options available. A more subtle influence goes, indirectly, through the decision process. A change in the economic environment may not affect either the preferences or the budget opportunities but still have an impact on the decision process. This idea is incorporated into the collective model by introducing distribution factors. Any variable that has an impact on the decision process but affects neither preferences nor budget constraints is termed a *distribution factor*. In theory, a large number of variables fit this description. Factors influencing divorce, either directly (for example, the legislation governing divorce settlements and alimony payments) or indirectly (for example, the probability of remarriage, which itself depends on the number of available potential mates - what Becker calls the'marriage market') should matter, at least insofar as the threat (or the risk) of divorce may play a role in the decision process. Individuals' income or wealth can also be used as distribution factors. Suppose, for example, that earned and unearned income is given for any individual and let Y^s denote the total income of person s. Then total household income, given by $Y = Y^a + Y^b$, is all that matters for the budget constraint. For any *given* level of Y, the individual contribution of a to total income, measured for instance by the ratio $\frac{Y^a}{Y}$, can only influence the outcome through its impact on the decision process; it is thus a distribution factor.

In the collective framework, changes in distribution factors typically lead to variations in outcomes while the set of efficient allocations remains unchanged; as such, it provides very useful information on the decision process actually at stake in the household. For that reason, it is in general crucial to explicitly take then into account in the formal model. In what follows, therefore, $\mathbf{z}$ denotes a vector of distribution factors.

5.4 Modeling efficiency

5.4.1 The basic framework

The characterization of efficient allocations can follow the standard approach. The basic definition is that an allocation is Pareto efficient if making one person better off makes the other worse off:

Definition 2 *An allocation* $\left(\mathbf{Q}, \mathbf{q}^a, \mathbf{q}^b\right)$ *is* Pareto efficient *if any other allo-*

cation $(\bar{\mathbf{q}}^a, \bar{\mathbf{q}}^b, \bar{\mathbf{Q}})$ *that is feasible:*

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

and is such that $u^a\left(\bar{\mathbf{Q}}, \bar{\mathbf{q}}^a, \bar{\mathbf{q}}^b\right) > u^a\left(\mathbf{Q}, \mathbf{q}^a, \mathbf{q}^b\right)$ *must be such that* $u^b\left(\bar{\mathbf{Q}}, \bar{\mathbf{q}}^a, \bar{\mathbf{q}}^b\right) < u^b\left(\mathbf{Q}, \mathbf{q}^a, \mathbf{q}^b\right)$ *(and conversely).*

In practice the basic definition is not very tractable and we often use one of two alternative characterizations. A first characterization is:

Definition 3 *For any given vector* $(\mathbf{P}, \mathbf{p}, x, \mathbf{z})$ *of prices, total expenditure and distribution factors, an allocation* $\left(\mathbf{Q}, \mathbf{q}^a, \mathbf{q}^b\right)$ *is* Pareto efficient *if there exists a feasible* $\bar{u}^a$*, which may depend on* $(\mathbf{P}, \mathbf{p}, x, \mathbf{z})$*, such that* $\left(\mathbf{Q}, \mathbf{q}^a, \mathbf{q}^b\right)$ *solves the problem:*

$$\max_{\mathbf{Q}, \mathbf{q}^a, \mathbf{q}^b} u^b\left(\mathbf{Q}, \mathbf{q}^a, \mathbf{q}^b\right) \tag{38}$$

$$\textit{subject to } \mathbf{P}'\mathbf{Q} + \mathbf{p}'\left(\mathbf{q}^a + \mathbf{q}^b\right) \leq x \tag{39}$$

$$\textit{and } u^a\left(\mathbf{Q}, \mathbf{q}^a, \mathbf{q}^b\right) \geq \bar{u}^a\left(\mathbf{P}, \mathbf{p}, x, \mathbf{z}\right) \tag{40}$$

and $\bar{u}^a\left(\mathbf{P}, \mathbf{p}, x, \mathbf{z}\right)$ *is some given function of prices income and distribution functions .*

Thus the Pareto efficient allocation can be derived from maximizing the utility of one partner holding the utility of the other at a given level: among all allocations providing a with exactly $\bar{u}^a$, the efficient one(s) generate the maximum possible utility for b. It goes without saying that this approach - just like most microeconomics - should not be taken literally. No one believes that agents actually write and solve a program such as (38). Our claim is simply that when a decision process, whatever its exact nature, always lead to efficient outcomes, then for any choice of prices, income and distribution factors there exists a $\bar{u}^a$ such that the household behaves *as if* it was solving program (38).

The solution to (38), when it exists (that is, if $\bar{u}^a$ is feasible), depends on prices, total expenditure and on the arbitrary level $\bar{u}^a$; it can be denoted as $u^b = \Upsilon\left(\mathbf{P}, \mathbf{p}, x, \bar{u}^a\right)$.[12] The set of all pairs $\left(\bar{u}^a, \Upsilon\left(\mathbf{P}, \mathbf{p}, x, \bar{u}^a\right)\right)$ when $\bar{u}^a$ varies over a domain of feasible allocations for a is the set of all efficient allocations; it is also called the *Pareto frontier* or *utility possibility frontier*, UPF. Under the assumption that the utility functions are strictly concave it is straight forward to show that the function $\Upsilon(.)$ is strictly concave in $\bar{u}^a$. This allows us to write Program (38) in a different but equivalent way. Let μ denote the Lagrange multiplier for constraint (40); note that μ is always nonnegative. Then the program is equivalent to:

[12] Here we are drop distributions factors to make the exposition easier.

$$\max \mu u^a \left(\mathbf{Q}, \mathbf{q}^a, \mathbf{q}^b\right) + u^b \left(\mathbf{Q}, \mathbf{q}^a, \mathbf{q}^b\right) \tag{41}$$

under the constraint ((39)). The coefficient μ is known as the *Pareto weight* for a. That is, a Pareto efficient outcome always maximizes a weighted sum of the two individual utilities. A slightly modified form that keeps the formal symmetry of the problem is sometimes used:

$$\tilde{\mu} u^a \left(\mathbf{Q}, \mathbf{q}^a, \mathbf{q}^b\right) + (1 - \tilde{\mu})\, u^b \left(\mathbf{Q}, \mathbf{q}^a, \mathbf{q}^b\right), \tag{42}$$

where $\tilde{\mu} \in [0, 1]$. The Pareto weight plays a critical role in all that follows.

5.4.2 Ordinal versus cardinal representation of preferences

It is important to understand what, in the previous discussion, requires a *cardinal* representation of preferences, and what can be defined using only a standard, ordinal representation. The *set* of Pareto efficient *allocations* is an ordinal concept; it is not modified when u^s is replaced with $F(u^s)$ for a strictly increasing mapping $F(.)$. Under smoothness conditions, the set is one-dimensional, and therefore can be described by one parameter. However, the parametrization entails cardinality issues. For instance, a natural parametrization is through the weight μ. But μ depends on the particular cardinal representation that has been adopted for u^a and u^b: if u^s is replaced with $F(u^s)$, then the parameter μ corresponding to a particular efficient allocation has to be modified accordingly. Moreover, the convexity properties of the Pareto set are also of a cardinal nature. Assuming smooth preferences, for any given price-income vector, one can find cardinal representations of preferences such that the Pareto frontier is convex, linear or concave. In most of what follows, we adopt the convention of always using a strictly concave representation of utilities. In this case, the Pareto set is strictly convex. Indeed, for a given price-income vector, take two points $\left(\bar{u}^a, \bar{u}^b\right)$ and $\left(u'^a, u'^b\right)$ on the Pareto frontier, and let $\left(\mathbf{Q}, \mathbf{q}^a, \mathbf{q}^b\right)$ and $\left(\mathbf{Q}', \mathbf{q}'^a, \mathbf{q}'^b\right)$ be the corresponding consumption vectors. The vector

$$\left(\mathbf{Q}'', \mathbf{q}''^a, \mathbf{q}''^b\right) = \frac{1}{2}\left(\mathbf{Q}, \mathbf{q}^a, \mathbf{q}^b\right) + \frac{1}{2}\left(\mathbf{Q}', \mathbf{q}'^a, \mathbf{q}'^b\right)$$

satisfies the budget constraint, and by strict concavity,

$$u^s \left(\mathbf{Q}'', \mathbf{q}''^a, \mathbf{q}''^b\right) > \frac{1}{2} u^s \left(\mathbf{Q}, \mathbf{q}^a, \mathbf{q}^b\right) + \frac{1}{2} u^s \left(\mathbf{Q}', \mathbf{q}'^a, \mathbf{q}'^b\right) = \frac{1}{2}\bar{u}^s + \frac{1}{2} u'^s$$

for $s = a, b$. We conclude that the point $\frac{1}{2}\left(\bar{u}^a, \bar{u}^b\right) + \frac{1}{2}\left(u'^a, u'^b\right)$ belongs to the interior of the Pareto set.

