

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

August 2007

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 household¹ 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 the chapters below.

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. To keep things simple, we assume that any particular commodity is either purely public or purely private. 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 $(\mathbf{Q}, \mathbf{q}^a, \mathbf{q}^b)$. The associated market prices are given by the N -vector \mathbf{P} and the n -vector \mathbf{p} for public and private goods respectively.

We assume that each married person has her or his own preferences over the allocation of family resources. Denote a 's utility function by $U^a(\mathbf{Q}, \mathbf{q}^a, \mathbf{q}^b)$ and b 's by $U^b(\mathbf{Q}, \mathbf{q}^a, \mathbf{q}^b)$. 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. This is a very mild restriction.

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

¹Children will however be introduced at some point.

²Sometimes the term *altruistic* is used for preferences taking this form. Pollak has recently suggested the term *deferential* since each person defers to the judgement of the

functions $u^a(\mathbf{Q}, \mathbf{q}^a)$ and $u^b(\mathbf{Q}, \mathbf{q}^b)$ respectively. The most general form has

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

Thus a cares about b 's allocation only through the private pleasure that b actually derives from his consumption. Caring rules out externalities and entails weak separability between \mathbf{q}^a and \mathbf{q}^b conditional on \mathbf{Q} . That is, a 's evaluation of her private consumption \mathbf{q}^a does not depend on the private goods b consumes (and *vice versa*). A more commonly used form is the restricted version:

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

Generally we take the weights δ^a and δ^b to be non-negative and such that each person cares for the other but not as much as they care for themselves. Some authors use a slightly different representation of altruism, namely

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

The intuition here is that a 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:

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

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

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

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.

A special case of these preferences is that in which there is no caring, $\delta^a = \delta^b = 0$; this is known as the *egotistic* case. We also note that it is possible to have one person caring and not the other; for example, $\delta^a > 0$ and $\delta^b = 0$.

other regarding their consumption.

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 good theorem. In that case we can consider the unique private good to be ‘money’.

In many contexts below we shall use 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 remuneration of the private goods) put into a form that is similar to the Gorman polar form:

$$\begin{aligned} u^a(\mathbf{Q}, \mathbf{q}^a) &= f^a(q_2^a, \dots, q_n^a) + F^a(\mathbf{Q}) + G(\mathbf{Q})q_1^a \\ u^b(\mathbf{Q}, \mathbf{q}^b) &= f^b(q_2^b, \dots, q_n^b) + F^b(\mathbf{Q}) + G(\mathbf{Q})q_1^b \end{aligned} \quad (6)$$

where $G(\mathbf{Q}) > 0$ for all \mathbf{Q} . Note that the $G(\cdot)$ function is identical for both members, whereas the $f(\cdot)$ and $F(\cdot)$ functions can be individual-specific. In words, the transferable utility assumption implies that, for some well chosen cardinalizations 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:

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

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

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

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 $\Omega(\mathbf{q}, t^a, t^b)$ 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.

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:

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

The associated constraints are:

$$\begin{aligned} \sum \mathbf{q}^j &= \mathbf{q} \\ \sum t_j^s &= t^s, \quad s = a, b \end{aligned} \quad (10)$$

Each person has preferences defined over household produced goods and the vectors of time use, $U^s(\mathbf{C}, \mathbf{c}^a, \mathbf{c}^b, \mathbf{t}^a, \mathbf{t}^b)$ 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 $f^j(\cdot)$) and may also enjoy spending time with the child (captured by the presence of t_j^b in $U^b(\cdot)$). 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:

$$C = f(t^a, t^b) \quad (11)$$

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

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

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

$$C = f(t^a, t^b) = t^b \phi(r) \quad (13)$$

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

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

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,

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

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 given identification.

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 ???. 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:

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

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

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

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

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

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 descendants, 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

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

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.

Finally, 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 Distribution factors

Among the various factors that can influence household behavior, many have a ‘direct’ impact on either preferences or the budget constraint. A change in price, an increase in wages or 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.

