June 1, 1995

A Survey of Theories of the Family

Ted Bergstrom, *University of California, Santa Barbara*

Chapter 2

A SURVEY OF THEORIES OF THE FAMILY

THEODORE C. BERGSTROM*

University of Michigan

Contents

*Thanks to Maria Cancian, Matz Dahlberg, John Laitner, David Lam, Robert Pollak, Robert Schoeni, Oded Stark, Yoram Weiss, and Robert Willis for welcome encouragement and useful advice.

Handbook of Population and Family Economics. Edited by M.R. Rosenzweig and O. Stark

1. Introduction

To a labor economist or an industrial organization economist, a family looks like "a little factory". To a bargaining theorist, a husband and wife are "two agents in a relation of bilateral monopoly". To an urban economist or a public choice theorist, a family looks like "a little city", or perhaps "a little club". To a welfare economist, a family is an association of benevolently interrelated individuals. Each of these analogies suggests useful ways in which the standard tools of neoclassical economics can aid in understanding the workings of a family. The second section of this review draws on the analogies to a little factory and to a little city. It explores the theory of household technology and the household utility possibility frontier. The third section concerns decision theory within the household. This discussion applies standard consumer decision theory as well as bargaining theory and the theory of public choice. The fourth section of this paper deals with family formation and the choice of mates. This theory is analogous to "Tiebout theory" in urban economics, where the objects of choice include not only the public goods supplied in each city, but which individuals live together. An aspect of family life that has fewer parallels in the economics of market economies is intrafamilial love and altruism. The final section of this paper reviews a growing theoretical literature on love, altruism and the family.

2. Household technology and utility possibility frontiers

2.1. Household production functions

In his *Treatise on the Family*, Becker (1981) emphasizes the importance of division of labor and gains from specialization. Drawing on his 1965 paper, "A Theory of the Allocation of Time", Becker endows households with *household production functions* which describe the possibilities for producing "household commodities". Becker's household commodities are nonmarket goods that are the outputs of production processes that use market goods and the labor time of household members as inputs. His examples of household commodities include "children, prestige and envy, health, and pleasures of the senses". He suggests that the number of household commodities is typically much smaller than the number of market goods.

The concept of production function, borrowed from the theory of the firm, has been a fruitful source of insight into the workings of families. Becker exploits this analogy as he examines such issues as specialization within the household, comparative advantage, returns to scale, factor substitution, human capital and assortative mating. Each individual in Becker's household can use time either for household labor or market labor. The family can purchase *market goods* and either consume them directly or use them as inputs into household production.

Pollak and Wachter (1975) develop the formal structure of Becker's household production model and show that if household commodities are produced with constant returns to scale and no joint production, then "shadow prices" for these commodities are determined by the prices of market goods and the wage rates for market labor – independently of the quantities demanded. This means that the household's production possibility set, like an ordinary competitive budget set, has a linear boundary with marginal rates of transformation that are independent of the quantities chosen.[1] If production possibility sets have this property, then the household production model allows a neat separation of production and consumption activities. Pollak and Wachter observe that with joint production or with nonconstant returns to scale, this separation of production and consumption is lost since the shadow prices of commodities depends on the quantities produced and the boundary of the production possibility set is "curved". Pollak and Wachter argue that unless the production of "household commodities" permits separation of production and consumption activities, there is little to be gained from adding unobservable household commodities to the model. Instead they recommend studying the demand for market goods and leisure directly as functions of wages and the prices of market goods. Pollak and Wachter also maintain that tastes and technology are likely to be confounded by treating nonmeasurable aggregate variables such as "child quality" as commodities. They recommend more narrowly defined child-related commodities such as "scores on standardized tests" or "number of dental cavities".

As Pollak and Wachter point out, even with very general technologies, there will be well-defined demand functions for market goods and supply functions of labor which could in principle be determined from household utility. Rader (1964) establishes general conditions under which "induced preferences" for trades inherit such properties as convexity, continuity and homogeneity from the production functions and the preferences for produced commodities. Muth (1966) shows that even if the output of a household commodity is not directly observable, the assumption that this commodity is produced with constant returns to scale can have interesting testable implications. If, for example, two or more market goods are used as inputs for this good and no other goods, then it must be that the income elasticities of demand for the two goods will be identical. A detailed discussion of the theory of household production and additional references can be found in Deaton and Muellbauer (1980).

2.2. *Household public goods*

A unit of private goods consumed by one person cannot be consumed by another. But some goods, such as living space, household heating and lighting and shared auto-

[1] This result is essentially the generalized Samuelson non-substitution theorem for the "small-country case" where factors can be purchased at constant prices (see Samuelson, 1961; Varian, 1984).

mobile trips are jointly consumed and are best modeled as local public goods which enter simultaneously into the utility functions of all family members.

While Becker (1981) did not explicitly distinguish *household public goods* from private goods, his household technology model could certainly be used to describe production of household public goods as well as ordinary private goods. Manser and Brown (1980) and McElroy and Horney (1981) were among the first to introduce household public goods as an integral part of their models of family behavior. These authors emphasize the benefits of shared public goods as a reason that marriage yields a utility surplus over living separately. Weiss and Willis (1985), in their economic study of divorce and child-support payments, treat the well-being of children as a household public good that enters the utility of both parents whether these parents are married or divorced. Lam (1988) suggests that the presence of household public goods will favor positive assortative mating by income.

Household public goods are modeled as follows. Consider a household with h members. Each household member i has a utility function $U_i(x_i, y)$ where x_i is the vector of private goods consumed by i and where y is the vector of household public goods. A *household allocation* is a vector $(x_1, \ldots, x_h, y)$ that specifies the consumption of private goods by each household member and the vector of household public goods. The household budget and household technology determine a *household production possibility set S* which specifies all possible aggregate consumptions of public and private goods for the household. A *feasible allocation* for the household is an allocation $(x_1, \ldots, x_h, y)$ such that $(x, y) \in S$, where $x = \sum_{i=1}^{h} x_i$.

2.3. *The household utility possibility frontier*

The *utility possibility frontier* is an analytic tool that illuminates many issues in the theory of the family.[2] Consider a household with h members and a household production possibility set S. To each feasible allocation $(x_1, \ldots, x_h, y)$, corresponds a distribution of utilities among household members in which person i gets utility $U_i(x_i, y)$. The set of all utility distributions that can be constructed in this way is called the utility possibility set. The "upper" boundary of the utility possibility set is known as the utility possibility frontier. By construction, the utility possibility frontier consists of all utility distributions that are Pareto optimal for the household.

For any fixed vector y of public goods, it is possible to construct a conditional utility possibility set $UP(y)$, such that $UP(y)$ corresponds to all of the distributions of utility that can be achieved by some feasible allocation in which the vector of public goods is y. The *conditional utility possibility frontier* corresponding to y is the upper

[2] The utility possibility frontier seems to have been introduced to the economic literature by Samuelson (1950), who presents this idea in a section of his paper titled *The Crucially Important Utility Possibility Function.*

boundary of $UP(y)$. In general, the conditional utility possibility frontiers corresponding to different public goods vectors may cross each other. The utility possibility frontier for the household will be the outer envelope of the conditional utility possibility frontiers corresponding to all feasible choices of y.

2.3.1. Examples of utility possibility frontiers for households with public goods

Example 1. A household has two members. There is one private good and one household public good. Total household income is \$3. The quantity of the household public good must be either zero or one unit. The price of the private good is \$1 per unit and the cost of a unit of the public good is \$2. Person 1 has utility function $U_1(X_1, Y) = X_1(Y + 1)$ and Person 2 has utility function $U_2(X_2, Y) = X_2(Y + 1)^2$.

If $Y = 0$, then $U_1 = X_1$ and $U_2 = X_2$. Since in this case, the household budget constraint is $X_1 + X_2 = 3$, the conditional utility possibility set $UP(0)$ is the set $\{(U_1, U_2) \mid U_1 + U_2 \leq 3\}$. In Fig. 1, $UP(0)$ is bounded by the line CD.

If $Y = 1$, then $U_1 = 2X_1$ and $U_2 = 4X_2$. Since the cost of one unit of the public good is \$2, the household budget constraint is now $X_1 + X_2 = 3 - 2 = 1$, the utility possibility set $UP(1)$ is the set $\{(U_1, U_2) \mid \tfrac{1}{2}U_1 + \tfrac{1}{4}U_2 \leq 1\}$. This is the set bounded by the line AB in Fig. 1.

The utility possibility set for the household is the union of the sets $UP(0)$ and $UP(1)$, and the household utility possibility frontier is the thick broken line running from A to E to D. Notice that in this example, some Pareto optimal allocations for the household are achieved by supplying no household public goods and others are achieved by supplying one unit of household public goods.

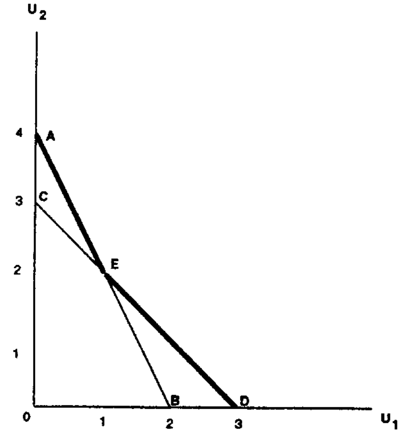

Fig. 1.

Example 2. A household has two members. There is one private good and one household public good. Total income available to the household is \$10. The price of private goods is \$1 per unit and the price of public goods is \$1 per unit. The household can choose any allocation $(X_1, X_2, Y) \geq 0$ such that $X_1 + X_2 + Y \leq 10$. Person 1 has the utility function $U_1(X_1, Y) = X_1 + Y^{1/2}$ and Person 2 has the utility function $U_2(X_2, Y) = X_2 + 3Y^{1/2}$.