Graphically, on Figure 4, the Pareto set is indeed strictly convex. We see that any point on the UPF can be defined either by its coordinate on the horizontal axis, here u^a, as in program (38), or by the slope of the Pareto frontier at that point, here μ as in program (41). Given that the UPF is strictly concave there is an increasing correspondence between $\bar{u}^a$ and μ: a larger $\bar{u}^a$ (or μ) corresponds to an allocation that is more favorable to a (hence less favorable for a). The slope of the tangent at any point on the UPF is given by $-\mu$. Note that the correspondence between $\bar{u}^a$ and μ is one-to-one; i.e., for any $\bar{u}^a$, there exists exactly one μ that picks up the efficient point providing a with exactly $\bar{u}^a$, and conversely for any μ there is only one allocation that maximizes (41) under budget constraint, therefore only one corresponding utility level $\bar{u}^a$.

Finally, we may briefly discuss two particular cases. One obtains when the cardinal representations of utilities are concave but not strictly concave. In that case, the UPF may include 'flat' (i.e., linear) segments (Figure 5). Then Program (38) is still equivalent to Program (41), but the relationship between $\bar{u}^a$ and μ is no longer one-to-one. It is still the case that for any $\bar{u}^a$, exactly one μ picks up the efficient point providing a with $\bar{u}^a$. But the converse property does not hold; i.e., to some values of μ are associated a continuum of utility levels $\bar{u}^a$; graphically, this occurs when $-\mu$ is exactly the slope of a flat portion of the UPF. This case is particularly relevant for two types of situations, namely transferable utility on the one hand (then the cardinalization is usually chosen so that the entire UPF is a straight line) and explicit randomization on the other hand.

The second particular case relates to local non differentiability of utility functions (Figure 6). Then the UPF may exhibit a kink, and the one-to-one relationship breaks down for the opposite reason - namely, many values of μ are associated with the same $\bar{u}^a$.

5.4.3 Stability and the Pareto weight

In what follows, we concentrate on deterministic decision processes. Then the stability axiom has a very simple implication - namely that in program (38), the coefficient $\bar{u}^a$ is a well-defined *function* of prices, income and possibly distribution factors, denoted $\bar{u}^a(\mathbf{P}, \mathbf{p}, x, \mathbf{z})$. It follows that, for given fundamentals and price-income bundle, the outcome of the decision process is such that the utility of a is $\bar{u}^a(\mathbf{P}, \mathbf{p}, x, \mathbf{z})$, and that of b is $\Upsilon(\mathbf{P}, \mathbf{p}, x, \bar{u}^a(\mathbf{P}, \mathbf{p}, x, \mathbf{z}))$. Note that under strict quasi concavity, these utility are reached for only one consumption bundle.

If, in addition, we adopt a strictly concave cardinalization of individual utilities, then the Pareto weight is also a well-defined function of prices, income and possibly distribution factors, denoted $\mu(\mathbf{P}, \mathbf{p}, x, \mathbf{z})$. For analytic tractability, we often add some structure to the problem by assuming that the function μ has convenient properties such as continuous differentiability. Such assumptions will be stated wherever they are needed.

In summary: under our two assumptions of stability and efficiency, using a strictly concave cardinalization of preferences, the behavior of the household can be modelled in a simple way; i.e., there exists a function $\mu(\mathbf{P}, \mathbf{p}, x, \mathbf{z})$ such that the household solves:

$$\max \mu\left(\mathbf{P},\mathbf{p},x,\mathbf{z}\right)u^{a}\left(\mathbf{Q},\mathbf{q}^{a},\mathbf{q}^{b}\right)+u^{b}\left(\mathbf{Q},\mathbf{q}^{a},\mathbf{q}^{b}\right) \tag{43}$$

under the budget constraint (39).

5.5 Pareto weights and 'power'

A major advantage of the formulation in (41) is that the Pareto weight has a natural interpretation in terms of respective decision *powers*. The notion of 'power' may be difficult to define formally, even in a simplified framework like ours. Still, it seems natural to expect that when two people bargain, a person's gain increases with the person's power. This somewhat hazy notion is captured very effectively by the Pareto weights. Clearly, if μ in (41) is zero then it is as though b is a dictator, while if μ is large then a effectively gets her way. A key property of (42) is precisely that increasing μ will result in a move along the Pareto set, in the direction of higher utility for a (and lower for b). If we restrict ourselves to economic considerations, we may thus consider that the Pareto weight μ 'reflects' a's power, in the sense that a larger μ corresponds to more power (and better outcomes) being enjoyed by a.

The empirical implications of this remark are quite interesting. For instance, when a reform is known or predicted to improve the relative situation of a particular member (say, paying some family benefits to the wife instead of the husband), we should find that the reform increases the member's Pareto weight. More generally, the intuitive idea that a specific distribution factor z is favorable to member a can readily be translated by the fact that μ is increasing in z. Conversely, we shall see later on that it is sometimes possible to recover the Pareto weights from a careful analysis of the behavior of the households at stake. Then one can find out which factors increase or decrease the power of each member - quite a fascinating outcome indeed.

Another important insight of the analysis is that, broadly speaking, *cooperation does not preclude conflict*. In other words, the Pareto efficiency assumption by no means implies that the members always agree on what to do. On the contrary, each agent will plausibly try to obtain that, among the continuum of possible, Pareto efficient outcomes, the one ultimately selected is favorable to her. In other words, who gets what is a crucial but difficult and potentially conflictual issue, that the efficiency assumption leaves totally open. It can be resolved in a number of different ways - bargaining, legally binding contracts, tradition, social norms or less formal ways that reflect the feelings of the two partners towards each other. Pareto efficiency does not preclude any of these aspects; it just imposes that whichever solution is found, no money is ultimately left on the bargaining table. In a sense, the collective approach provides the tools needed to concentrate on the interesting issue of who gets what - or, technically, what do the Pareto weights look like as functions of prices, income and distribution factors.

5.6 Household utility

If the Pareto weights are, in fact, constant (that is, not a function of prices and income) then we have a unitary model and we can define a household utility function as a function of household public and private goods. It turns out that for the collective model we can also define a household utility function over household purchases of public and private goods but this function has one extra argument as compared to the unitary model. We define the *household utility function* by:

$$\begin{aligned} & u^h\left(\mathbf{Q},\mathbf{q},\mu\left(\mathbf{P},\mathbf{p},x,\mathbf{z}\right)\right) = \\ \max_{\mathbf{q}^a,\mathbf{q}^b} & \left\{\mu\left(\mathbf{P},\mathbf{p},x,\mathbf{z}\right)u^a\left(\mathbf{Q},\mathbf{q}^a,\mathbf{q}^b\right)+u^b\left(\mathbf{Q},\mathbf{q}^a,\mathbf{q}^b\right)\right\} \\ & \text{subject to } \mathbf{q}^a+\mathbf{q}^b=\mathbf{q} \end{aligned} \tag{44}$$

With this definition of the household utility u^h, program (41) is equivalent to the maximization of u^h under the budget constraint. This looks a lot like standard utility maximization in a unitary model. However, the critical feature of this household utility function is that it depends on the Pareto weight $\mu\left(\mathbf{P},\mathbf{p},x,\mathbf{z}\right)$. This remark is important for two reasons. First, it explains why an efficient household needs not (and will not in general) behave like an individual: since the utility u^h is price-dependent, the demand derived from its maximization under budget constraint needs not (and will not in general) satisfy the standard conditions of consumer theory. Secondly, while the idea of introducing prices into the utility function is an old one, the important feature in our case is how it is done. Following standard demand theory we do not allow prices to enter the individual utility functions; *prices and income can only affect the respective weights* given to individual utilities. As we shall see below, this gives very specific predictions for household demands. Additionally, it makes analysis using a collective model almost as easy as using a unitary model which is an important consideration when considering non-unitary alternatives.