Divorce

Although we cannot give a more formal description of this idea for the moment, some examples can be given. Haddad and Kanbur (1992) analyze the welfare impact of the introduction of a guaranteed employment program in some Indian states. According to these programs, any applicant is offered a temporary job in some public work. Haddad and Kanbur point out that the impact of such a measure can be much stronger than suggested by the (relatively small) number of actual applicants to such jobs. They argue that the availability of an external source of income can dramatically alter the bargaining position of women within the household, say because the threat of leaving the household becomes more credible (it no longer means starvation). If we believe that the ‘bargaining powers’ of the members may affect the outcomes of the decision process, then the measure may have a significant impact even on individuals (or couples) that will ultimately never apply to the program. After all, a standard conclusion of bargaining theory is that threats matter even though they are rarely carried out.

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.

Individual incomes or wealths 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. This idea has been used in a number of empirical works. For instance, Duflo (1999), studying a reform of the South African social pension program that extended the benefits to a large population that was not previously covered, finds that the gender of the recipient of this windfall gain has a significant impact on children's health (with a considerable improvement when the recipient is female, but no visible effect for males). The recipient's gender is a typical distribution factor; since the household's total budget constraint is the same whether the amount is paid to one member or the other, Duflo's finding is strongly suggestive of an impact of the distribution factor on the decision process.

In general, much can be learned about intrahousehold decision processes from a careful study of how distribution factors influence behavior. In practice, however, exogeneity is a difficult requirement, especially when the population studied is heterogenous. Current non labor income is likely to reflect past behavior - people who worked more and saved more will receive more. On cross sectional data, this creates a spurious correlation between non-labor income and preferences, and non-labor income cannot be considered as exogenous. Even legislation changes can be partly endogenous, either because laws tend to reflect the preferences of the individuals living in the relevant community, or because people may tend to take legal factors into account when deciding whether to migrate to another location. This suggests two caveats. First, the definition of a distribution factor may be quite delicate, and may depend on the type of data available (for example, cross sectional versus panel data). Secondly, an ideal theory should not only describe and formalize how distribution factors may influence the decision process, but also generate testable predictions upon their influence. In other words, whether a given variable is an acceptable distribution factor should ultimately be an empirical issue; it is incumbent *on the theory* to suggest the tests that will allow to answer that question.

4 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 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:

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

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 $\tilde{U}(\mathbf{Q}, \mathbf{q}^a, \mathbf{q}^b)$. In the next section we shall investigate when such an assumption is justified, but for now we simply consider the consequences. One natural assumption to impose on the household utility function is that it respects the individual preferences in the sense that:

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

where W is a *utility weighting function* which is strictly 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:

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

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

$$\begin{aligned} & \max_{(\mathbf{Q}, \mathbf{q})} U(\mathbf{Q}, \mathbf{q}) \\ \text{subject to } & \mathbf{P}'\mathbf{Q} + \mathbf{p}'\mathbf{q} \leq x \end{aligned}$$

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

$$\mathbf{Q} = \Xi(\mathbf{P}, \mathbf{p}, x) \quad (22)$$

$$\mathbf{q} = \xi(\mathbf{P}, \mathbf{p}, x) \quad (23)$$

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.

5 Noncooperative models

5.1 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 modelled noncooperatively, 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 non cooperative solution boils down to the following two programs:

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

In words, the non cooperative 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 $\xi^s(\mathbf{p}, Y^s)$. Note that the allocation $(\xi^a(\mathbf{p}, Y^a), \xi^b(\mathbf{p}, Y^b))$ 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

$$\xi(\mathbf{p}, Y^a, Y^b) = \xi^a(\mathbf{p}, Y^a) + \xi^b(\mathbf{p}, Y^b) \quad (25)$$

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 condition s hold. That is, if the two agents have linear Engel curves with each partner having the same slope for any good:

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

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

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

⁴Several 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.