An allocation in which both consumers consume positive amounts of the private good will be Pareto optimal if and only if the sum of their marginal rates of substitution between public and private goods equals the price ratio of public to private goods. (This is sometimes known as the Samuelson condition, in honor of Samuelson's (1954) construction of the theory of public goods.) This condition implies that $\frac{1}{2}Y^{-1/2} + \frac{3}{2}Y^{-1/2} = 1$ or equivalently that $Y = 4$. Therefore the set of Pareto optimal allocations in which both household members have positive consumption of private goods consists of all allocations $(X_1, X_2, 4) > 0$ such that $X_1 + X_2 = 10 - 4 = 6$. When $Y = 4$, it must be that $U_1 = X_1 + 2$ and $U_2 = X_2 + 6$. Therefore along the part of the utility possibility frontier corresponding to allocations where both consume positive amounts of private good, it must be that $U_1 + U_2 = X_1 + X_2 + 8$. Since for these allocations, $X_1 + X_2 = 6$, it follows that this part of the utility possibility frontier lies on the line $U_1 + U_2 = 14$. If both consumers are consuming positive amounts of the public good, then it is also true that $U_1 > 2$ and $U_2 > 6$. Therefore the part of the utility possibility frontier corresponding to positive consumption of private goods for both household members is the line segment *BC* in Fig. 2. There are other Pareto allocations where only one of the consumers consumes positive amounts of private goods. These allocations correspond to the curved lines *AB* and *CD* in Fig. 2. At these points, the Samuelson conditions do not apply and the amount of public good supplied is less

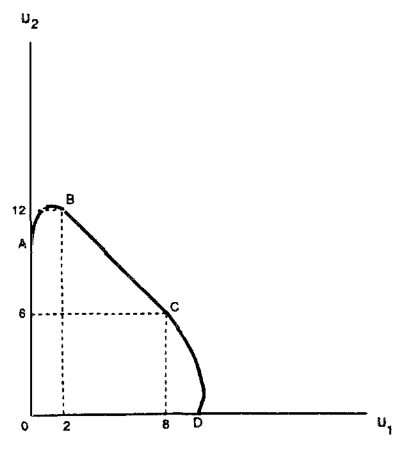

Fig. 2.

than four units. This part of the story is explained in detail by Bergstrom and Cornes (1983) and by Campbell and Truchon (1988).

In Example 2, all Pareto optima such that both consumers have positive consumption of private goods must have the same output of public good. Moreover, for all such allocations, the utility possibility frontier is a straight line. Every Pareto optimal household allocation has the property that (X_1, X_2, Y) could be found by maximizing the sum of utilities, $U_1(X_1, Y) + U_2(X_2, Y)$, subject to the household feasibility constraint $X_1 + X_2 + Y = 10$.

2.4. *Transferable utility in the household*

The term *transferable utility* seems to have originated in game theory, but the idea that it represents is familiar to all economists. Roughly speaking, transferable utility means that utility can be "redistributed", like apples or bananas. If one distribution of utility is possible then so is any other distribution of utilities where individual utilities sum to the same number. The assumption of transferable utility has powerful and interesting implications in the theory of the household. Therefore it is worthwhile investigating this assumption carefully.[3]

This rough definition of transferable utility needs to be extended and clarified. One issue arises in Example 2, where the utility possibility frontier is not a straight line over its entire range, but is a straight line for all utility allocations that are achieved with a positive amount of private good for each consumer. Since this is the part of the utility possibility frontier that is relevant for most analysis, economists usually define transferable utility to include cases like Example 2, where the utility possibility frontier is linear over the "relevant" range of utility distributions.

Another, more subtle aspect of the definition needs clarification. For the utility functions that are specified in Example 1, the utility possibility frontier is not a straight line, but it would be possible to make a monotonic transformation of each person's utility in such a way that with the new utility representation, the utility possibility frontier is a straight line. In fact this is possible whenever the utility possibility frontier is downward-sloping.[4] Transferable utility becomes a nontrivial property

[3] The domain of applicability of transferable utility is not as widely understood as it should be. Some economists and game theorists think there is more transferable utility than there really is and some think there is less. Some believe that transferable utility obtains whenever preferences are continuous and there exists a divisible and fully exchangeable good that is desired by all consumers. Others believe that transferable utility applies only in the very special case of "quasilinear utility", where utility is linear in the quantity of some commodity. As we will show, neither view is correct.

[4] Consider any utility possibility frontier in the positive orthant, defined by an equation $U_2 = F(U_1)$, where F is a strictly decreasing function. Define the functions $g_1(U) = U/(U + F(U))$ and $g_2(U) = U/(U + F^{-1}(U))$. The functions g_1 and g_2 are well defined and strictly monotone increasing. Let $V_1(x) = g_1(U_1(x))$ and $V_2(x) = g_2(U_2(x))$. Then the V_i's are increasing transformations of the U_i's, and for the utility possibility frontier with utilities measured by the functions, V_i is described by $V_1(x) + V_2(x) = 1$.

when we assume that a utility function can be found for each consumer such that the utility possibility frontier is linear for all family budget constraints within some broad class of budgets. In Example 1, it is possible to find utility representations of each person's preferences so that there is a linear utility possibility frontier when household wealth is three. But with this utility function, when household wealth is four, the utility possibility frontier would once again have a kink.

For the utility functions in Example 2, the linear portion of the utility possibility frontier is described by an equation of the form $U_1 + U_2 = C$. This is true for all prices and family incomes. Changes in prices or in family income will change only the constant C, causing a parallel shift of the linear utility possibility frontier. Thus Example 2 has transferable utility and Example 1 does not.

2.4.1. Transferable utility with private goods: Hicksian and Gorman aggregation

The simplest example of a household with transferable utility is a household with two selfish people, one private good, and no public goods. Then preferences can be represented by the utility function $U(x_i) = x_i$ where x_i is i's consumption of the private good. If the household has a fixed wealth W and the price of the good is 1, then the utility possibility frontier is just the set of vectors (U_1, U_2) such that $U_1 + U_2 = W$. Changes in household wealth would shift this utility possibility frontier in a parallel fashion, but it would remain a straight line.

The single-private-good model is more general than may first appear. If the only goods consumed are marketable private goods, then the "Hicks composite commodity theorem" (Hicks, 1956) allows one to model the household as if there were just one good, "wealth". The composite commodity theorem can even accommodate household labor supply if each household member is a price-taker in the labor market. This can be done with the usual trick of treating each person i's leisure as a commodity which is marketable at some price w_i. Total household income includes the value of each person's potential amount of leisure evaluated at that person's wage. Individuals are allocated shares of total household income and are able to "buy" their leisure at the market wage. Therefore if the relative prices of private goods are assumed constant throughout the analysis and if there are no household public goods, there is no loss of generality in assuming transferable utility.

If one wants to analyze the responses of households to changes in relative prices or to study possible changes in the amount of household public goods, then one cannot appeal to Hicks' aggregation as a justification for transferable utility. But there is a nontrivial special class of preferences for which the utility possibility frontier is linear and changes in prices induce parallel shifts of the utility possibility frontier. Bergstrom and Varian (1985) showed that in a pure exchange economy with private goods and convex preferences, there is transferable utility if and only if preferences of all individuals can be represented by indirect utility of the Gorman polar form. Indirect utility is of the Gorman polar form if it can be expressed in the form $V_i(p, m_i) = \alpha(p)m_i + \beta_i(p)$ for each individual i for some function $\alpha(p)$ and functions $\beta_i(p)$.

If indirect utility is of the Gorman polar form, then the income–consumption paths of all consumers are parallel straight lines, the slopes and intercepts of which may depend on prices (see Gorman, 1953). Although different people may allocate their total income differently among goods, in the Gorman case it must be that everyone allocates a marginal dollar of income in the same way. Therefore aggregate demand is unchanged by marginal income redistributions and behaves as if it were the demand of a single consumer who controlled all of the income.

The Gorman class includes preferences represented by a *quasilinear* direct utility function of the form $U(x_1, x_2, \ldots, x_n) = x_1 + f(x_2, \ldots, x_n)$. For many applications, quasilinearity is an unsatisfactory assumption, since if preferences are quasilinear, the demand for all goods except for good 1 is independent of income. But the Gorman class also includes identical homothetic preferences and more generally, preferences which, like the Stone–Geary utility function, are represented by utility functions of the form $U(x - e_i)$, where U is a homogeneous function and e_i is some vector that "displaces the origin" and which may be different for different people. Thus it is possible to have transferable utility, but to have preferences differing among individuals and to have income-responsive demand for all goods.

In the Pollak–Wachter model cited above, there is transferable utility in a household that produces nonmarket goods. The assumptions needed for this result are that household production occurs with single-output production functions and constant returns to scale and that no more than one nonmarketed commodity is used as an input. With these assumptions, there exist "shadow prices" for the produced household commodities. These shadow prices are determined by technology and the price of marketable goods, independently of the quantities demanded. With this setup, transferable utility is guaranteed by the same assumptions on preferences that have been discussed for the case where all consumption goods are marketable. If, on the other hand, there are nonconstant returns to scale in household production or if there is more than one nonmarketable (and nonproduced) input, there will typically not be transferable utility.

2.4.2. *Transferable utility with household public goods*

Example 1 shows that if there are household public goods, there might not be transferable utility, even if there is only one private good. There is, however, a useful special class of preferences over private and public goods for which there is transferable utility. This class bears a neat duality to preferences of the Gorman polar form. Recall that the Gorman polar form requires that *indirect* utility be representable in the form $V_i(p, m_i) = \alpha(p)m + \beta_i(p)$. Bergstrom and Cornes (1981, 1983) found that a necessary and sufficient condition for transferable utility when there is one private good and n public goods is that preferences of each household member are representable by a *direct* utility function of the form $U_i(x_i, y) = f(y)x_i + g_i(y)$, where x_i is the amount of the private good received by person i and y is the n-vector of public goods in the household. One special case is the quasilinear utility representation, $U_i(x_i, y) = x_i + f_i(y)$ (see, for instance, Example 2). In this example, a consumer's willingness to pay for public goods is independent of his income. But there are also examples in this class of utility functions where willingness to pay depends on income. For example, let each consumer i have a utility function of the generalized Cobb–Douglas form $U_i(x_i, y) = x_i(y + b_i)^c$ where $c > 0$ and where the b_i's are arbitrary constants.

The Bergstrom–Cornes result can be combined with the Bergstrom–Varian results on transferable utility with several private goods to characterize transferable utility when there is any number of private goods and any number of public goods. Let us define an *extended indirect utility function* with public goods in the same way that indirect utility is defined when there are no public goods except that utility depends on the amount of public goods as well as on the price vector and on income. Define $V_i(p, m_i, y)$ as the maximum of $U_i(x_i, y)$ subject to the constraint that $px_i = m_i$. It can be shown that there will be transferable utility in a household if and only if this indirect utility function is of the form $V_i(p, m_i, y) = \alpha(p, Y)m_i + \beta_i(p)$.

The utility possibility frontiers in households where utility is of the Bergstrom–Cornes form have the following properties:

- The conditional utility possibility frontiers $UPF(y)$ and $UPF(y')$, corresponding to any two different quantities of public goods y and y' will not cross each other. One of these two frontiers lies strictly to the northeast of the other.
- The utility possibility frontier is a straight line segment with slope −1. All points on the utility possibility frontier that correspond to interior Pareto optima are achieved with the same amount of public goods, $y = y^*$. The end points of the utility possibility frontier correspond to the two utility distributions in which one consumer gets $x^* = W - py^*$ units of private goods, and the other gets no private goods.