This approach allows us to decompose changes in the utility levels of the two partners following a change in the environment into changes that would follow in a unitary model and the additional effect due to the collective framework. This is illustrated in figure 7, where we ignore distribution factors. Here we consider a change in prices and incomes that moves the UPF from $UPF\left(\mathbf{P},\mathbf{p},x\right)$ to $UPF\left(\mathbf{P}',\mathbf{p}',x'\right)$. Initially the point I is chosen on $UPF\left(\mathbf{P},\mathbf{p},x\right)$. If we hold μ constant when prices and income change (the unitary assumption) then the utility levels move to point II at which point the tangent to $UPF\left(\mathbf{P}',\mathbf{p}',x'\right)$ is parallel to the tangent at point I on $UPF\left(\mathbf{P},\mathbf{p},x\right)$. However, a change in the economic environment may also lead to a change in the Pareto weight. This is the 'collective' effect, illustrated by the move around $UPF\left(\mathbf{P}',\mathbf{p}',x'\right)$ from II to III.

Finally, the collective formalization provides a natural way of introducing distribution factors within the framework of household decision process. If some distribution factors $\mathbf{z}$ influence the process by shifting the individual weights, then μ will depend on these variables. The fact that distribution factors matter only through their impact on μ plays a key role in the results of Chapter 4. As we shall show, efficiency can be tested using cross equation restrictions that arise from the fact that the same function $\mu(\mathbf{P}, \mathbf{p}, x, \mathbf{z})$ appears in the demand for all goods. Moreover, there is an important difference between prices and total income, on the one hand, and distribution factors on the other hand. A change in prices or total income will affect not only the weight μ, but also the shape of the Pareto set; hence it final impact on individual welfare may be difficult to assess. On the contrary, a change in a distribution factor can by definition *only influence the weight* μ.In general, its effect on welfare is not ambiguous. In terms of figure 7 a distribution factors shift the tangent point but not the frontier itself.

As an illustration of this point, we may briefly come back to the example discussed in subsection 4.2 of the impact of individual incomes Y^a and Y^b on household behavior. From a collective perspective, this impact should be decomposed into two components. One is the resulting change in total income $Y = Y^a + Y^b$ (hence of total expenditures x in our static framework); this affects the shape of the Pareto frontier as well as the weight μ, and its effect is *a priori* ambiguous. The second component is the change in *relative* incomes, say $z = \frac{Y^a}{Y^b}$, keeping the sum constant. The latter should be analyzed as a variation of a distribution factor, and its consequences are much easier to assess. For instance, if we assume, as is natural, that increasing the relative income of a increases a's weight, then it must increase a's welfare. However, how this improvement in a's situation will be translated into observable household behavior (for example, which consumptions will increase) is a difficult issue, for which a more precise formalization is needed.

Finally, let us mention that one could define a 'household utility' as some function $W\left(u^a, u^b, \mathbf{P}, \mathbf{p}, x, \mathbf{z}\right)$ which is increasing in its first two arguments. One can readily check that, for any given value of $(\mathbf{P}, \mathbf{p}, x, \mathbf{z})$, the maximization of W under the budget constraint will always generate Pareto efficient outcomes; therefore, this formalization would be equivalent to the previous ones, although the details of this equivalence may be somewhat tricky. Again, note that this household utility will, in general, depend on prices, incomes and distribution factors. As such, it is much more general than the unitary version (20).

5.7 Caring

The way in which partners care about each other may also affect the Pareto utility frontier. To take a simple example, consider the caring preferences introduced in 1:

$$\begin{aligned} U^a\left(\mathbf{Q}, \mathbf{q}^a, \mathbf{q}^b\right) &= u^a\left(\mathbf{Q}, \mathbf{q}^a\right) + \delta^a u^b\left(\mathbf{Q}, \mathbf{q}^b\right) \\ U^b\left(\mathbf{Q}, \mathbf{q}^a, \mathbf{q}^b\right) &= u^b\left(\mathbf{Q}, \mathbf{q}^b\right) + \delta^b u^a\left(\mathbf{Q}, \mathbf{q}^a\right) \end{aligned} \tag{45}$$

The maximand is now

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

which can also be written as:

$$\left(\mu+\delta^b\right)u^a\left(\mathbf{Q},\mathbf{q}^a\right)+\left(1+\mu\delta^a\right)u^b\left(\mathbf{Q},\mathbf{q}^b\right)=\mu' u^a\left(\mathbf{Q},\mathbf{q}^a\right)+\mu' u^b\left(\mathbf{Q},\mathbf{q}^b\right) \tag{46}$$

Assuming that $(\mu\delta^a+1)>0$ we can then represent household preferences by:

$$\mu' u^a\left(\mathbf{Q},\mathbf{q}^a\right)+u^b\left(\mathbf{Q},\mathbf{q}^b\right) \tag{47}$$

where

$$\mu'=\frac{\mu+\delta^b}{1+\mu\delta^a}$$

Formally, (47) is identical to the egotistic case ($\delta^a=\delta^b=0$), indicating that *any allocation that is Pareto efficient for the caring preferences is also Pareto efficient for the egotistic ones*. The argument underlying this conclusion is quite general, and goes as follows: if an allocation fails to be efficient for egotistic preferences, there exist another allocation that entails higher values of both u^a and u^b. But then it also entails higher values of both U^a and U^b, showing that the initial allocation was not efficient for caring preferences as well. In other words, *the Pareto set for caring preferences is a subset of the Pareto set for egotistic preferences*. Note, however, that the converse is not true: some allocations may be efficient for egotistic preferences, but not for caring ones. Indeed, an allocation that gives all resources to one member may be efficient for egotistic agents, but not for caring persons - a redistribution in favor of the'poor' member would then typically be Pareto improving. Technically, when μ varies from 0 to infinity, μ' only varies from δ^b to $1/\delta^a$, and the new Pareto set is a strict interior subset of the initial one.

A variant of this is if the two partners care for each other in the following way:

$$\begin{aligned}U^a\left(\mathbf{Q},\mathbf{q}^a,\mathbf{q}^b\right)&=\min\left\{u^a\left(\mathbf{Q},\mathbf{q}^a\right),u^b\left(\mathbf{Q},\mathbf{q}^b\right)\right\}\\U^b\left(\mathbf{Q},\mathbf{q}^a,\mathbf{q}^b\right)&=\min\left\{u^a\left(\mathbf{Q},\mathbf{q}^a\right),u^b\left(\mathbf{Q},\mathbf{q}^b\right)\right\}\end{aligned} \tag{48}$$

This formalizes the maxim that 'no man can be happier than his wife'. In this very special case the utility possibility frontier shrinks to a single point at which $u^a=u^b$.

5.8 Children

Finally, let us briefly come back to the distinction sketched above between children being modeled as public goods or genuine decision makers. In the first case, using parental utilities of the form $u^s + \kappa^s u^k$ described above, the maximand in (41) becomes

$$\mu\left(u^a + \kappa^a u^k\right) + \left(u^b + \kappa^b u^k\right)$$

which is equivalent to

$$\frac{1}{1 + \mu\kappa^a + \kappa^b + \mu}\left[\mu u^a + u^b + \left(\mu\kappa^a + \kappa^b\right) u^k\right] \tag{49}$$

the initial fraction in (49) being needed to maintain the normalization that the weights sum to unity.

Alternatively, we may model the child as a decision maker. Then (s)he is characterized by an additional Pareto weight, say μ^k, and the household maximizes the weighted sum:

$$\mu u^a\left(\mathbf{Q}, \mathbf{q}^a, \mathbf{q}^b\right) + u^b\left(\mathbf{Q}, \mathbf{q}^a, \mathbf{q}^b\right) + \mu^k u^k \tag{50}$$

Although the two forms (49)and (50) look similar, they are, in fact, quite different. Recall that the key insight of collective models is that Pareto weights may depend on prices, wages, incomes and distribution factors, and that this fact explains why collective households do not generally behave as unitary ones. In (49) all Pareto weights are defined by the knowledge of the function μ ; in (50), however, μ and μ^k can be defined independently, and the location of the final outcome on the Pareto frontier now depends on *two* parameters. Broadly speaking, the deviation from the unitary model is one dimensional in the first case (it is summarized by a unique function μ) whereas it is two-dimensional in the second case. As it turns out, this distinction has testable implication; that is, we shall see later on that a household with three decision makers does *not* generally behave as a couple, pretty much in the same way as couples do not generally behave as singles. Another fascinating implication is that, in principle, *one can assess the number of actual decision makers in a household from the sole examination of the household's behavior*, even in a fairly general context!