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

5.2 One public good.

Whenever direct interaction between members are 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, non cooperative 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:

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

$$\max_{Q^b, q^b} \left\{ u^b \left(Q^a + Q^b, q^b \right) \text{ subject to } PQ^b + pq^b = Y^b \right\} \quad (29)$$

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:

$$\frac{u_Q^a}{u_q^a} \left(\hat{Q}, \hat{q}^a \right) = \frac{P}{p} \quad (30)$$

$$\frac{u_Q^b}{u_q^b} \left(\hat{Q}, \hat{q}^b \right) = \frac{P}{p} \quad (31)$$

If we sum the budget constraints we have:

$$P\hat{Q} + p \left(\hat{q}^a + \hat{q}^b \right) = Y^a + Y^b \quad (32)$$

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

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

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 (31) holds for her. Person b spends all of his income on the private good, so that:

$$\frac{u_Q^b}{u_q^b} \left(\hat{Q}, \frac{Y^b}{p} \right) \leq \frac{P}{p} \quad (34)$$

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:

$$\begin{aligned} \hat{q} &= \hat{q}^a + \frac{Y^b}{p} = \xi^a(P, p, Y^a) + \frac{Y^b}{p} \\ \hat{Q} &= \hat{Q}^a = \Xi(P, p, Y^a, Y^b) \end{aligned} \quad (35)$$

Finally we note that the allocations this non-cooperative procedure leads to are always inefficient (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.⁶

We now present an example to illustrate some of the points made here.

Example 1 *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 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:*

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

(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

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

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 ($\alpha > 1$ or $\alpha < 1$). 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 from 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 good 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.




Figure 1: The demand for public goods.

5.3 Many public goods.

When we turn to the more realistic case with more than one public good, the important features we saw in the last subsection persist but some new ones emerge. The main points can be seen in a model with N public goods, a single private good and prices normalised to unity. The voluntary contributions model has:

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

$$\max_{\mathbf{Q}^b, q^b} \left\{ u^b \left(\mathbf{Q}^a + \mathbf{Q}^b, q^b \right) \text{ subject to } \mathbf{e}' \mathbf{Q}^b + q^b = Y^b \right\} \quad (37)$$

where \mathbf{e} is an N -vector of ones. Let $(\hat{Q}_1^s, \dots, \hat{Q}_N^s)$ 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) 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 equalise the mrs's:

$$\frac{u_i^a}{u_j^a} = \frac{u_i^b}{u_j^b} \quad (38)$$

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 (38) 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.

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

⁸This 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.

We illustrate with an example 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 2 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 her 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 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). 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).




Figure 2: Household demands for public goods.

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.

6 Cooperative models

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 behaviour (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 whatever resources are allocated, none are left on the table.

In general there are very many ways to model cooperative behaviour. 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. Following Chiappori we refer to such models as *collective* and refer to households that always have Pareto efficient outcomes as *collective households*. As we shall show shortly, the collective model assumes, in addition to efficiency, a stable rule that determines the selection of one point on the efficiency frontier. More specific representations, based on bargaining theory, are briefly discussed at the end of this section.

In the remainder of this chapter we briefly introduce the collective model. Chapters 4 and 5 expand on this discussion.

6.1 Assuming efficiency

Natural as it seems, the efficiency assumption should however be carefully discussed. The main motivation supporting it is that households represent perhaps the most typical example of repeated interactions. A large literature in game theory, based on several ‘folk theorems’, suggests that in such situations, efficiency should prevail.⁹ There are however two situations (at least) in which the efficiency assumption may fail to apply. One is when existing

⁹Note, however, that technically most folk theorems only apply to infinitely repeated interactions.

social norms impose patterns of behavior that may conflict with efficiency. For instance, Udry (1996) finds 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.¹⁰ 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 into 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 stick to the most common approach, which maintains the efficiency assumption throughout static models in which commitment is less of an issue, but relaxes it in the more dynamic approaches.