2.4.3. *Transferable utility and mate selection*

An important conceptual building block for economic theories of marriage is the utility possibility frontier that any couple would face if they were to marry each other. Shapley and Shubik (1972) and Becker (1974) suggested that this problem could be modeled as a linear programming assignment problem. The linear programming assignment model requires that there be transferable utility within each possible marriage. Here we explore the degree of generality that can be accommodated by transferable utility within this framework.

The notion of household public goods is well-suited for analyzing the issues of compatibility that arise in possible marriages. The public goods model of course applies to such jointly consumed household commodities as heat, light, and a well-tended garden. It is also suitable to the many important joint decisions that a married couple must make. If two people marry each other, they must marry on the same date. They must also agree about where to live, how many children to have, how to educate their children and what size of estate to leave them. Each of these variables can be modeled as a pure public good that enters the utility function of both partners.

Suppose that male i marries female j and they choose private consumption X_i for him, X_j for her, and the vector Y of household public goods. Assume that their utility functions are of the following functional form: $U_i = A(Y)X_i + B_i(Y, j)$ and $U_j = A(Y)X_j + B_j(Y, i)$.

These utility functions indicate three things that a person must consider when he or she contemplates a potential marriage: (i) the vector Y of public choices that would be made in this marriage; (ii) the amount of private goods that he or she would get to consume in that marriage; and (iii) his or her feelings about the intrinsic desirability of the other person as a marriage partner. The third effect is registered by the fact that the functions $B_i(Y, j)$ and $B_j(Y, i)$ depend not only on Y but also on who one has as a partner.

The consumption options available to a potential pair of spouses depend on their joint economic productivity. In particular, let $F_{ij}(Y)$ be the total amount of private good that would be available to male i and female j if they chose to have the vector Y of household public goods. (The function $F_{ij}(Y)$ would incorporate the effect on the household budget of public goods that must be purchased. It could also include the effects of household public goods that influence household income, like location or education of household members.) The set of affordable combinations of private consumptions and public choices is the set of vectors (X_i, X_j, Y) which satisfy the equation $X_i + X_j < F_{ij}(Y)$.

In this case, for any couple, i and j, the part of the utility possibility frontier corresponding to allocations where both persons get positive consumptions of private goods is described by the linear equation $U_i + U_j = A_{ij}$ where A_{ij} is the maximum of $A(Y)X + B_i(Y, j) + B_j(Y, i)$ subject to the constraint $X \leq F_{ij}(Y)$. In this case, the problem of finding Pareto efficient allocations within a potential marriage reduces to a constrained maximization problem in the aggregate quantities X and Y. If they were to marry, there would be transferable utility between any pair i and j, with utility possibility frontier for this pair consisting of utilities that add to A_{ij}.

3. Decision-making in the family

After proving that in general there do not exist "social indifference curves" that rationalize aggregate demand, Samuelson (1956) worries that

> But haven't I in a sense proved too much? Who after all is the consumer in the theory of consumer's (not consumers') behavior? Is he a bachelor? A spinster? Or

is he a "spending unit" as defined by statistical pollsters and recorders of budgetary spending? In most of the cultures actually studied by modern economists, the fundamental unit on the demand side is clearly the "family" and this consists of a single individual in but a fraction of cases.

Economists are not the only social scientists to be concerned about whether the family should be treated as a decision-making agent. The issue is nicely posed by the sociologist, James Coleman (1990: p. 580):

> The family has always been an entity within which multiple activities are carried out: economic production, joint consumption, procreation, socialization of children, and leisure pursuits. Generally it cannot be regarded as a purposive actor ... for it cannot usually be described as having a purpose for which it acts. It is, like society as a whole but on a smaller scale, a system of action composed of purposive actors in relation. Yet in some capacities the family may be usefully regarded as a purposive actor for it is an entity in terms of whose perceived interests natural persons act; for example some persons say they are acting to "uphold the honor of the family." And in some cases a family does act as a unit, to attain ends that can be described as purposes or goals of the family. It may be useful to clarify when and for what purposes a system of actions should be called an actor. For example, in a swarm of insects hovering in the summer air, each insect is darting this way and that, apparently either randomly or in pursuit of its own ends. But the swarm as a whole will move this way or that, hover expand or contract, and then fly off no less coherently than if it were a single organism. ... Thus just as a swarm of insects may be considered an actor, the family may – sometimes – be considered an actor.

3.1. Unitary theories of family decision-making

At least until recently, empirical studies of household demand have routinely assumed that a family acts as if it were maximizing a "family utility function".[5] That they should have done so is understandable since most of the available cross-sectional data on household consumption consists of household aggregates which do not distinguish either the incomes or the consumptions of individual family members. The reservations that Samuelson (1956) raised about this approach were amplified by the work of Manser and Brown (1980) and McElroy and Horney (1981), who showed that if the allocation of resources within the family is determined by Nash bargaining, then

[5] Good surveys of this work can be found in Deaton and Muellbauer (1980) and Browning (1992).

household aggregate demand functions cannot in general be rationalized by maximization of a family utility function.

Theories in which household demand can be rationalized by a family utility function have at various times been called "single-agent", "common preference", "consensus", "altruistic", or "benevolent dictator" theories.[6] But this does not mean that the natural course of neoclassical economics is to assume that families act like single agents. We follow the suggestion of Chiappori et al. (1992), who designate these models as unitary models. The classification "unitary" is sufficiently broad to encompass the several different models of family structure that predict that a family in aggregate behaves "as if" it is maximizing a family utility function.

3.1.1. Unitary models with transferable utility

Consider a household with private goods but no public goods, where indirect utility functions are of the Gorman polar form. If each family member is given an income and all face the same competitive prices, then the total amount of each good consumed by household members is determined by prices and total family income. Changes in the way that income is distributed within the household would have no effect on total household consumption. Gorman (1953) described this situation as the presence of "community preference fields". In this case, household demand can be rationalized as maximizing the utility of a single consumer. Suppose that an econometrician had access to a time series of commodity prices and to this household's time path of total income and total consumption, but could not observe consumption by individual household members. The econometrician would not be able to reject the hypothesis that all household decisions were made by a single rational consumer (who spends the entire household budget on himself).

Conversely, if preferences are not of the Gorman form, total household consumption will, in general, depend on the distribution of income within the household. Therefore if changes in the distribution of household income occur during the course of the time series, an econometrician might detect a violation of the weak axiom of revealed preference in household consumption data. This would enable her to con-

[6] Some economists refer to single-agent theories of the family as the "neoclassical theory of the family". Chiappori (1992) suggests that this is not an appropriate name. Although neoclassical economists sometimes treat the family as a single maximizing agent, more pluralistic theories of the family fall squarely within the neoclassical tradition. A distinctive feature of neoclassical microeconomics is "methodological individualism". When faced with fundamental questions about group behavior, neoclassical economists typically take a reductionist approach that seeks to explain group behavior not as a choice of a single rational agent, but rather as the result of the interplay of actions by group members with distinct objectives. It is true that neoclassical economists such as Gorman (1953), Samuelson (1956), and Becker (1974) have explored special assumptions on preferences under which families act as maximizing agents. It is also true that econometricians have often tried to simplify the task of applying neoclassical theory to household data by assuming that they act like a single agent.

clude that the household does not act like a single rational decision-maker. (Of course the data might not be rich enough so that the econometrician could detect violations of the weak axiom, even though the actual household decisions are not consistent with the unitary model.)

It is a consequence of the second fundamental theorem of welfare economics that if preferences are convex and there are no consumption externalities, then any household that allocates marketable private goods efficiently among its members will act as if each household member is given a personal income and is allowed to spend it as he or she wishes (see Chiappori, 1988, 1992). Combining this fact with our previous discussion, we see that if preferences in a household are convex and if indirect utility is of the Gorman polar form, then an efficiently operating household will act as if all household decisions were made by a single utility-maximizing consumer.

The results of Bergstrom and Cornes (1981, 1983) allow us to extend this result to the case where there are public goods as well as private goods. In particular, if indirect utility is of the form $V(p, m_i, y) = \alpha(p, y)m_i + \beta_i(p, y)$ and if the household chooses a Pareto optimal allocation in which all household members consume some private goods, then the vector of public goods selected and the vector of total household consumption of private goods is independent of the distribution of income within the household. As in the case of private goods with Gorman polar form utility, an observer of the response of household aggregates respond to prices and household income would not be able to reject the hypothesis of decision-making by a single rational consumer.

3.1.2. Unitary models with a household social welfare function

Samuelson (1956) and Varian (1984) point out that if income distribution within the family is itself the result of an optimizing choice rather than arbitrarily determined as in our previous discussion, then even for very general individual preferences, aggregate household demand will behave as if it is the demand of a single maximizer. Suppose that each household member i has a quasi-concave utility function of the form $U_i(x_i, y)$ and that income distribution within the family is decided by a benevolent dictator who has a utility function of the form $W(U_1(x_1, y), \ldots, U_n(x_n, y))$. The dictator could solve for the allocation $(x_1^*, \ldots, x_n^*, y^*)$ which maximizes $W(U_1(x_1, y), \ldots, U_n(x_n, y))$ subject to $px\sum_{i \in H} x_i + p_y y \le W$ and implement this outcome by providing the family with the vector y^* of public goods and giving the family member an income of px_i^*, which i would use to purchase x_i^*. If we define the function $V(x, y)$ to be the maximum of $W(U_1(x_1, y), \ldots, U_n(x_n, y))$ subject to $\sum x_i = x$, then it will be the case that aggregate demand in this family is always chosen in order to maximize $V(x, y)$ subject to the family budget constraint. Samuelson suggests that even if a family is not a dictatorship, it might be that

> preferences of the different members are interrelated by what might be called a 'consensus' or 'social welfare function' which takes into account the deservingness

or ethical worths of the consumption levels of each of the members. The family acts *as if* it were maximizing their joint welfare function.

The joint welfare function that Samuelson has in mind is a function of the form $W(U_1(x_1, y), \ldots, U_n(x_n, y))$ where W is an increasing function of each family member's utility. A utility function of this kind is known as a Bergson–Samuelson social welfare function.[7]

If a family chooses allocations to maximize a Bergson–Samuelson social welfare function subject to a family budget constraint, then the family's aggregate consumption is rationalized by some utility function $V(x, y)$ where $V(x, y)$ is the maximum of $W(U_1(x_1, y), \ldots, U_n(x_n, y))$ subject to the constraint that $\sum_i^N x_i = x$. Therefore it is impossible using data on total family consumption to distinguish the behavior of a family that maximizes a social welfare function from the behavior of a single rational consumer. If W is a concave function of the x_i's then it will also be true that the consumption of each family member is uniquely determined by aggregate family income.