5.9 The unitary model as a special case

It is clear, from the discussion above, that the unitary model is a special case of the collective framework. An obvious example obtains when the household utility (or the Pareto weight μ) does not depend directly on prices, incomes and distribution factors. As a matter of fact, the unitary assumption is far and away the most common assumption in the modelling of household decisions. This, however, is certainly due to its very great convenience, rather than any intuitive plausibility. If one is to take seriously the idea of a decision process actually taking place between the members, it hard to believe that neither prices (including respective wages), nor respective incomes, nor any exterior factor will influence the 'weights' of individual agents in the decision process.

Nonetheless there are circumstances under which the household will act as though it has a single utility function. One obvious example is if custom or strong social traditions give all the power to one person (usually the husband) in the household.

An alternative is given in Samuelson (1956). Samuelson considers the family to be the fundamental unit on the demand side of the economy. However, because such a unit consists of several members, we cannot expect a consensus (that is, consistent family indifference curves). He recognizes that preferences within a family are interrelated and external consumption effects (*a la* Veblen and Dusenberry) are the "essence of family life". Nevertheless, if such external effects are put aside, and a restricted form of altruism is assumed, families may behave *as if* they maximize a single social utility. In particular Samuelson considers a common social welfare function for the family that is restricted to depend on the individual consumptions of family members only through the preferences of these members. This restriction, together with the assumption of no external consumption effects and no public goods, implies that all family decisions can be decentralized via a distribution of income.[13] The important point is the distribution of income depends on prices and income and will not be constant. Thus schemes such as a receives 60% of income and b receives 40% can never be consistent with the maximization of a family SWF. The main result that Samuelson provides is that if income is redistributed so as to maximize a given social welfare function, then the family aggregate consumptions will satisfy the Slutsky conditions. That is a family will act in the same manner as a single person.

Becker (1991) criticizes Samuelson for not explaining how a social welfare function arises. In the context of moral judgements, each person can have a private utility that is defined on outcomes affecting them directly, and a social utility function that reflects preferences on the outcomes for all family members. So it is unclear how partners agree on a single common social welfare function. One mechanism suggested by Becker is that one person, the 'head', has most of the family resources and is sufficiently altruistic that they will transfer resources to the other member. If the dependents' consumption is a normal good for the head, all family members will align their actions with the head's preferences, as any improvement in the income under the command of the head raises their utilities. It is then the case that the family as a group acts as if a single objective is being maximized. This is the Rotten Kid Theorem mechanism outlined in subsection 4.3. In that noncooperative voluntary contributions model, one of the partners may effectively be a dictator if they control most (but not necessarily all) of household resources. In that case a unitary model obtains locally if one partner is wealthier and they are the sole contributor to the public good.

Another important case is when the preferences display *transferable utility*

[13] If we have public goods and externalities then we also need Lindahl prices and Pigouvian taxes to decentralize.

(TU); see subsection 1. Indeed, under TU, program (41) becomes

$$\max \mu \left(f^a \left(q_2^a, ..., q_n^a\right) + F^a \left(\mathbf{Q}\right) + G \left(\mathbf{Q}\right) q_1^a\right) + \left(f^b \left(q_2^b, ..., q_n^b\right) + F^b \left(\mathbf{Q}\right) + G \left(\mathbf{Q}\right) q_1^b\right) \tag{51}$$

under the budget constraint 39. The first surprising feature of the TU assumption is that if the optimum has q_1^a and q_1^b both positive, then μ is necessarily equal to unity. To see this, set the price of the first good to unity and substitute for q_1^b using the budget constraint:

$$\max \mu \left(f^a \left(q_2^a, ..., q_n^a\right) + F^a \left(\mathbf{Q}\right) + G \left(\mathbf{Q}\right) q_1^a\right) + \left(f^b \left(q_2^b, ..., q_n^b\right) + F^b \left(\mathbf{Q}\right) + G \left(\mathbf{Q}\right) \left(x - \mathbf{P}'\mathbf{Q} - \mathbf{p}'_{-1} \left(\mathbf{q}_{-1}^a + \mathbf{q}_{-1}^b\right) - q_1^a\right)\right) \tag{52}$$

where $\mathbf{p}_{-1}$ denotes the vector of prices for all goods except the first (and similarly for $\mathbf{q}_{-1}^s$). Taking the derivative with respect to q_1^a we see that:

$$\mu G \left(\mathbf{Q}\right) - G \left(\mathbf{Q}\right) = 0 \tag{53}$$

which implies $\mu = 1$ so that the UPF is a line with a constant slope of -1. Thus the Pareto weight cannot depend on prices, income or any distribution factors. Therefore the partners will always agree to act in a manner which shifts the frontier out as far as possible by the choice of $\left(\mathbf{Q}, \mathbf{q}_{-1}^a, \mathbf{q}_{-1}^b\right)$. In fact they will agree to maximize the sum of their individual utilities given by:

$$f^a \left(q_2^a, ..., q_n^a\right) + f^b \left(q_2^b, ..., q_n^b\right) + F^a \left(\mathbf{Q}\right) + F^b \left(\mathbf{Q}\right) + G \left(\mathbf{Q}\right) \left(x - \mathbf{P}'\mathbf{Q} + \mathbf{p}'_{-1} \left(\mathbf{q}_{-1}^a + \mathbf{q}_{-1}^b\right)\right) \tag{54}$$

Thus, under transferable utility and assuming efficiency, married partners will agree on almost all consumption choices. The only conflict will be in how to divide the private good q_1 which is often referred to as 'money' but in the family context it may be interpreted more broadly as a medium of exchange.

5.10 The Rotten Kid Theorem revisited

As we have just seen, under transferable utility and efficiency, a couple acts as a single decision unit in the sense that both partners would agree on the set of actions which maximizes the *joint* marital output, defined as the sum of the partners's individual utilities. In contrast to the case of dictatorship, where the issue of implementation does not arise, for the case of transferable utility we also need to ask how the actions that maximize the joint output are actually enforced. One possibility is that bargaining takes place at the outset of marriage, and some sort of binding agreement is signed and then carried out. However, if the partners are altruistic towards each other, these emotional ties generate commitments that can replace legally binding contracts. In particular, commitments that arise from altruistic preferences can be exploited in the design of a mechanism that implements the maximization of the total output (sum of utilities) and is self enforcing. One such scheme (see Becker, 1974) is to select a principal (a family head) who is given control over family resources and can make transfers as she or he sees fit. The only requirement is that the principal should care about all family members in the sense that their utilities enter his or her own preferences as normal goods. Once this scheme is put in place, each person is allowed to choose their own actions selfishly. It had been observed by Becker that such a mechanism is efficient and each participant voluntarily acts in the interest of the group. The reason is that any productive action which increases total output is rewarded by an increased transfer from the principal. Conversely, any destructive action is punished by reduced transfers. In this way the interests of the group are internalized by every member. Although the allocation of income depends on who is the head, family decisions will be invariant to his or her preferences. The crucial aspect is that every partner should trust the principal to truly care about all family members and that the principal should be able to fully control the distribution of income (in the sense that his resources are such that he gives everyone a positive transfer that can be reduces or increased).[14]

To illustrate the working of the 'family head mechanism', let each spouse have two private actions: consumption and work. Time not spent at work is used to produce a household good which is a public good (for example, child quality). Let us assume transferable utility and write the person specific utility as

$$U^s(Q, q^s, l^s) = Qq^s + v^s(l^s), \;\; s = a, b \tag{55}$$

where where, q^s denotes private consumption, l^s is leisure time and Q is a public good produced at home. The household production function is

$$Q = \phi(t^a, t^b) \tag{56}$$

where t^s denotes time spent by s in home on the production of the public good. The the family budget constraint is

$$q^a + q^b = \;(1 - t^a - l^a)w^a + (1 - t^b - l^b)w^b, \tag{57}$$

where w^s is the wage of person s. Applying the results on transferable utility, it is easy to verify that any Pareto efficient allocation must maximize the sum of private utilities given by

$$\pi = [(1 - t^a - l^a)w^a + (1 - t^b - l^b)w^b]\phi(t^a, t^b) + v^a(l^a) + v^b(l^b) \tag{58}$$

[14] Becker has two slightly different versions of the Rotten Kid Theorem. The early one stated is Becker (1974, page 1080) is "If a head exists, other family members are also motivated to maximize family income and consumption, even if their utility depend on their consumption alone " The later version in Becker (1991, p288) is set in context of mutual altruism where each person is a potential contributor to the other and states that "Each beneficiary, no-matter how selfish, maximizes the family income of his benefactor and thereby internalizes all effects of his actions on other beneficiaries." In both versions, there is only one good that is distributed. Following Bergstrom (1989) we consider here a problem with two goods and show that under transferable utility similar results are obtained.