6.2 Modeling efficiency

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 $(\mathbf{Q}, \mathbf{q}^a, \mathbf{q}^b)$ is Pareto efficient if any other allocation $(\bar{\mathbf{q}}^a, \bar{\mathbf{q}}^b, \bar{\mathbf{Q}})$ that is feasible:*

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

gives that if $u^a(\bar{\mathbf{Q}}, \bar{\mathbf{q}}^a, \bar{\mathbf{q}}^b) > u^a(\mathbf{Q}, \mathbf{q}^a, \mathbf{q}^b)$ then $u^b(\bar{\mathbf{Q}}, \bar{\mathbf{q}}^a, \bar{\mathbf{q}}^b) < u^b(\mathbf{Q}, \mathbf{q}^a, \mathbf{q}^b)$ (and conversely).

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

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

¹⁰ 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. From such a perspective, norms resulting in systematic deviations from efficiency should be scarce, at least in ‘open’ enough societies.

the problem:

$$\max_{\mathbf{Q}, \mathbf{q}^a, \mathbf{q}^b} u^b(\mathbf{Q}, \mathbf{q}^a, \mathbf{q}^b) \quad (39)$$

$$\text{subject to } \mathbf{P}'\mathbf{Q} + \mathbf{p}'(\mathbf{q}^a + \mathbf{q}^b) \leq x \quad (40)$$

$$\text{and } u^a(\mathbf{Q}, \mathbf{q}^a, \mathbf{q}^b) \geq \bar{u}^a \quad (41)$$

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 (39). 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 (39).

The solution to (39), 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 $u^b = \Upsilon(\mathbf{P}, \mathbf{p}, x, \bar{u}^a)$. The set of all pairs $(\bar{u}^a, \Upsilon(\mathbf{P}, \mathbf{p}, x, \bar{u}^a))$ 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 concave it is straight forward to show that the function $\Upsilon(\cdot)$ is strictly concave in \bar{u}^a . This allows us to write Program (39) in a different but equivalent way. Let μ denote the Lagrange multiplier for constraint (41); note that μ is always nonnegative. Then the program is equivalent to:

$$\max \mu u^a(\mathbf{Q}, \mathbf{q}^a, \mathbf{q}^b) + u^b(\mathbf{Q}, \mathbf{q}^a, \mathbf{q}^b) \quad (42)$$

under the constraint ((40)). 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:

$$\mu u^a(\mathbf{Q}, \mathbf{q}^a, \mathbf{q}^b) + (1 - \mu) u^b(\mathbf{Q}, \mathbf{q}^a, \mathbf{q}^b), \quad (43)$$

The Pareto weight plays a critical role in all that follows.

Graphically, on Figure ??, 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 (39), or by the slope of the Pareto frontier at that point, here μ as in program (42). Given that the UPF is strictly concave there is a one-to-one, 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 $-\mu$.

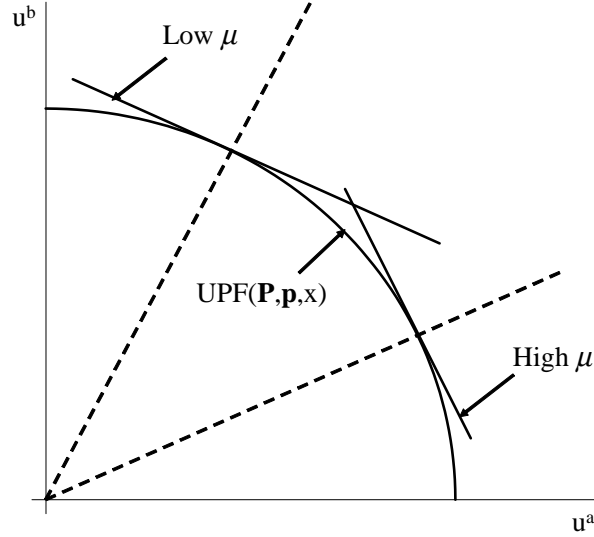


Figure 3: The utility possibility frontier.