Eisenberg (1961) discovered that if each family member always gets the same fraction of income and if all family members have homothetic, but not necessarily identical, preferences then family demand can be rationalized as the choice of a single individual. This idea was clarified and generalized by Chipman (1974), Shafer (1977), and Shapiro (1977).

3.1.3. Unitary models with the Rotten Kid Theorem

Becker's "Rotten Kid Theorem" (1974, 1981) establishes another set of circumstances under which households act as if they were governed by a single, utility-maximizing decision maker. In Becker's model, there is one consumption good and no public goods. There is a single, benevolent parent and n selfish "kids", who care only about their own consumptions. The parent's utility is given by $U(x_0, \ldots, x_n)$ where x_0 is the parent's consumption and x_i is the ith kid's consumption. Each kid i has an income, m_i. The parent's income m_0 is much larger than that of the kids and he chooses to make "gifts" to each of them. Since the parent wants to make gifts to each kid, the post-gift distribution of consumption in the family is the vector $(x_0^*, \ldots, x_n^*)$, that maximizes the parent's utility, $U(x_0, \ldots, x_n)$ subject to $\sum_{i=0}^n x_i = \sum_{i=0}^n m_i$. If each kid's consumption enters the parent's utility function as a "normal good", then each kid's consumption is an increasing function of total family income. If each person in the family chooses an action a_i that influences the income of other family members but does not influence their utility directly, then it follows that all persons in the family seek to maximize total family income. The problem of choosing the actions a_i is there-

[7] A social welfare function of this type was introduced to economics by Bergson (1938) and further developed by Samuelson (1947).

fore of the type that Marshak and Radner (1972) describe as a problem in "team theory".

It should be recognized that Becker's results are not a trivial consequence of the household head being a "dictator". As Becker (1974) remarks, although the head is able to choose consumption distributions, he is *not* able to dictate the actions a_i that determine individual incomes. Nevertheless, because of the head's distributional actions, all individuals in the family will agree on the same objective function to pursue in their choices of the a_i's, namely maximization of total family income.

Lindbeck and Weibull (1988) and Bruce and Waldman (1990) show that parental altruism can lead to inefficiency in a multi-period model. A similar source of inefficiency in the context of government welfare programs was called the "Samaritan's Dilemma" by Buchanan (1975).[8] In the Bruce–Waldman model, if parents make transfers to their children in the second period, then children will do too little saving. But if parents confine their transfers to the first period, then children will have no incentive to maximize joint family income in the second period. Lindbeck and Weibull extend the analysis to cases where there is mutual benevolence between "parent" and "child". They also point out that an important instance of this problem is the case of children supporting indigent parents. Parents will have too little incentive to save for their old age if the support they will receive from their children is a decreasing function of the aged parents' resources.

Bergstrom (1989a) showed that the Rotten Kid Theorem depends critically on an implicit assumption of transferable utility.[9] The need for this assumption is illustrated by the following example taken from Becker – *the case of the controversial nightlight.* A husband likes to read at night, but the light interferes with his wife's slumber. The husband controls the family budget, but he loves his wife and gives her a generous bundle of consumption goods. He is aware that the nightlight annoys her and because he wants her to be happy, he turns the light out earlier than he otherwise would. But he still uses the nightlight more than she would like him to. One day, while the husband is away, an electrician drops by and offers to disconnect the nightlight in such a way that the husband would not be able to use it again. The wife is convinced that the husband would never know the reason that the nightlight was disconnected. But although she is entirely selfish and dislikes the nightlight, she decides to refuse the electrician's offer. Becker reasons as follows. Although the husband will not *blame* the wife for the loss of the nightlight, he will be made worse off. This will change his utility-maximizing gift to his wife in such a way as to make her worse off, despite her gain in utility from elimination of the nightlight. The effect, according to Becker, is like a loss in family income. If the wife's utility is a "normal good" for the husband, then the

[8] The biblical New Testament parable of the "Prodigal Son", Luke XV 11–32, seems more appropriate to this problem than the parable of the "Good Samaritan".

[9] Johnson (1990) has independently obtained similar results.

effect of a loss in family income is to choose a gift level that leaves her with a lower utility than she had before the nightlight was eliminated.

Bergstrom shows by an example that Becker's conclusion is not in general correct. For a reasonable choice of utility functions for husband and wife, it turns out that even after the husband adjusts his behavior in response to the loss of the nightlight, the wife is better off than before it was disconnected. The reason this happens, is that removal of the nightlight does not necessarily shift the household utility possibility frontier representing private preferences toward the origin in a parallel fashion. Elimination of the nightlight may also change the slope of this utility possibility frontier in such a way that there is a "substitution effect" which induces the husband to give her a higher total utility than she had when the nightlight was available.

3.2. The family with pluralistic decision-making

Just as it is possible in general competitive equilibrium theory to construct useful models without assuming that there is only one consumer, there are many ways to build interesting economic theories of "pluralistic" households. In this section we consider examples of such theories.

3.2.1. Proportional sharing rules

One of the simplest possible models of household consumption assumes that household income is always divided in prespecified proportions between household members and that there are no public goods. Each household member chooses his or her own consumption bundle to maximize utility subject to the resulting budget constraint. Samuelson (1956) calls this division rule an example of a shibboleth and points out that in general, dividing income in proportions that do not change when prices change would be inconsistent with maximizing the utility of a benevolent parent or with maximizing a well-defined social welfare function.[10]

3.2.2. Cooperative Nash bargaining solutions

Manser and Brown (1980) and McElroy and Horney (1981) applied the Nash cooperative bargaining model to marriages. These authors modeled marriage as a static bilateral monopoly in which a married couple can either remain married or divorce. There are potential gains for both parties from remaining married rather than getting

[10] As Samuelson acknowledges, demands resulting from proportional income division would be rationalizable as the demand of a single consumer if preferences were identical and homothetic. As remarked above, Eisenberg (1961) showed that this would be the case even if different household members had different preferences so long as everyone's preferences are homothetic.

divorced. These authors propose that the division of potential gains from marriage is determined by the symmetric Nash bargaining model, where the *threat point* is that they dissolve the marriage.

Specifically, they propose the following model. If they remain married, each partner has a utility function $U_i(x_i, y)$ where x_i is i's vector of consumption of private goods (including leisure) and y is the vector of household public goods that they share. There is a vector of prices p_x for private goods and p_y for public goods. The set of possible household allocations consists of all vectors (x_1, x_2, y) such that $p_x x_1 + p_x x_2 + p_y y = W_1 + W_2$, where W_i is the "full income" of household member i.[11] Given this information, it is possible to construct the utility possibility set within this marriage. The utility that each person can achieve if the marriage is dissolved depends on prices and on his or her full income. Where V_i is i's utility if the marriage is dissolved, the symmetric Nash bargaining solution is the utility distribution (U_1^*, U_2^*) that maximizes $(U_1 - V_1)(U_2 - V_2)$ on the utility possibility set.[12]

If the threat points V_1 and V_2 were independent of prices and individual incomes, then in a household governed by Nash bargaining, aggregate demand could be rationalized by the family utility function $W(x_1, x_2, y) = (U_1(x_1, y) - V_1)(U_2(x_2, y) - V_2)$. Aggregate demand would obey the Slutsky conditions and the revealed preference axioms. But in the model proposed by Manser and Brown and by McElroy and Horney, the threat points V_i represent the utility levels that each person could achieve if he or she were not married. These threat points generally will depend both on prices and on individual incomes. Since the parameters V_j depend on prices and incomes, the family aggregate utility function $\tilde{U}$, unlike the utility function of ordinary neoclassical consumers, depends on prices and on the distribution of incomes as well as on consumption. McElroy and Horney work out the "generalized Slutsky matrix" that corresponds to this situation and show that it will not in general be symmetric; thus total demand would not satisfy the revealed preference axioms, and family demand could not be explained as the behavior of a single rational consumer.

Woolley (1988) questions the assumption that divorce is the appropriate threat point for Nash bargaining between spouses. Woolley examines a model in which the threat point is a noncooperative Nash equilibrium within marriage and another model in which the threat point is a "consistent conjectural equilibrium". Lundberg and Pollak (1993) propose a threat point that is not necessarily a noncooperative equilibrium, but a "division of labor based on socially recognized and sanctioned gender roles". Lundberg and Pollak point out that their model (and Woolley's model) predict an em-

[11] When we treat leisure as a commodity, full income is the value at market prices of a person's initial endowment of nonhuman wealth plus the value of the total amount of labor the person could supply to the market.

[12] This expression is sometimes known as the Nash product. Nash (1950) proposed a set of axioms for resolution of static two-person bargaining games such that the only outcomes that satisfy the axioms maximize the Nash product on the utility possibility set.

pirical outcome that differs strikingly from divorce-threat models and from Becker's Rotten Kid model. If government child-care allowances are paid to mothers rather than to fathers in two-parent households, the threat point envisioned by Lundberg–Pollak and by Woolley is likely to shift in the mothers' favor. Accordingly, the outcomes of cooperative bargaining within households are likely to be more favorable to women. By contrast, in divorce-threat models, the outcome of bargaining depends only on the total resources available to the household and on the utilities that each would receive if they divorced. Whether the nominal recipient of child-care allowances in a marriage is the husband or wife would have no effect on total resources available to the married couple nor would it change the resources available to either spouse if they were to divorce. Therefore the divorce-threat model predicts that such differences would have no effect on allocation within married households. Similarly in Becker's Rotten Kid model, the well-being of each household member is determined by total family income, independently of intrafamily income distribution.

3.2.3. Empirical tests of the unitary model based on private-goods consumption

McElroy (1990) observed that it is possible in principle to test the unitary model of family decision-making, even if one cannot observe consumption of individual household members, using observations of aggregate household consumption and of other variables that could affect an individual's threat point; for example the wage rates and unearned incomes of each household member. If, holding prices and incomes constant, the distribution of income within the household has significant effects on demand, then one would reject the unitary hypothesis.

Stronger tests will be possible if it is possible to determine which household member is the ultimate consumer. For most commodities, this is very difficult to obtain, but there are some interesting exceptions. Sometimes data are available about the amount of leisure consumed by each household member. For a useful discussion and an extensive description of the data available from time-allocation studies in Canada, Europe, Japan, see Juster and Stafford (1991). Ingenious use has also been made of studies in which the nature of the goods strongly suggests the gender of the ultimate consumer.