To show that this is an equilibrium outcome of the 'family head mechanism', we consider a two stage game, such that in the first stage each agent s chooses independently the amount of work at home, t^s and in the market $1 - t^s - l^s$. In the second stage, the head, say partner a, chooses the level of the private good, q^s that each partner receives based on a's *social* preferences, $W^a(U^a(Q, q^a, l^a), U^b(Q, q^b, l^b))$. We can solve this problem backwards. In the last stage, the levels of work at home, t^s and l^s are given to a and she can only transfer private goods. This means that the head faces a *linear* Pareto frontier (see Figure) and will select the best point for her on this frontier. Now assume that the two private utilities appear as 'normal goods' in a's social welfare function so that whenever the Pareto frontier shifts up (down) the head reallocates private goods to raise (reduce) the private utilities of *both* agents. Anticipating that, each agent who chooses actions selfishly in the first stage will realize that their private utility is a monotone increasing function of the total resources available for the head for redistribution (π in equation (58)). Therefore, each agent will choose the actions under their control to maximize the pie and the outcome is the same efficient outcome that would arise if the head could directly control all family decisions.

The family head mechanism was first proposed by Becker and is discussed in detail in Becker (1991, chapter 8). One application of the analysis is parent child relationship and the main result is that selfish children can act in a manner that internalizes the consequences of their actions, yielding an efficient outcome. This result was coined by Becker as 'the rotten kid theorem'. His analysis, however, was much more general, dealing with various forms of altruism and preference dependence. The subsequent literature addressed the generality of the efficiency head mechanism. Bergstrom (1991) shows that this result generally fails in the absence of transferable utility, because agents can then affect not only the location of the Pareto frontier but also its slope, destroying the monotonicity result required for the theorem to hold. Another issue is the precise sequence of events. Suppose that the children can consume in both periods 1 and 2. Then, efficiency requires that, for each child, the ratio of the marginal utilities of consumption in the two periods is equated to the cost of transferring goods over time that is facing the household, $1 + r$. However, in choosing consumption, the child will take into account that his first period consumption also influences the transfer from the head. Being poor in the second period causes the parent to transfer more, causing the child to under-save. This pattern of behavior, where giving leads to under provision, is referred to as the Samaritan Dilemma (Bruce and Waldman, 1990). This example shows that altruism can also be a constraint on mechanism design. The parent could in principle impose the efficient outcome by conditioning the payment on past performance. However, an altruistic parent cannot commit to punish a deviating child. This restriction is captured in modeling the stages of game and seeking a subgame perfect equilibrium.

5.11 Bargaining models[15]

Throughout the chapter, we have stressed that the collective model, in its fully general version, is agnostic about the specific decision process provide that the latter generates Pareto efficient outcomes. Because of this generality, it is thus compatible with a host of more specific models that further specify the way a particular point on the Pareto frontier is selected. For instance, we shall show in detail in chapter 8 that under some conditions, this choice can be fully determined by the competition in the marriage market, where considerations such as what are the individual characteristics that generate marital surplus, what is the matching process and does a person have a close substitute play a crucial role. However, much of the literature pursues a more partial view and concentrates on the relative strengths of two individuals who are already matched and use tools from cooperative game theory to derive the bargaining outcome.[16]

Any bargaining model requires a specific setting: in addition to the framework described above (two agents, with specific utility functions), one has to define a *threat point* T^s for each individual s. Intuitively, a person's threat point describes the utility level this person could reach in the absence of an agreement with the partner. Then resources are allocated between public and private consumption, resulting in two utility levels $\bar{u}^a$ and $\bar{u}^b$. Typically, bargaining models assume that the outcome of the decision process is Pareto efficient. Bargaining theory is used to determine how the threat points influence the location of the chosen point on the Pareto frontier. Clearly, if the point $T = \left(T^a, T^b\right)$ is outside of the Pareto set, then no agreement can be reached, since at least one member would lose by agreeing. However, if T belongs to the interior of the Pareto set so that both agents can gain from the relationship, the model picks a particular point on the Pareto utility frontier.

Before describing in more detail some of the standard solutions to the bargaining problem, however, it is important to note that the crucial role played by threat points - a common feature of all bargaining models - has a very natural interpretation in terms of distribution factors. Indeed, *any variable that is relevant for threat points only is a potential distribution factor.* For instance, the nature of divorce settlements, the generosity of single parent benefits or the probability of re-marriage do not directly change a household's budget constraint (as long as it does not dissolve), but may affect the respective threat points of individuals within it. Then bargaining theory implies that they will influence the intrahousehold distribution of power and, ultimately, household behavior. This intuition is perfectly captured in the collective framework by the idea that the Pareto weight depends on distribution factors. Moreover, it provides a clear idea of the direction of these effects. That is, a variable that ameliorates the wife's threat point should

[15] For a more complete discussion of two person bargaining, see Myerson (1991, ch.8).

[16] Bargaining approaches to household decision making were first introduced by Manser and Brown (1980) and McElroy and Horney (1981).

always positively affect her Pareto weight. These notions potentially provide a number of powerful tests, which are moreover independent of the particular bargaining concept at stake.

5.11.1 Nash bargaining

The most commonly used bargaining solution was proposed by John Nash in the early 1950's. Nash derived this solution as the unique outcome of a set of axioms that any "reasonable' solution must satisfy. Some of the axioms are uncontroversial. One is individual rationality: an agent will never accept an agreement that is less favorable than her threat point. Another is Pareto efficiency, as discussed above. A third mild requirement is invariance with respect to affine transformations[17]: if both the utility and the threat point of an agent are transformed by the same increasing, affine mapping the prediction about the equilibrium outcome of cooperation does not change. Note, however, that a non linear transform will change the outcome; that is, Nash bargaining requires a *cardinal* representation of preferences.

The last two axioms are more specific. One is symmetry; it states that if utilities and threat points are permuted between members (u^a and T^a are replaced with u^b and T^b, and conversely) then the outcomes are simply switched ($\bar{u}^a$ is replaced with $\bar{u}^b$ and conversely). Natural as it may sound, this assumption may still sometimes be too strong. In many socioeconomic contexts, for instance, male and female roles are by no means symmetric. Fortunately, Nash bargaining can easily be extended to avoid the symmetry assumption.

The last and crucial axiom is independence. It can be stated as follows. Assume that the set of available opportunities (the Pareto set) shrinks, so that the new Pareto set is within the old one, but the initial equilibrium outcome is still feasible; then the new equilibrium outcome will be the same as before. In other words, the fact that one member misses some opportunities that he had before does not affect his bargaining position towards the other member. This requirement alone implies that the Nash solution maximizes some function of the utilities of the two partners.

If one accepts these axioms, then only one outcome is possible. It given by the following rule: find the pair $\left(\bar{u}^a, \bar{u}^b\right)$ on the Pareto frontier that maximizes the product $\left(u^a - T^a\right)\left(u^b - T^b\right)$. That is, the Nash bargaining allocation $\left(\mathbf{Q}, \mathbf{q}^a, \mathbf{q}^b\right)$ solves

$$\max_{\mathbf{Q},\mathbf{q}^a,\mathbf{q}^b} \left(u^a\left(\mathbf{Q}, \mathbf{q}^a, \mathbf{q}^b\right) - T^a\right)\left(u^b\left(\mathbf{Q}, \mathbf{q}^a, \mathbf{q}^b\right) - T^b\right) \tag{NB}$$

under the budget constraint 40. Thus the product $\left(u^a - T^a\right)\left(u^b - T^b\right)$ can be considered as a household utility function, that is maximized on the Pareto set. Note that $\left(u^s - T^s\right)$ is the surplus derived from the relationship by agent s. The main implication of Nash bargaining is that the product of surpluses should be maximized.[18]

[17] An affine mapping is of the form $f(x) = ax + b$.