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 $F(\cdot)$. Under smoothness conditions, the set is one-dimensional, in the sense that it 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.

6.3 The collective axiom

Although the assumption of efficiency is useful, to go further we need to impose more structure. Specifically, the assumption of efficiency specifies that the household chooses some point on the UPF, but it does not specify which one. Indeed, there is no necessity that faced with the same prices and income in two different time periods, the household chooses the same point. If we are to make any use of the efficiency assumption we also need to assume that the Pareto weight is a well-defined *function* of prices, income and possibly distribution factors. That is, for any vector $(\mathbf{P}, \mathbf{p}, x, \mathbf{z})$ there exists exactly *one* value for μ :

Axiom 4 (Collective axiom) *The Pareto weight μ is a function of prices,*

income and possibly distribution factors : $\mu(\mathbf{P}, \mathbf{p}, x, \mathbf{z})$

A fundamental assumption in unitary demand theory is that individual preferences are stable, in the sense of not change capriciously from moment to moment. This is not a logical requirement but a necessary one if we wish to make predictions based on our models. By the same token, the collective model requires stability for the Pareto weight function.¹¹ One way to justify the collective axiom is to assume that the household uses a well specified bargaining model but this is not strictly necessary. The great advantage of the collective assumption is that we do not have to specify the mechanism households use, but only assume that there is one. In addition, we often assume that the function μ has convenient properties such as continuously differentiability.

6.4 Pareto weights and ‘power’

A major advantage of the formulation in (42) 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 (42) is zero then it is as though b is a dictator, while if μ is large then a effectively gets her way. In general, a key property of (43) 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. 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, *co-operation 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

¹¹In empirical work we would not necessarily assume that the econometrician observes all the distribution factors in which case the Pareto weight would be a stochastic function.

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 and legally binding contracts, tradition and 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.

6.5 Household utility

If the Pareto weights are, in fact, fixed (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:

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

With this definition of the household utility u^f , program (42) is equivalent to the maximization of u^f 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 $\mu(\mathbf{P}, \mathbf{p}, x, \mathbf{z})$. 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^f 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 of how we do it. Following standard demand theory we do not allow prices to enter the individual utility functions. However, we allow all prices and income to enter through *one* index, the Pareto weight. 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 4, where we ignore distribution

factors. Here we consider a change in prices and incomes that moves the UPF from $UPF(\mathbf{P}, \mathbf{p}, x)$ to $UPF(\mathbf{P}', \mathbf{p}', x')$. Initially the point I is chosen on $UPF(\mathbf{P}, \mathbf{p}, x)$. 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 $UPF(\mathbf{P}', \mathbf{p}', x')$ is parallel to the tangent at point I on $UPF(\mathbf{P}, \mathbf{p}, x)$. 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 $UPF(\mathbf{P}', \mathbf{p}', x')$ from II to III .




Figure 4: The effects of changes in prices

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. 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 its 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 4 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 3 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 $\frac{Y^a}{Y^b}$, 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 $W(u^a, u^b, \mathbf{P}, \mathbf{p}, x, \mathbf{z})$ 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 $W(\cdot)$ 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).

6.6 Caring

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

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

The maximand is now

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

which can also be written as:

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

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

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

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 unity, 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:

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

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 $u^a = u^b$.

6.7 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 (42) becomes

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

which is equivalent to

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

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

Modeling the child as a decision maker gives

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

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. 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!

6.8 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 $\mu(\cdot)$) 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 agents 'weights' in the decision process.