Schultz (1990) found that in Thailand, an increase in a woman's unearned income from outside the household will have a larger negative effect on the probability that she joins the labor force than does an equal increase in her husband's unearned outside income. According to Browning et al. (1992), in Canada the shares of the family budget devoted to men's clothing and to women's clothing are positively related to the shares of family income earned respectively by men and women. Using data from a household survey from the Cote d'Ivoire, Haddad and Hoddinott (1994) and Hoddinott and Haddad (1995) report that "increases in the proportion of cash income accruing to women significantly raise the budget share of food and lower those of alcohol and cigarettes".

The results found by Schultz, Browning et al., and Hoddinott and Haddad would not be observed in a unitary model of household demand. Schultz is able to peek inside the family black box and observe separate consumptions of leisure by husbands and by wives. Hoddinott and Haddad do not directly observe which household members consume the food or the alcohol and cigarettes, nor do Browning and his coworkers know who wears the trousers in Canadian families. The finding that an increase in the wife's share of family income tends to increase consumption of food and women's clothes and an increase in the husband's share tends to increase consumption of cigarettes, alcohol, and men's clothes is, however, indirect evidence that men in Cote d'Ivoire consume more cigarettes and alcohol than women, and that people in Canada tend to wear gender-appropriate clothing.

3.2.4. The well-being of children, household public goods, and Pareto efficiency

Thomas (1990) found evidence that in Brazilian families, unearned income of the mother has a much stronger positive effect on fertility and on measures of child health such as caloric intake, weight, height, and survival probability than unearned income of the father. For fertility and measures of caloric intake, the effect of mothers' income is about eight times as large as that of fathers' income. For survival probability, the effect of mothers' income is nearly 20 times as large. Hoddinott and Haddad find that in Cote d'Ivoire, children's height for age is positively related to the share of family wealth controlled by their mothers. Schultz finds that in Thailand, an increase in a woman's unearned income tends to increase her fertility while an increase in her husband's unearned income does not.

The results of Thomas, Hoddinott and Haddad, and Schultz on child welfare and fertility, suggest that the distribution of control of resources within the family influences the composition of household demand. But unlike the commodities, leisure, clothing, and alcohol and tobacco discussed above, child health and fertility are *household public goods* jointly "consumed" by both the husband and wife. In the case of private goods, the theory suggested as an alternative to the unitary theory is that the distribution of earnings within the household determines the way in which household expenditure budgets are divided between household members. As Woolley (1988) observed, where the commodity in question is a household public good, even if the wife values the commodity more than her husband, it does not follow that more of the public good will be supplied at Pareto optimal allocations that distribute more utility to the wife.

Suppose, for example, that the household always chooses a Pareto efficient allocation and that utility functions are of the form $U_i(x_i, y) = A(y)x_i + B_i(y)$, where y is child welfare and x_i is i's private consumption. According to Bergstrom and Cornes (1981, 1983), when utility functions are of this form, the Pareto efficient choice of child welfare will be independent of the distribution of income within the household. This will be the case, even if $B_{\mathrm{w}}'(y) > B_{\mathrm{h}}'(y)$ for all y, an assumption that implies that the wife is more concerned about child welfare than the husband, and even if $A'(y) > 0$, an assumption that implies that child care is a normal good for both husband and wife.[13] This is not to say that the results of Thomas, Schultz, and Haddad–Hoddinott are inconsistent with Pareto efficient allocations within the household. Alternative utility functions can be found such that Pareto efficient allocations that give higher utility to the wife are allocations with greater amounts of child welfare. But it is important to realize that if this is the explanation, it rides on stronger assumptions on household preferences than the assumptions that child welfare is a normal good and the wife is more concerned about child welfare than the husband.

An alternative to assuming that allocation is efficient within households is the hypothesis that public goods like child care are provided by voluntary contributions in a noncooperative equilibrium. Woolley (1988) proposed this model of family decision-making as the threat point for Nash bargaining and investigated its comparative statics under the assumption of Stone–Geary utility. Bergstrom et al. (1986) explore the general comparative statics of Nash equilibrium in voluntary public goods supply. Weiss and Willis (1985) suggest that divorced couples, because they are not able to monitor each other's activities, are likely to reach an inefficient, noncooperative equilibrium in supplying resources to their children. In contrast, married couples (who plan to stay married) are likely to be able to sustain an efficient outcome because of the repeated nature of their interaction and because they are able to observe each others' actions closely.

If a household public good is supplied as a noncooperative Nash equilibrium between husband and wife, there are three possible equilibrium regimes: (i) the wife supplies a positive amount and the husband supplies none; (ii) the husband and wife each supply a positive amount; (iii) the husband supplies a positive amount and the wife supplies none. According to Bergstrom et al. (1986), if equilibrium is in regime (i) an increase in wife's wealth relative to husband's would increase expenditures on the public good. If equilibrium is in regime (iii) an increase in wife's wealth relative to husband's would decrease expenditures on the public good, and if the equilibrium is in regime (ii), income redistribution within the family would have no effect on the equilibrium supply of the public good.

One way to explain the finding that redistribution of income toward wives has a strong positive effect on child welfare is to argue that the families observed by Thomas, Shultz, and Haddad and Hoddinott are in noncooperative Nash equilibria where the wife is the only contributor to the public good, child welfare. It would be interesting to find direct evidence that bears on whether the explanation lies here or whether

[13] The reason for this rather puzzling result is that a transfer of private consumption from husband to wife would typically increase her marginal willingness to pay for the public good, child care, and would decrease his. To assume that her marginal rate of substitution between child care and private consumption is higher than his is not sufficient to imply that a transfer of income from him to her will *increase* her marginal rate of substitution by more than it will reduce his.

outcomes in the families reflect a difference in the amount of child care appropriate to different points on the utility possibility frontier.

3.2.5. *Testing the hypothesis of Pareto efficiency with private goods*

The evidence of Schultz, Browning et al., and Haddad and Hoddinott on private goods tends to support rejection of the unitary hypothesis on household demand, but these results provide no direct evidence about whether households allocations are Pareto efficient. Chiappori (1988, 1992) points out that a cooperative Nash solution with divorce as the threat point is not the only alternative to the unitary hypothesis of family decision-making, even if the assumption of Pareto efficiency within households is maintained. Chiappori proposes to test the weaker hypothesis that a family chooses some efficient point on the household utility possibility frontier, using only data on household aggregate consumption. He studies a model in which each member of a married couple consumes a "Hicksian composite" private good and leisure. The two household members are assumed to be "price-takers" both in the goods market and the labor market and are free to work as many hours as they choose. Household aggregate demand is therefore formally the same as aggregate demand in a competitive economy with two consumers and three commodities. According to Diewert (1977), competitive equilibrium in an economy with more commodities than consumers must obey certain empirically falsifiable restrictions. In his 1988 paper, Chiappori spells out Diewert's restrictions as applied to his two-consumer, three-commodity model, both in parametric form and in a nonparametric, revealed-preference form.

3.2.6. *Noncooperative bargaining theory*

The Nash cooperative solution predicts that the outcome of a static, two-person bargaining game will be the outcome that maximizes the product of the two persons' *utility gains* over the *threat point* that would obtain in the absence of agreement. But deciding the appropriate threat point is problematic. Should the threat point be divorce as in Manser and Brown (1980) and McElroy and Horney (1981)? Should it be an uncooperative marriage in which spouses revert to socially sanctioned gender roles for uncooperative spouses as in Lundberg and Pollak (1993, 1994)? If either party to a marriage has the right to divorce the other, should the threat point for each person be the maximum of his or her utility from divorce and from a noncooperative marriage?

The Nash axioms are of no direct help in deciding the appropriate threat point in specific models. But recent work on the noncooperative foundations of bargaining theory offers useful guidance on this question. Rubinstein (1982) developed an extensive form multiperiod bargaining game for two agents in which a cake is to be partitioned only after the players reach agreement. Players alternate in proposing how to divide the cake, with one time period elapsing between each offer. Both agents are impatient; player i discounts future income by the discount factor δ_i. Thus, the utility to player i of receiving w units of cake in period t is $w\delta_i^t$. Rubinstein proved that in the limit as the time between proposals becomes small, the only subgame perfect equilibrium is for the cake to be divided in the first period with player i's share of the cake being $\alpha_i = \delta_i/(\delta_1 + \delta_2)$. More generally, if agent i's utility from receiving w_i units of cake in period t is $u_i(w_i)\delta_i^t$ where u_i is a concave function, then the only perfect equilibrium is the allocation that maximizes the "generalized Nash product", $u_1^{\alpha_1}u_2^{\alpha_2}$ on the utility possibility set $\{(u_1(w), u_2(1-w) \mid 0 \leq w \leq 1)\}$. In the case where the two agents have equal discount rates, this outcome is the same as the symmetric Nash equilibrium corresponding to the threat point (0, 0).

Binmore (1985) shows how the Rubinstein model can be extended to the case where each of the bargaining agents has access to an "outside option". Binmore's model is like the Rubinstein model, except that each agent i has the option of breaking off negotiations at any time and receiving a payoff of m_i units of cake, in which case the other player receives no cake. If agreement is not reached and neither agent exercises an outside option, the utility outcome would be (0, 0) as in the standard Rubinstein model. One might expect that the introduction of outside options would move the threat point to (m_1, m_2). (If negative values of m_i are considered, this conjecture might be amended to $(\max\{0, m_1\}, \max\{0, m_2\})$. Binmore finds, however, that this is not the answer. Instead, it turns out that the only perfect equilibrium for the game with outside options is an agreement in the first period on the utility distribution (u_1, u_2) that maximizes the Nash product $u_1^{\alpha_1}u_2^{\alpha_2}$ on the utility possibility set $\{u_1(w), u_2(1-w) \mid 0 \leq w \leq 1\}$ subject to the constraint that $u_i \geq m_i$ for each i. In general, this solution is not the same as maximizing $(u_1 - m_1)^{\alpha_1}(u_2 - m_2)^{\alpha_2}$ on the utility possibility set, which would be the outcome of shifting the threat point to (m_1, m_2). A similar argument is made by Sutton (1987) and the argument is presented in more detail in a paper by Binmore et al. (1989). The latter paper reports on laboratory tests of a Rubinstein bargaining game with outside options. The laboratory results were better predicted by the Binmore model than by a competing model in which the outside option is the threat point.