[18] Note that simply maximizing the sum of surpluses, $\left(u^A - T^A\right) + \left(u^B - T^B\right)$, would violate the invariance axiom.

Clearly, if the threat points do not depend on prices, incomes and distribution factors, the maximand can be seen as a standard, unitary utility and the Nash bargaining solution boils down to a unitary model; the outcome satisfies in particular the properties of a regular consumer demand. This case, however, is of little interest. Typically, threat point depends on a number of parameters, and the previous formalization allows us to study how these effects translate into behavioral patterns. An important result is that at the Nash bargaining equilibrium $\left(\bar{u}^a, \bar{u}^b\right)$, $\bar{u}^a$ is increasing in T_a and decreasing in T_b (while, obviously, $\bar{u}^b$ is decreasing in T_a and increasing in T_b). Hence, any change that increases a member's threat point without changing the Pareto frontier (the typical impact of a distribution factor) will ameliorate this member's situation.

Finally, the symmetry axiom can be relaxed. Then the general form is a straightforward generalization of the previous one: instead of maximizing the sum of log surpluses, one maximizes a weighted sum of the form $\gamma_a \log\left(u^a - T^a\right) + \gamma_b \log\left(u^b - T^b\right)$. In this form, the weights γ_s introduce an asymmetry between the members' situations.

5.11.2 Kalai-Smorodinsky

An alternative concept has been proposed by Kalai and Smorodinsky (1975). It relies on the following, *monotonicity* property. Consider two bargaining problems such that (i) the range of individually rational payoffs that player a can get is the same in the two problems, and (ii) for any given, individually rational utility level for player a, the maximum utility that player b can achieve (given the Pareto frontier) is never smaller in the second problem than in the first. Then player b should do at least as well in the second problem than in the first. In other words, if one enlarges the Pareto set by inflating b's opportunities while keeping a's constant, this change cannot harm b.

Kalai and Smorodinsky prove that there exists a unique bargaining solution that satisfies all the previous axioms except for independence, which is replaced with monotonicity. Moreover, the solution has an interesting interpretation. Define the *aspiration level* A^s of player s as the maximum utility he/she can get that is compatible with feasibility and the partner's individual rationality constraint; this corresponds to the point on the Pareto frontier that leaves the partner, say s', at their threat point utility $T^{s'}$. Define, now, the *ideal point* as the point $\left(A^a, A^b\right)$; obviously, this point lies outside of the Pareto frontier. The solution, now, is to chose a point $U = \left(u^a, u^b\right)$ on the Pareto frontier so that

$$\frac{u^a - T^a}{u^b - T^b} = \frac{A^a - T^a}{A^b - T^b}$$

In words, the bargaining is here influenced, in addition to the threat points, by the players' aspirations about what they might receive within marriage. The surplus share received by player s, $u^s - T^s$, is directly proportional to the maximum gain s could aspire to, $A^s - T^s$.

5.11.3 Non cooperative foundations

Finally, an on going research agenda, initially proposed by Nash, is to provide noncooperative foundations to the bargaining solutions derived from axioms. The most influential framework is the model of Rubinstein (1982), in which players make alternating offers until one is accepted. When time matters through a constant discount factor, there exists a unique, subgame-perfect equilibrium of this noncooperative game, which is characterized by the requirement that each player should be indifferent between accepting the current offer and waiting to an additional round and making an offer that the opponent would accept. Binmore (1987) has shown that when the discount factor is close to 1, the outcome of the subgame perfect equilibrium is close to the Nash bargaining solution. Moreover, the asymmetric Nash bargaining solution can be justified by introducing differences in discount factors of the two bargaining agents.

5.11.4 Empirical content of bargaining models

Since the bargaining models just described all assume (or imply) Pareto efficiency, their solutions will satisfy the general properties generated by the collective model; these will be detailed in the next Chapter. But do these models allow us to go one step further? That is, which additional insights (if any) can we obtain from the use of bargaining concepts?

The answer to that question depends on what is known on the threat points. Indeed, a first result (Chiappori and Donni 2007) is that *any* Pareto efficient allocation can be derived as the Nash bargaining solution for an *ad hoc* definition of the threat points. Hence the additional information provided by the bargaining concepts (with respect to the sole efficiency assumption) must come from specific hypotheses on the threat points - that is, on what is meant by the sentence: 'no agreement is reached'.

Several ideas have been used in the literature. One is to refer to divorce as the 'no agreement' situation. Then the threat point is defined as the maximum utility a person could reach after divorce. Such an idea seems well adapted when one is interested in the effects of laws governing divorce on intrahousehold allocation. Another interesting illustration would be public policies such as single parent benefits, or the guaranteed employment programs that exist in some Indian states; Kanbur and Haddad convincingly argue that the main impact of the program was to change the opportunities available to the wife outside marriage (or cohabitation). It is probably less natural when minor decisions are at stake: choosing who will walk the dog is unlikely to involve threats of divorce.[19]

A second idea relies on the presence of public goods, and the fact that noncooperative behavior typically leads to inefficient outcomes. The idea, then,

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

is to take the non-cooperative outcome as the threat point: in the absence of an agreement, both members provide the public good(s) egotistically, not taking into account the impact of their decision on the other member's welfare. This version captures the idea that the person who would suffer more from this lack of cooperation (the person who has the higher valuation for the public good) is likely to be more willing to compromise in order to reach an agreement. Interestingly, in this context some of the results derived in the non-cooperative case extend to the cooperative context as well. For instance, the income pooling result for interior solutions, derived in subsection 4.4, applies here as well: total income being kept constant, a change in respective incomes that does not affect the noncooperative consumption pattern leaves the threat point unchanged and hence has no impact on the bargaining outcome. Thus local income pooling is inherited by the bargaining solution.

Finally, it must be remarked that assumptions on threat points tend to be strong, somewhat *ad hoc*, and often not independently testable. Given this cost, models based on bargaining should be used parsimoniously, and preferably when there is good evidence that the actual structure of the decision process is close to what is implicitly assumed by the concept referred to. An alternative approach is to concentrate on more general assumptions, the implications of which hold for a large class of models. Efficiency is one natural example. Another is that some distribution factors, whatever the distribution process, can only be favorable to one partner (hence unfavorable to the other) - an intuition that can often be documented using sociological or ethnographic studies. This point should be kept in mind for the next chapters.

5.12 Other approaches

Finally, we may briefly review three approaches that have been proposed for analyzing household and family behavior. Two of them (the equilibrium models of Grossbard-Shechtman and Haller, and the 'separate spheres' model of Lundberg and Pollak) lead to efficient outcomes, therefore are consistent with the cooperative/collective model; the third (Basu's 'inefficient bargaining') is not, although it relies on a bargaining framework.

5.12.1 Equilibrium models

Following the seminal contributions of Becker[20], several papers by Grossbard-Schechtman[21] analyze marriage in a general equilibrium framework, in which intrahousehold allocation are directly driven by the competitive constraints that exist on the marriage market. In some of these models, women's role is essentially to produce domestic commodities. Men employ women to produce for them, and compensate them with transfers (which, in developing societies, may take the form of provision of basic needs) and/or non-pecuniary benefits. From this perspective, marriage can essentially be analyzed as an employment relationship, which allows to apply the standard concepts of labor economics.

[20] See Becker (1991) for a general overview.

[21] See Grossbard-Schechtman (1993) for a unified presentation.

The framework is then generalized to situations where both men and women engage in household production work. In all these models, the emphasis is put on a general equilibrium analysis, and specifically on the impact of the economy on intrahousehold decisions. One may remark, at this stage, that the outcome of the decision process thus described is efficient; therefore these models belong to the cooperative/collective family.[22]

In a related line, Gerbach and Haller (1999) and Haller (2000) study the general equilibrium implications of competitive exchange among multimember households, in a context in which consumptions are exclusively private but consumption externalities may exist within the household. They compare two benchmark cases: one in which decision making within households is always efficient (therefore households can be described using the collective representation), and one ('individual decentralization') in which each household member 'goes shopping on his or her own, following his or her own interests, after receiving a share of household income' (Haller, 2000, p.835). They first analyze whether competitive exchange among efficient households lead to a Pareto-optimal allocation at the global level. The answer is positive as long as each household's demand exhausts its budget.[23] They then ask whether such an optimal allocation can be 'individually decentralized' in the sense just defined. They show that, generically on preferences, the answer is now negative; they conclude that some specific household decision processes are needed to internalize the externalities.