Nonetheless there are circumstances under which whatever the particular decision process, the household will act as though it has a single utility function. Specifically, this is the case when the preferences display *transferable utility* (TU); see subsection ???. Indeed, under TU, program (42) becomes

$$\begin{aligned} \max \mu & (f^a(q_2^a, \dots, q_n^a) + F^a(\mathbf{Q}) + G(\mathbf{Q})q_1^a) \\ & + (f^b(q_2^b, \dots, q_n^b) + F^b(\mathbf{Q}) + G(\mathbf{Q})q_1^b) \end{aligned} \quad (51)$$

under the budget constraint 40. 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

$$\begin{aligned} \max \mu & (f^a(q_2^a, \dots, q_n^a) + F^a(\mathbf{Q}) + G(\mathbf{Q})q_1^a) \\ & + (f^b(q_2^b, \dots, q_n^b) + F^b(\mathbf{Q}) + G(\mathbf{Q}) (x - \mathbf{P}'\mathbf{Q} + (\mathbf{q}_{-1}^a + \mathbf{q}_{-1}^b) - q_1^a)) \end{aligned} \quad (52)$$

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

$$\mu G(\mathbf{Q}) - G(\mathbf{Q}) = 0 \quad (53)$$

which implies $\mu = 1$ implying 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 $(\mathbf{Q}, \mathbf{q}_{-1}^a, \mathbf{q}_{-1}^b)$. In fact they will agree to maximise the sum of their individual utilities given by:

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

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.

6.9 The Rotten Kid Theorem

To follow.

6.10 Bargaining models

While the efficiency assumption restricts the set of possible allocations, we are still left with ‘many’ of them - as many, in fact, as possible values for the Pareto weight μ . If we want to make more specific predictions on household behavior, we need to specify more precisely the decision process itself. A standard tool, at this stage, is cooperative game theory, and in particular bargaining models. Here we present a brief summary of bargaining models, included only for the sake of completeness. For a detailed presentation, including a survey of noncooperative foundations of cooperative equilibria, the reader is referred to one of the numerous textbooks on game theory.

Any bargaining model require 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 will be used to determine how the threat points influence the location of the chosen point on the Pareto frontier. Note that if the point $T = (T^a, T^b)$ is outside of the Pareto set, then no agreement can be reached, since at least one member would lose by agreeing. If T sits on the Pareto frontier, then only one agreement is possible, and both agents are indifferent between this agreement and no agreement at all. In the following, we consider the more interesting case in which T belongs to the interior of the Pareto set so that both agents can gain from the relationship. The previous arguments can explain household dissolution. If, because of some external shock, the Pareto set shrinks (or external opportunities expand) in such a

way that T is outside the new feasible set, then no agreement is possible, which optimally results in separation.

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 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.

Nash bargaining The most commonly used bargaining concept was proposed by John Nash in the early 1950's. Some of the axioms on which the notion relies 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 transformations¹²: 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 the independence axiom. 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

¹² An affine mapping is of the form $f(x) = ax + b$.

solution maximises 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 (\bar{u}^a, \bar{u}^b) on the Pareto frontier that maximizes the product $(u^a - T^a)(u^b - T^b)$. That is, the Nash bargaining allocation $(\mathbf{Q}, \mathbf{q}^a, \mathbf{q}^b)$ solves

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

under the budget constraint 41. Thus the product $(u^a - T^a)(u^b - T^b)$ can be considered as a household utility function, that is maximized on the Pareto set. Note that $(u^s - T^s)$ 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.¹³

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 (\bar{u}^a, \bar{u}^b) , \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 $\gamma_a \log(u^a - T^a) + \gamma_b \log(u^b - T^b)$. In this form, the weights γ_s introduce an asymmetry between the members' situations.

Kalai-Smorodinsky To follow.

6.11 Empirical content of bargaining models

Which additional insights can we obtain from the use of bargaining concepts? This depends on what is known on the threat points. Indeed, a first result 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

¹³Note that simply maximizing the sum of surpluses, $(u^A - T^A) + (u^B - T^B)$, would violate the invariance axiom.

maximum utility a person could reach after divorce. Such an idea seems well adapted to situations such as the guaranteed employment program discussed in subsection 3, in which 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.¹⁴

A second idea relies on the presence of public goods, and the fact that non-cooperative behaviour typically leads to inefficient outcomes. The idea, then, is to take the non-cooperative outcome as the threat point: in the absence of an agreement, both members provide the public good(s) egotistically, not taking into account the impact of their decision on the other member's welfare. This version captures the idea that the person who would suffer more from this lack of cooperation (the person who has the higher valuation for the public good) is likely to be more willing to compromise in order to reach an agreement. 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 5.3, applies here as well: total income being kept constant, a change in respective incomes that does not affect the non cooperative 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.