Binmore's model of bargaining with an outside option has an interesting interpretation for bargaining models of marriage. Consider a married couple who expect to live forever in a stationary environment. Suppose that in any period, there is transferable utility with the utility possibility frontier $\{(u_1, u_2) \mid u_1 + u_2 = 1\}$. Each spouse has an intertemporal utility function that is a discounted sum of the period-by-period utility flows. Spouse i evaluates the time path $(u_1, \ldots, u_t, \ldots)$ of period utilities by the utility function $\sum_{t=1}^{\infty} u_t\delta_i^t$, where $\delta_i < 1$ is i's discount factor. Let b_i be the utility that spouse i would get in any period where the couple stays married, but does not reach agreement and suppose that if they divorce, then spouse i will get a utility of v_i in every subsequent period. Assume that $b_1 + b_2 = b < 1$ and $m_1 + m_2 = m < 1$. This means that there are potential gains for both persons in reaching an agreement about how to divide utility.

As in Rubinstein, the spouses alternate in making offers of feasible utility distributions. Following Binmore's argument, one finds that in the limit as the time between offers approaches zero, the only subgame perfect equilibrium is one in which the spouses agree immediately to distribute utility in every period in such a way as to maximize the Nash product $(u_1 - b_1)(u_2 - b_2)$ subject to $u_1 + u_2 = 1$ and subject to $u_i \geq m_i$ for $i = 1,2$. Depending on the parameters m_i and b_i, there are three possible types of solution:

(i) Neither of the outside option constraints $u_i \geq v_i$ is binding. In this case, the outcome is $u_1 = b_1 + (1 - b)/2$ and $u_2 = b_2 + (1 - b)/2$. Neither outside option is binding if $b_i + (1 - b)/2 \geq m_i$ for $i = 1$ and $i = 2$.

(ii) The outside option is binding for person 1, but not for person 2. In this case the solution is $u_1 - v_1$ and $u_2 = 1 - v_1$. This happens if $b_1 + (1 - b)/2 < m_1$.

(iii) The outside option is binding for person 2, but not for person 1. In this case the solution is $u_2 = v_2$ and $u_1 = 1 - v_2$. This happens if $b_2 + (1 - b)/2 < m_2$.

The first of these cases corresponds to the Lundberg–Pollak cooperative solution where the threat point is not divorce, but a noncooperative marriage. In the other two cases, the divorce threat is relevant, but notice that the outcome is never the outcome predicated by the Manser–Brown and McElroy–Horney models. When the divorce threat is relevant, there is not an equal split of the gains from being married rather than divorced. Instead one partner enjoys all of the surplus and the other is indifferent between being divorced and being married.

In the absence of agreement one might expect harsh words and burnt toast until the next offer is made. If the couple were to persist in noncooperative behavior forever, the outcome might be worse for one or both persons than being divorced. But divorce (as we have modeled it) is irrevocable, while a bargaining impasse need last only as long as the time between a rejected offer and acceptance of a counter offer. So long as the gains from marriage are divided in such a way that both are better off being married than being divorced, a threat of divorce is not credible. But for many divisions of utility, a threat of delayed agreement and a later counter proposal is credible. In fact the Rubinstein theorem tells us that there is only one equilibrium division of utility in which no such threat is credible.

Rubinstein's original bargaining model can be relaxed in the direction of realism without altering the main results. Binmore (1985) shows that one can relax the assumption that the two parties take turns making offers and that the period between offers is of fixed length. Qualitatively similar results obtain when the length of time between offers and the person whose turn it is to make the next offer are randomly determined after every refusal. It is straightforward to add a constant probability of death for each partner without seriously changing the model. On the other hand, stationarity of the model seems to be necessary for Rubinstein's beautifully simple result. This stationarity is lacking in a model where children grow up and leave the family and where the probability of death increases with age. It would be useful to know more about the robustness of the Rubinstein results to more realistic models of the family. Some interesting beginnings for such an investigation are found in Lundberg and Pollak (1994).

4. Theories of the marriage and household membership

4.1. Matching models

4.1.1. The Gale–Shapley stable marriage assignment

Gale and Shapley (1962) introduced the concept of a *stable marriage assignment* and presented a courtship algorithm that leads to a stable assignment of marriage partners for arbitrary configurations of preference rankings of the opposite sex as possible marriage partners. This model has been extended and developed by several authors (see Roth and Sotomayor (1990) for an excellent survey of this work). Here we follow Roth and Sotomayor in presenting a slightly modified version of Gale and Shapley's model that allows the option of remaining single to persons who do not want anyone who will have them.

Consider a population consisting of n men and p women. Each person i in the population is able to rank all members of the opposite sex as possible marriage partners and also to determine which members of the opposite sex he or she would be willing to marry if remaining single were the only alternative. All persons satisfying the latter condition are said to be *acceptable to i*. A monogamous assignment of marriage partners is said to be *stable* if no two people of opposite sexes would prefer each other to their assigned partners, if no married person would prefer being single to being married to his or her spouse and if no two single people would prefer being married to each other over being single.

The men-propose version of the Gale–Shapley algorithm has each man propose to his favorite woman. Each woman rejects the proposal of any man who is unacceptable to her, and if she receives more than one proposal, she rejects the proposal of any but the most preferred of these. To her most preferred suitor she says "maybe". At each step in the procedure, men who have been rejected move to their next choice so long as there are any acceptable women to whom they have not proposed. Women reject proposals from unacceptable men and from any but the best of their current suitors, including any man to whom she said "maybe" on the previous step. The algorithm continues until a step is reached where no man is rejected. At this point, all women marry the last man to whom they said "maybe".

In general, the assignment resulting from the men-propose version of the Gale–Shapley algorithm is different from that produced by the women-propose version. The difference between these outcomes reveals a remarkable polarity of interest between men and women. When all men and women have strict preference orderings over the opposite sex, the men-propose assignment turns out to be at least as good for *every* man as any other stable marriage assignment and to be at least as bad for every woman as any other stable assignment. Conversely, the women-propose assignment is at least as good for every woman and at least as bad for every man as any other stable assignment. In general there can be other stable assignments besides the men-propose and the women-propose assignments. Where there are more than two stable assignments, the binary relations "is at least as good for all women" and "is at least as good for all men", defined over the set of stable matchings, have the following "lattice property". For any two stable matchings A and B, there is a stable matching C that is at least as good for all women as A and as B. This matching C will be no better for any man than either A or B.[14] Roth and Sotomayor prove some interesting general comparative statics results, including the proposition that adding more women to the marriage market or enlarging some women's lists of acceptable men will (if it has any effect at all) help some men and harm no men and harm some women and help no women.

4.1.2. Marriage markets as a linear programming assignment

Shapley and Shubik (1972) and Becker (1974) suggest that the market for marital partners can be posed formally as the classic *linear programming assignment problem*.[15] The assignment problem is one of the early showcase applications of linear programming techniques (Danzig, 1951). Not only does linear programming offer powerful algorithms for finding an optimal assignment, it also yields dual variables which can serve as "shadow prices" to guide decentralized implementation of the optimum as a market solution. Koopmans and Beckmann (1957) developed the assignment problem model as an economic tool. Their model has n workers and n jobs. Each worker can be assigned to one and only one job. The value of output from worker i in job j would be a specified amount, a_{ij}. An efficient assignment maximizes the total value of output from all workers subject to the constraint that each worker can only have one job and each job must be done by only one worker. The dual solution to this linear program yields a vector of shadow prices for workers and jobs. If w_i is the shadow price for worker i and r_j is the shadow price for job j, it happens that $w_i + r_j \geq a_{ij}$ for all workers i and jobs j and if the optimal solution assigns i to j, then $w_i + r_j = a_{ij}$.

The dual variables to the assignment problem can be given a market interpretation as follows. Imagine that each job has an owner who wants to maximize his profits net of wages and who has to pay any worker his shadow price. If the owner of job j hires worker k, her net profit is $a_{ik} - w_k$. Since $w_k + r_j \geq a_{kj}$ for all k and $w_i + r_j = a_{ij}$ for the worker assigned to firm j by the assignment problem algorithm, it must be that the profit maximizing choice of worker for firm j is the same one assigned to her by the assignment problem algorithm. A similar argument establishes that if workers rather than firms were the residual claimants after paying the shadow prices r_k for jobs, the

[14] A proof of this result, and a nice illustrative example is found in Roth and Sotomayor (1990: p. 37).

[15] A thorough discussion of the assignment model and its extensions is found in Roth and Sotomayor (1990).

profit maximizing choice of job for each worker would be the same one that is assigned to him by the linear programming solution.

The Koopmans–Beckmann model can also be interpreted as a problem of optimal marriage assignments, where one sex plays the formal role of workers and the other the role of firms. For each male, i, and each female, j, there is a number a_{ij} which measures the amount of "marital bliss" that would be produced if i married j. Each male is only allowed to marry one female and each female is only allowed to marry one male. The solution to the assignment problem determines not only who is assigned to whom, but also how the jointly produced marital bliss is "divided" between the partners of the marriages that form. This division is determined by the "prices" in the dual linear program, just as wages and rents are found in the Koopmans–Beckman interpretation.

Shapley and Shubik developed a cooperative game-theoretic interpretation of the assignment model of marriage. Consider the transferable utility game in which the only coalitions that yield nonzero payoffs consist of exactly one man and exactly one woman. The payoff from the coalition consisting of man i and woman j is a_{ij}. An allocation in which partners are assigned by solution of the linear programming assignment problem and payoffs are equal to the corresponding dual shadow prices will be in the *core* of this game. The core of this game consists of the set of allocations such that no two people could improve on their current situation by abandoning their current partners, forming a new partnership and dividing their joint payoff in some way. Shapley and Shubik show that the set of core allocations has the same lattice property that was found for the stable marriage assignments in the Gale–Shapley model.

4.1.3. Marriage assignments with and without transferable utility

The special ingredient that makes it possible to model the matching of marital partners as an assignment problem is transferable utility. The interpretation of marriage as an assignment problem has it that for each male, i, and female, j, there is a number a_{ij} such that if i marries j they will produce a_{ij} units of "bliss" which can be divided between them in any way such that the sum of i's bliss and j's bliss is a_{ij}. On the face of it, the assumption that the total utility from a marriage could be redistributed between the partners just like money or jelly beans seems crude and unreasonable. But, as we showed in the model of transferable utility and mate selection in Section 2.4, a transferable utility framework can accommodate a wide range of interesting and subtle interactions between marital partners.