5.12.2 Separate spheres

The 'separate sphere' approach of Lundberg and Pollack (1993) considers a model with two public goods and assume that each partner is assigned a public good to which they alone can contribute; this is their 'sphere' of responsibility or expertise. These spheres are determined by social norms. The question Lundberg and Pollak address is how the contributions to the individual spheres are determined. If the partners cooperate, they pool their incomes and set the levels of all goods at the Nash bargaining solution, which is efficient. The Nash solution is enforced by a binding agreement. The resulting allocation then depends on the respective threat points of the husband and wife. They consider the threats of continued marriage in which the partners act non-cooperatively and each chooses independently the level of public good under their domain. In this case, the outcome is inefficient. Specifically, if the partners' individual utilities are additively separable in the two public goods (implying no strategic interactions) each partner will choose the level desired by him\her given their respective incomes. If the wife is poor and the child is under her sphere, the outcome will be under provision of child services. This

[22] The relationship between intrahousehold decision processes in a collective framework and equilibrium on the market for marriage will be the main topic of the second part of the present book.

[23] That the household should spend its entire budget may seem an obvious implication of efficiency, et least in the static context under consideration here. However, the authors show that the property may be violated in the presence of negative externalities.

solution can be modified, however, by transfers that the husband voluntarily commits to pay his wife (before incomes are known) or by a direct purchase of child services in the market.

5.12.3 Inefficient bargaining

Basu (2006) considers a model in which agents bargain in a cooperative way, but the respective threat points depend in part on *endogenous* decisions. For instance, when deciding on labour supply and consumption, a spouse's bargaining position may depend not only on her wage and non labor income, but also on the labor income she generates. Basu analyzes the corresponding model, and shows in particular that multiple equilibria may coexist; moreover, decisions may not be monotonic in a member's power (for instance, child labour can decline and rise as the wife's power increases). It is important to note, here, that although it uses a bargaining framework, Basu's model leads to Pareto inefficient decisions, because of the noncooperative ingredient implicit in the framework. Typically, linking a person's weight to that person's labor income leads to oversupply of labor: once an efficient allocation has been reached, it is individually rational for each spouse to marginally boost her Pareto weight through additional labor supply. The outcome has therefore a prisoner's dilemma flavor. Both members could benefit from a *simultaneous* reduction of their labor supply that would leave Pareto weights unchanged, but strategic incentives prevent this Pareto improvement from taking place.

A similar intuition had actually been proposed earlier by Brossolet (1993) and Konrad and Lommerud (1995). In the two-period model of Konrad and Lommerud, individuals first invest in education, then marry; when married, their decisions are derived from a Nash bargaining framework. Since investments in human capital are made noncooperatively and current investments will serve to improve future bargaining power, there is again inefficient over investment in human capital. Unlike Basu, the second period outcome is efficient in the static sense (i.e., labor supply choices, conditional on education, are ex post Pareto efficient); the inefficiency, here, is dynamic, and can be seen in the initial overinvestment.

In both cases, efficiency could be restored through adequate commitment devices. In practice, such devices are likely to exist in Basu's setting (since the Pareto improvement could be reached during marriage) but not in Konrad and Lommerud's framework (because investments are made before the spouses meet). All in all, these models emphasize the key role of commitment, a point that has been evoked earlier and that will be extensively discussed in Chapter 6.

References

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

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

[3] Becker, Gary S. and Kevin M. Murphy, **Accounting for Tastes,** (Cambridge Mass: Harvard University Press, 1996).

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

[5] Becker, Gary S., 'A Theory of Social Interactions', **Journal of Political Economy**, 82 (1974), 1063-93.

[6] Becker, Gary S., **Treatise on the family,** (Cambridge Mass: Harvard University Press, 1991).

[7] Becker, Gary S. and Robert J. Barro, 'A Reformulation of the Economic Theory of Fertility', **The Quarterly Journal of Economics,** 103 (1988), 1-25.

[8] Bergstrom, Theodore C., 'Puzzles: Love and Spaghetti, The Opportunity Cost of Virtue', **The Journal of Economic Perspectives,** 3 (1989), 165-173.

[9] Bergstrom, Theodore C., Blume, Larry and Hal R. Varian, 'On the Private Provision of Public Goods', **Journal of Public Economics,** 29 (1986), 25-49.

[10] Bergstrom, Theodore C., 'A Fresh Look at the Rotten Kid Theorem–and Other Household Mysteries', **Journal of Political Economy**, 97 (1989), 1138-1159.

[11] Bernheim, B. Douglas and Kyle Bagwell, 'Is Everything Neutral?', **Journal of Political Economy,** 96 (1988), 308-338.

[12] Bernheim, B. Douglas and Oded Stark, 'Altruism within the Family Reconsidered: Do Nice Guys Finish Last?', **The American Economic Review,** 78 (1988), 1034-1045.

[13] Brossolet Cecile, **Femmes au travail : l'introuvable égalité**, (Paris : L'Harmattan, 1993).

[14] Browing, Martin, Chiappori, Pierre-André and Valérie Lechene, 'Distributional Effects in Household Models: Separate Spheres and Income Pooling', Economics Series Working Papers No. 293, University of Oxford, Department of Economics, (2006).

[15] Browning, Martin, Chiappori, Pierre-André and Arthur Lewbel, 'Estimating Consumption Economies of Scale, Adult Equivalence Scales, and Household Bargaining Power', CAM Working Paper No. 2003-12, (2006).

[16] Bruce, Neil and Michael Waldman, 'The Rotten-Kid Theorem Meets the Samaritan's Dilemma', **The Quarterly Journal of Economics,** 105 (1990), 155-165.

[17] Chen, Zhiqi and Frances Woolley, 'A Cournot-Nash Model of Family

Decision Making', **The Economic Journal,** 111 (2001), 722-748.

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

[19] Chiappori, Pierre-André, 'Collective Labor Supply and Welfare', **Journal of Political Economy,** 100 (1992), 437-67.

[20] Chiappori, Pierre-André, 'Nash-Bargained Household Decisions: A Rejoinder', **International Economic Review,** 32 (1991), 761-62.

[21] Chiappori, Pierre André, 'Random Labor Supply', Columbia University, (2009).

[22] Chiappori, Pierre-André and Ivan Werning, 'Comment on "Rotten Kids, Purity, and Perfection"', **Journal of Political Economy,** 110 (2002), 475-504.

[23] Cornes, Richard C. and Emilson C.D. Silva, 'Rotten Kids, Purity and Perfection', **Journal of Political Economy,** 107 (1999), 1034-1040.

[24] Gersbach, Hans and Hans Haller, 'Allocation among Multi-Member Households: Issues, Cores and Equilibria', in Alkan, Ahmet (eds.), Aliprantis, Charalambos D. (eds.) and Yannelis, Nicholas C. (eds.), **Current Trends in Economics: Theory and Applications (Studies in Economic Theory, 8)**, (New York: Springer-verlag Berlin Heidelberg, 1999).

[25] Gronau, Reuben, 'Home Production and the Macro Economy-Some Lessons from Pollak and Wachter and from Transition Russia', NBER Working Paper No. 12287, (2006).

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

[27] Haddad, Lawrence and Ravi Kanbur, 'Intrahousehold Inequality and the Theory of Targeting', **European Economic Review,** 36 (1992), 372-378.

[28] Haller, Hans, 'Household Decisions and Equilibrium Efficiency', **International Economic Review,** 41 (2000), 835-847.

[29] Konrad, Kai A. and Kjell Erik Lommerud, 'Family Policy with Non-Cooperative Families', **The Scandinavian Journal of Economics,** 97 (1995), 581-601.

[30] Lundberg, Shelly and Robert A. Pollak, 'Separate Spheres Bargaining and the Marriage Market', **Journal of Political Economy,** 101 (6), 988-1010.

[31] Lundberg, Shelly and Robert A. Pollak, 'Efficiency in Marriage', **Review of Economics of the Household,** 1 (2003), 153-167.