A variant of this idea is developed in the 'separate sphere' model of Lundberg and Pollack (1993). They consider 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 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.

Finally, it must be remarked that assumptions on threat points tend to be strong, somewhat *ad hoc*, and often not independently testable. Given this

¹⁴ An additional difficulty is empirical. The estimation of utility in case of divorce is delicate, since most data set allow us to estimate (at best) an ordinal representation of preferences, whereas Nash bargaining requires a cardinal representation.

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.

References

- [1] Becker, G. (1996). *Accounting for Tastes*. Cambridge, Mass, Harvard University Press.
- [2] Becker, G. S. (1965). "A Theory of the Allocation of Time." *Economic Journal* 75(299): 493-517.
- [3] Becker, G. S. (1991). *A treatise on the family*. Cambridge, Mass ; London, Harvard University Press.
- [4] Becker, G. S. and R. J. Barro (1988). "A Reformulation of the Economic-Theory of Fertility." *Quarterly Journal of Economics* 103(1): 1-25.
- [5] Bergstrom, T. (1989). "Puzzles - Love and Spaghetti, the Opportunity Cost of Virtue." *Journal of Economic Perspectives* 3(2): 165-173.
- [6] Bergstrom, T., L. Blume, et al. (1986). "On the Private Provision of Public-Goods." *Journal of Public Economics* 29(1): 25-49.
- [7] Bergstrom, T. C. (1989). "A Fresh Look at the Rotten Kid Theorem - and Other Household Mysteries." *Journal of Political Economy* 97(5): 1138-1159.
- [8] Bernheim, B. D. and K. Bagwell (1988). "Is Everything Neutral." *Journal of Political Economy* 96(2): 308-338.
- [9] Bernheim, B. D. and O. Stark (1988). "Altruism within the Family Reconsidered - Do Nice Guys Finish Last." *American Economic Review* 78(5): 1034-1045.
- [10] Browning, M., Pierre-André Chiappori and Valerie Lechene (2006). "Distributional effects in household models: separate spheres and income pooling." Oxford, University of Oxford.
- [11] Chen, Z. Q. and F. Wooley (2001). "A Cournot-Nash model of family decision making." *Economic Journal* 111(474): 722-748.
- [12] Gronau, R. (2006). *Home Production and the Macro Economy - Some Lessons from Pollak and Wachter and from Transition Russia* NBER.
- [13] Haddad, L. and R. Kanbur (1992). "Intrahousehold Inequality and the Theory of Targeting." *European Economic Review* 36(2-3): 372-378.
- [14] Lundberg, S. and R. A. Pollak (1993). "Separate Spheres Bargaining and the Marriage Market." *Journal of Political Economy* 101(6): 988-1010.
- [15] Lundberg, S. and R. A. Pollak (2003). "'Efficiency in marriage'." *Review of Economics of the Household* 1(3): 153-167.
- [16] Mas-Colell, A., Michael Whinston and Jerry Green (1995). *Microeco-*

conomic Theory. New York, Oxford University Press.

- [17] Matthews, S. (2006). Smooth Monotone Contributions Games. Penn Institute for Economic Research University of Pennsylvania.
- [18] Pollak, R. A. and M. L. Wachter (1975). "Relevance of Household Production Function and Its Implications for Allocation of Time." *Journal of Political Economy* 83(2): 255-277.
- [19] Udry, C. (1996). "Gender, agricultural production, and the theory of the household." *Journal of Political Economy* 104(5): 1010-1046.
- [20] Warr, P. G. (1983). "The Private Provision of a Public Good Is Independent of the Distribution of Income." *Economics Letters* 13(2-3): 207-211.