The Gale–Shapley model is at the opposite extreme from transferable utility. The model contains no allowance for altering the terms of marriage. Since there are no "side-payments", the utility possibility frontier available to any two potential spouses is just a single point. If it is possible for potential marriage partners to draw up premarital contracts which determine in advance the household's choice of public goods and division of private goods, then the Gale–Shapley model is very unrealistic. But if credible and binding premarital contracts are not possible then, as we will see, there are reasonable models in which the relevant part of the utility possibility frontier for any pair is a single point.

Crawford and Knoer (1981) present an ingenious extension of the Gale–Shapley algorithm that works when monetary side-payments can be made but which does not require transferable utility.[16] In their model side-payments are measured in discrete units. Individuals rank options that are specified not only by whom one mates but also by the size of promised side-payments. In the initial round, the side-payment of each member of the proposing sex is restricted to a low level. Proposals, refusals and maybes proceed as in the Gale–Shapley algorithm, but each time a member of the proposing sex is refused, the side-payment that he is allowed to offer to the refusing individual increases by one unit. This process continues until no proposal is rejected. The outcome is a stable assignment of partners. Demange and Gale (1985) show that general models with side-payments but without transferable utility share the lattice property found for the Gale–Shapley model and for the assignment model. These results are well summarized by Roth and Sotomayor (1990).

Kaneko (1982) also analyzes general models of "two-sided exchange economies" that include both the Gale–Shapley model and the transferable utility model as well as markets where there are side-payments without transferable utility. Kaneko shows that for his model, the core is nonempty and coincides with the set of competitive equilibria.

4.1.4. Strategic issues in stable matching

Roth and Sotomayor (1990) present several interesting results about the extent to which assignment mechanisms can be manipulated. If marriages were assigned by the Gale–Shapley algorithm with males proposing to females, then unless there is only one stable matching for the population, there will be at least one woman who will be better off if she misrepresents her preferences. More generally, there exists no stable matching algorithm which would make truth-telling a dominant strategy for all members of the population. In the Gale–Shapley algorithm, with males proposing, it is dominant strategy for the males to reveal their preferences truthfully. In general, a woman would need a great deal of information about others' preferences in order to know how to improve her outcome by deceptive play. But suppose that each woman knew who her match would be in the women-propose version of the Gale–Shapley algorithm and suppose that in the play of the men-propose algorithm, each woman declares a man to be unacceptable if she likes him less well than the mate that she

[16] The Crawford–Knoer results are generalized by Kelso and Crawford (1982).

would be assigned in the women-propose outcome. This configuration of strategies would be a strong Nash equilibrium.[17]

4.2. Household allocation in the shadow of the marriage market

4.2.1. Gifts, commitment, and divorce

Carmichael and MacLeod (1993) propose a theory of gifts as a commitment device in long term relationships. Their theory offers a partial answer to such questions as: Why are courting males expected to offer "inefficient" gifts such as cut flowers, or gift-wrapped, perishable chocolates? What explains seemingly wasteful expenditures on such commodities as engagement rings, wedding rings, and expensive weddings? Carmichael and MacLeod consider an overlapping generations model in which individuals find a partner with whom they will play repeated prisoners' dilemma. After each round in the game, each player has the option of abandoning his or her current partner or offering to play another round. If either partner chooses to abandon, then both must return to the matching market to acquire a new mate. All players who return to the marriage market are assumed to find a match for the next period. At the beginning of a new match, there is no information available regarding an individual's play in previous matches.

In games of repeated prisoners' dilemma, where abandonment is not a possible action, it is well known that for a large range of parameter values, cooperative behavior can be sustained in a subgame perfect equilibrium by "punishment strategies" that offer cooperation so long as one's partner cooperates, and defection for at least some period of time if the partner should defect (see, e.g., Axelrod, 1984). But, as Carmichael and MacLeod demonstrate, if abandonment is possible, there cannot be an equilibrium in which new relationships begin costlessly and all agents cooperate in every period. The reason is simple. If all other agents cooperate at each stage, a player could defect on his partner in the first round, abandon her, and play the same trick on a new partner in the next round.[18] Such a defector would receive the benefits of defecting against a cooperating partner in every round of the game. Carmichael and MacLeod suggest that in such an environment, lasting cooperation could be sustained by a convention in which at the beginning of a new relationship, each partner is expected to give a gift at the beginning of the relationship. In the equilibrium proposed by Carmichael and MacLeod, nobody will be willing to start a new relationship with

[17] A strong Nash equilibrium is a Nash equilibrium that has the additional property that no subgroup of the players would all benefit by changing their actions if the actions of players outside this subgroup are left unchanged.

[18] The model presented by Carmichael and MacLeod is a one-sex model. The main idea extends readily to the two-sex case, but it would be interesting to consider the effects of asymmetries between the sexes, especially asymmetries in the cost of being abandoned.

someone who does not offer the conventional gift. If the cost of the gift is large enough, cooperative behavior can be sustained in Nash equilibrium if everyone chooses the strategy of making the conventional gift and playing *cooperate* so long as his or her partner plays *cooperate* and by abandoning the partner if the partner ever plays *defect*. In such an environment, it does not pay to defect against your partner because if you do, the partner will leave you and you will have to reenter the matching market and present a new gift in order to attract a mate. If the required gift is large enough relative to the gains from defecting on your partner, defection will not be worthwhile. The authors point out that for the gifts to serve this purpose, they should be expensive, but of little benefit to the recipient and certainly not resellable.[19] Otherwise, the cost of buying a gift for your new partner would be nullified by the benefit of receiving a gift when you reenter the marriage market.

4.2.2. Household equilibrium without perfect information

Roth (1992) develops a model of a potentially long-lived partnership in which the partners do not know with certainty how "good" the partnership will be relative to their outside opportunities. If the partnership were known with certainty to be a "good one", the partners would understand that it will be long-lived and will both invest substantial resources in it. As time passes, the partners learn about the quality of their partnership by observing current and past outcomes. The greater the amount of investment in the partnership, the more likely that good realizations will occur in each period. The more good realizations that are observed, the more likely the partners are to invest. In the model as formulated by Roth, no matter how much is invested, "bad" partnerships will eventually be discovered to be bad, but some "good" partnerships are dissolved by rational decision-makers because a run of bad luck has led the partners to think that the partnership will not last and hence the partners do not invest. Roth develops an iterative procedure suggested by dynamic programming which enables him to compute and characterize sequential equilibria for this game.

4.2.3. Divorce as a threat point in bargaining

In the bargaining models discussed in Section 3.2, the payoff from being divorced is determined outside the model. For many people who consider divorce, the utility of the divorce option depends on the utility of forming a second marriage.[20] But the util-

[19] Brinig (1990) offers intriguing evidence that the demand for diamond engagement rings increased dramatically with the abolition in the 1930s of state laws enabling women to bring lawsuits for "breach of promise to marry". Brinig suggests that the engagement rings were "part of an extralegal contract guarantee", which were a prerequisite to premarital sexual intimacy for many couples.

[20] Weiss and Willis (1993) report that in a sample taken in 1985 of Americans who had graduated from high school in 1972, about 60% of those persons who divorced during the period since high school had remarried.

ity of a second marriage must be determined as part of the same theory that determines the distribution of utility within all marriages. Rochford (1984) defines an equilibrium which she calls a symmetrically pairwise-bargained allocation (SPB) that captures this idea. Rochford's model has transferable utility within households. She defines an SPB to be an assignment of partners and an allocation of payoffs within marriages such that the division of utility within each marriage is determined by bargaining, where the threat point is determined by the utility each spouse would get from divorce and remarriage. A person's threat point in a marriage is the highest utility that he or she could achieve as a Nash cooperative solution in some other marriage where the threat points in the other marriage are the utilities the two hypothetical partners get in their current marriages. Rochford proposes an iterative process that is guaranteed to converge to an SPB. Moldevanu (1990) considers a model of trading partners and formulates an equilibrium concept similar to Rochford's. He is able to extend her results to economies without transferable utility.

The models proposed by Rochford and by Moldevanu do not include explicit costs of divorce. While the costs of switching from one partner to another may be small for trading partners, this is not likely to be the case for marriage partners, who are likely to have invested significant amounts of "marriage-specific capital" that will be lost if they divorce. If the partners have children, then arrangements for sharing the costs and joys of child care become difficult and inefficient.[21] In societies where divorce is unusual, divorced people are sometimes ostracized or at least suspected of being unusually difficult to live with.

The introduction of costs of divorce will markedly affect the workings of the formal model. As Moldevanu points out, in a model like Rochford's, if each person has at least one "clone", then any core allocation, including the symmetrically pairwise-bargained allocations would have the property that identical people must be equally well off. In this case, the SPB allows no scope for bargaining within households. In equilibrium, each married couple would correspond to another couple just like them with the same payoffs. The threat point of each individual would be the same as the utility he or she obtains in equilibrium. In the absence of clones, if there were very close substitutes for each person in the society, spouses would not have much surplus to bargain over, once each is given at least his or her outside option.

4.2.4. Household bargaining with outside options

Shaked and Sutton (1984) discuss a model of labor and management which is formally similar to a model of a marriage in which divorce is costly. In their model, the firm has a current work force which it cannot replace immediately or costlessly. They impose these costs by putting a restriction on the timing of offers in a Rubenstein bar-

[21] See Weiss and Willis (1985) for a discussion of incentive problems that arise for child care in divorce settlements.

gaining model. This leads to an outcome that is intermediate between a bilateral monopoly outcome where neither side has an outside option and the "Walrasian" outcome, which would obtain if there were no costs to a firm from changing its workforce.

Rubinstein and Wolinsky (1985) have a model of transactions between pairs of agents who meet randomly and bargain if they meet. Their model, viewed as a model of marriage, offers an interesting interpretation of the "costs of divorce". There are two types of agents, buyers and sellers. All agents of a given type are identical. At the beginning of each time period, there is a matching stage, where each agent tries to find a new partner. Some agents will find partners, some will not. Any buyer and seller who meet will start to bargain according to a noncooperative iterative bargaining scheme. If these two agents reach agreement, a transaction occurs and they leave the market. If they do not reach agreement in this period, there is a chance that one or both of them will meet another agent of the opposite type. If this happens, the agent ceases bargaining with his or her current bargaining partner and starts bargaining with the newly met partner. If neither meets a new partner, the current partners proceed together to the next round of bargaining. The cost of not reaching agreement in the current period is now twofold. If agreement is ultimately reached with the current partner, there is a cost of delay. In addition there is the risk that one's current bargaining partner will meet someone else before the next round of offers. If one is abandoned by one's current partner, one will not be sure to meet anybody to bargain with in the next period. When the number of buyers does not equal the number of sellers, it takes longer on average for the abundant type to find a new partner than it does for the scarce type. Because of this, the abundant type will be willing to concede a larger share of the gains from agreement than will the scarce type.