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

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

[34] Matthews, Steven A., 'Smooth Monotone Contribution Games', PIER

Working Paper Archive 06-018, Penn Institute for Economic Research, Department of Economics, University of Pennsylvania, (2006).

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

[36] Myerson, Roger B., **Game Theory: Analysis of Conflict** , (Cambridge Mass: Harvard University Press, 1991).

[37] Pollak, Robert A. and Michael L. Wachter, 'The Relevance of the Household Production Function and Its Implications for the Allocation of Time', **The Journal of Political Economy,** 83 (1975), 255-278.

[38] Samuelson, Paul A., 'Social Indifference Curves', **The Quaterly Journal of Economics,** 70 (1956), 1-22.

[39] Samuelson, Paul A., 'The Pure Theory of Public Expenditure', **The Review of Economics and Statistics,** 36 (1954), 387-389.

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

[41] Warr, Peter G., 'The Private Provision of a Public Good is Independent of the Distribution of Income', **Economics Letters,** 13 (1983), 207-211.

Figure 1: Demand for public good

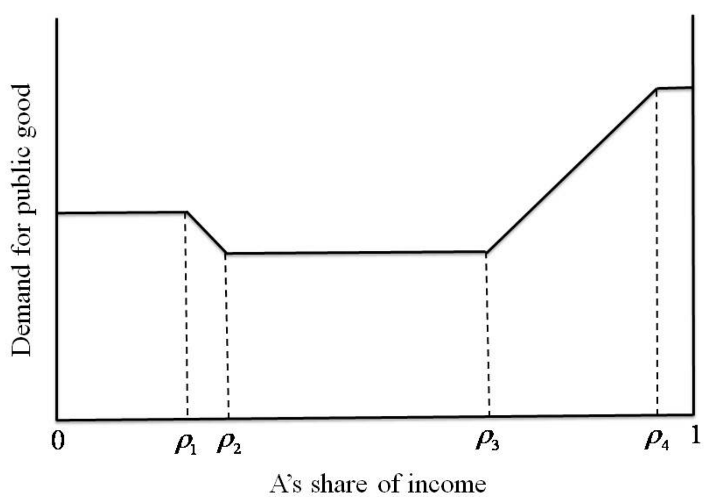

Figure 2: The demand for public goods with altruism.

Figure 3: Household demands for public goods.

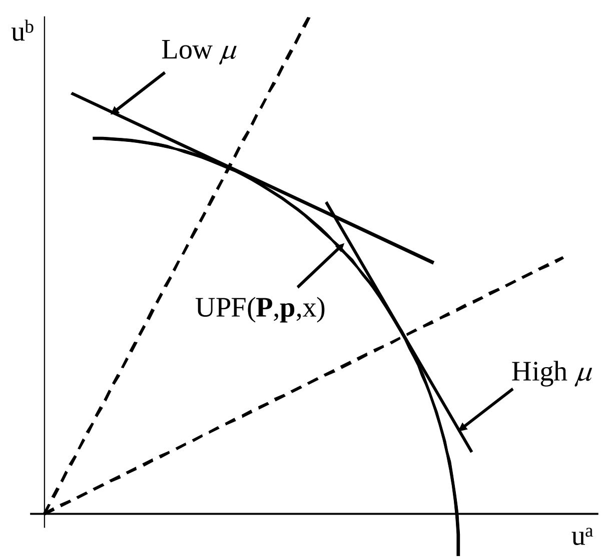

Figure 4: The utility possibility frontier.

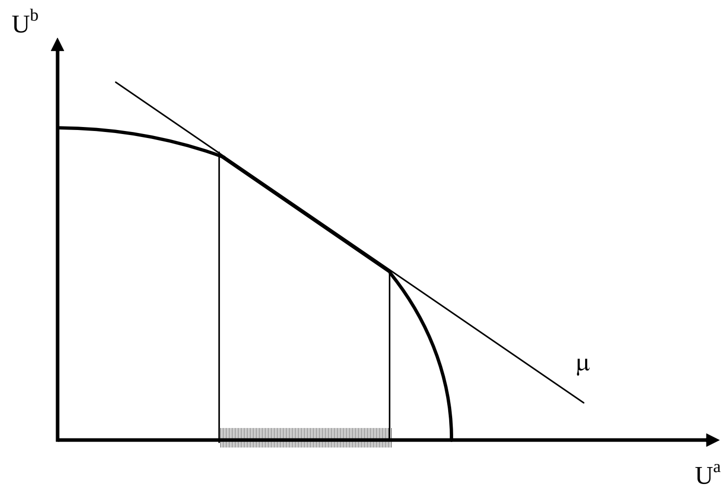

Figure 5: All utilities in the shaded area correspong to the same μ

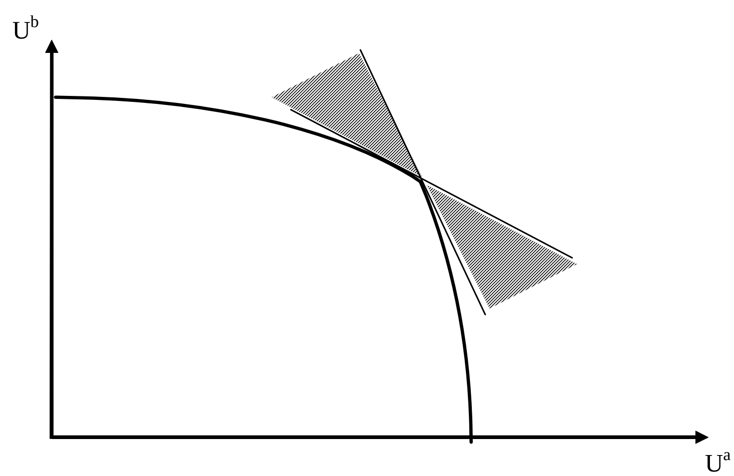

Figure 6: All μ in the shaded cone correspond to the same u^a

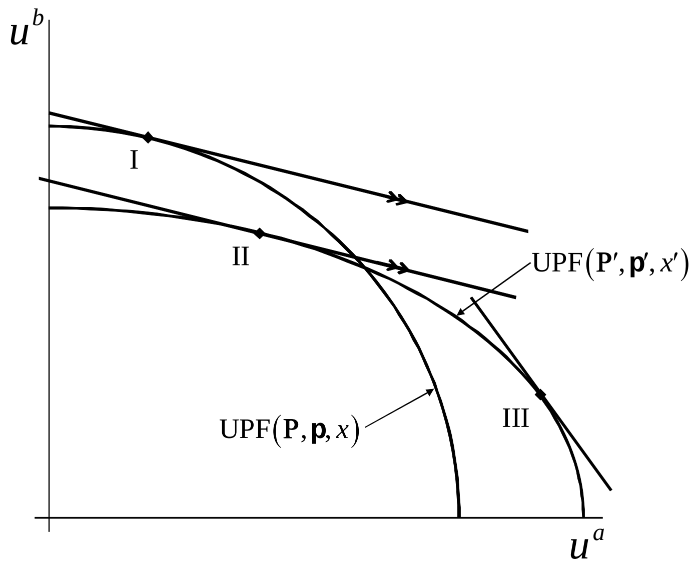

Figure 7: The effects of changes in prices

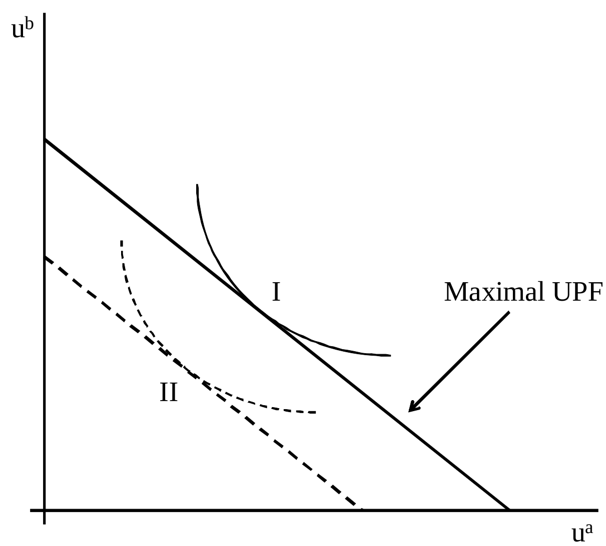

Figure 8: The RKT utility possibility frontier.