Binmore's version of the Rubinstein model with outside options, discussed in the previous section, has strong and interesting implications if all people who divorce eventually remarry, but face a transactions cost in the process. Consider the special case of a large population of identical males and of identical females. A male and female who marry and who reach agreement can achieve any constant flow of utility (u_m, u_f) such that $u_m + u_f = 1$. Utilities are normalized so that the utility flow while the partners are in disagreement is 0 for each. At any stage in the bargaining process, either spouse can either accept the other person's offer, reject the other person's offer and make a counter offer, or ask for a divorce. If the two spouses have equal time rates of discount, then in equilibrium, according to Binmore's results, the outcome will be an allocation of utility (u_m, u_f) that maximizes the Nash product $u_m u_f$ subject to the constraints that $u_m + u_f = 1$, and that each person gets a utility at least as high as his or her outside option. The utility distribution $(\bar{u}_m, \bar{u}_f) = (1/2, 1/2)$ maximizes $u_m u_f$ subject to $u_m + u_f = 1$. Given that the equilibrium distribution of utility in a marriage is $(\bar{u}_m, \bar{u}_f)$ a person who divorces and remarries will have to bear a divorce cost of c_m if he is male and c_f if she is female. Therefore the utility of a male who takes the outside option of divorce and remarriage is $\bar{u}_m - c_m < \bar{u}_m$ and the utility of a female who chooses this option is $\bar{u}_f - c_f < \bar{u}_f$. Therefore so long as divorce costs are positive for both parties, the presence of the outside option does not influence the bargaining outcome.

In this model, unlike the Rubinstein–Wolinsky model, the distribution of utility within marriages does not depend on the relative supplies of males and females, but only on their impatience and on the position of the utility possibility frontier relative to the noncooperative outcome within marriage. The difference seems to lie in the fact that in the Rubinstein–Wolinsky model, an individual who "meets a stranger" of the opposite sex can abandon his or her spouse without bearing any transaction cost (although the abandoned spouse may be in for a long wait before another offer appears). In the variant of the Binmore model just proposed, a threat by either party to abandon the current marriage is not credible because persons who divorce would have to pay the transactions cost of divorce and remarriage and when they are done with this, would be in no better bargaining situation than they were before divorcing. These ideas can readily be extended to a community with many types of males and females and very general utility possibility frontiers within each possible marriage.

When there are divorce costs to both parties, the Binmore argument implies the striking conclusion that even though technology and preferences allow transferable utility, the nature of bargaining determines that there is only one possible distribution of utility between any two people if they should marry. This outcome is the Nash cooperative solution with uncooperative marriage as the threat point. Despite the presence of side-payments and the availability of remarriage as an outside option, the original Gale–Shapley model without side-payments then applies. This means that the marriage market suffers from rigidity of a "price" which may be important for clearing the marriage market. If it were possible to settle the distribution of utility within possible marriages in advance by a binding contract, then the distribution of utility between males and females within marriages would respond to competitive forces in such a way as to tend to equilibrate the number of males and the number of females who choose to marry at any time. If, on the other hand, the distribution of utility within marriages is determined by a threat point such as uncooperative marriage, which is independent of market forces, then changes in the terms of marriage cannot be expected to equalize imbalances in supplies of the two sexes in the marriage market.

4.3. Age at marriage

4.3.1. Classical one-sex population theory

Classical Stable Population Theory as developed by Lotka (1922), shows that if current age-specific fertility and mortality rates for females remain unchanged, then, in the long run, the age distribution of the population would asymptotically approach some constant distribution, and therefore raw birth rates and population growth rates would approach constancy. This fact has encouraged demographers to project hypothetical long run population growth rates implied by current age-specific female fertility. Since every baby has a father as well as a mother, it is possible (in societies that have good data on paternity) to construct tables of age-specific rates at which the current population of males father children. These data can be used to make an alternative projection of long run population growth rates parallel to the projections made using a female one-sex model. Perhaps surprisingly, when the male one-sex model is applied to actual populations, the predictions are often quite different from those found by applying it to females. For example, using 1968 data in a one-sex model for US males would predict a long term population growth rate of 10.1 per 1000 population, while the female one-sex model would predict a long term population growth rate of 5.7 per 1000 population (Das Gupta, 1973). Since every child that is born, must have exactly one mother and one father, it is simply numerically impossible that both sexes would maintain the same age-specific fertility and mortality rates after 1968 as did the population surveyed in 1968. Since every wedding also involves one male and one female, the same logical difficulty is present in efforts to predict future marriage rates of males and of females separately, by projecting current age and sex-specific marriage rates into the future.

4.3.2. Two-sex theories of mating

Modern demographers (Keyfitz, 1971; Das Gupta, 1973) have responded to this discrepancy by building two-sex models based on "marriage functions" which ensure the necessary parity between male and female parents or wedding partners. As applied to marriages, these models predict that the number of marriages between a female of age i and a male of age j in year t should depend at least on the number of males and the number of females present in year t. McFarland (1972) criticizes these models because they do not adequately reflect the possibilities for substitution among various cohorts. To allow these possibilities, the number of marriages between a female of age i and a male of age j should depend on the numbers of males and females of other ages in the population as well. McFarland suggests an iterative procedure (which bears an interesting similarity to the Gale–Shapley model) for dealing with these effects.

Pollak (1986, 1987, 1990) reformulates the "two-sex problem" by replacing the constant age-specific fertility schedule of the classical theory with two more fundamental relationships. These are a "birth matrix" and a "mating rule". The birth matrix postulates an expected number of births per period from a marriage of an age i male to an age j female. The mating rule is a function that determines the number of marriages of type i males to type j females for all i and j as a function of the vector listing the numbers of males and females of each age in the population. Pollak shows that if these fundamental relationships remain constant over time and if the mating rule follows certain natural conditions, the resulting dynamical system will converge to a constant equilibrium growth rate, yielding a constant equilibrium age structure. Pollak imposes only certain very general conditions on the mating function such as nonnegativity, homogeneity, continuity, and that the number of persons of a given age and sex who marry must not exceed the number of persons of that age and sex in the population.

4.3.3. Transferable utility model suitable for empirical estimation

Pollak's mating rule is a "reduced form" description of the dependence on the outcome of a marriage market on supplies and demands of the two sexes from various cohorts. Bergstrom and Lam (1989a,b) construct a model of the marriage market that rationalizes Pollak's mating rule. Their work concentrates on reconciling the numbers of males and females who are willing to marry in any given year. In the absence of side-payments, two arbitrarily selected persons would usually disagree about their preferred wedding date. Suppose, for example that all males prefer marrying at age 25 and all females prefer marrying at age 23. A male and a female will agree about the best time for them to marry only if the male was born two years earlier than the female. But in a population where cohort sizes change over time, there will not always be an equality between the number of females of one cohort and the number of males of a cohort born two years earlier. If males prefer to marry at an older age than do females, then if there is a "baby boom", females in the boom generation will find a shortage of males who want to marry when they do. Males born at a time when the birth rate is falling will find a shortage of females two years younger, who will want to marry when these males are 25. When members of one sex and cohort are in excess supply relative to their "natural partners", there will be readjustments in which some of the abundant group postpone marriage and some of the scarce group marry earlier than they otherwise would.

Bergstrom and Lam propose a simple overlapping-generations model of the marriage market, designed to deal with this problem. This model has enough special structure so that its parameters can be empirically estimated. Utility is assumed to be linear in consumption and quadratic in age at marriage. Utility of a person whose preferred age at marriage is a^* and who consumes c units of consumption good and marries at age a is $c - (a - a^*)^2$. In the simplest form of this model, suppose that all males have preferred age of marriage a_{m}^* and all females have preferred age at marriage a_{f}^*. Suppose also that the income that each individual brings to a marriage is independent of whom he or she marries. Suppose that male i has income I_i and was born in year b_i while female j has income I_j and was born in year b_j. If they marry, they will both have to choose the same date of marriage, so the date of their marriage is a "household public good". The assumption of quasilinear utility implies that there is a unique Pareto optimal wedding date for this couple. Given the quadratic specifications of utility of age at marriage, this date is the midpoint between the two partners' preferred wedding dates. The preferred wedding date of male i is $b_i + a_{\mathrm{m}}^*$, the preferred wedding date of female j is $b_j + a_f^*$, and the Pareto optimal date for their wedding is $(b_i + a_m^*)/2 + (b_j + a_f^*)/2$.

Let us define d_{ij} to be the number of years that separate the preferred wedding dates of male i and female j. Then $d_{ij} = |(b_j + a_f^*) - (b_i + a_m^*)| = |(b_j - b_i) - (a_f^* - a_m^*)|$. Each partner's actual wedding date will differ from his or her preferred wedding date by $d_{ij}/2$. The feasible consumption allocations (c_i, c_j) for this couple must satisfy the equation $c_i + c_j = I_i + I_j$. Therefore the couple's utility possibility frontier is described by the equation $u_i + u_j = a_{ij}$ where $a_{ij} = I_i + I_j - d_{ij}^2/2$.

Let the numbers of surviving males and females born in year i be M_i and F_i. The linear programming assignment model predicts that the pattern of marriages will solve the following maximization problem: Where X_{ij} represents the number of marriages between males born in year i and females born in year j, solve for the values of X_{ij} that maximize $\sum_i\sum_j a_{ij}X_{ij}$ subject to the constraints, $\sum_i X_{ij} = F_j$ for all j and $\sum_j X_{ij} = M_i$ for all i. Since we have assumed that incomes are independent of whom one marries, the optimizing solution for the X_{ij}'s is independent of the distribution of incomes and can be determined by minimizing $\sum_i\sum_j d_{ij}^2 X_{ij}$ subject to the constraints $\sum_i X_{ij} = F_j$ for all j and $\sum_j X_j = M_i$ for all i.

In the simple model proposed here, the only parameters to be estimated are the preferred marriage ages a_m^* and a_f^* of males and females. Any specification of these parameters determines the matrix of d_{ij}'s. This information together with an empirically observed distribution of age-cohorts by sex will determine an optimal assignment of marriage partners by cohort. Estimation can proceed by choosing the values of a_m^* and a_f^* that best predict the patterns of actual marriages. More flexible functional forms and some variation of preferences among individuals can also be accommodated within this model, in fairly obvious ways. Bergstrom and Lam (1989a) applied this technique to Swedish historical data on marriage rates in the 19th and 20th centuries.

4.3.4. Why do women marry older men?

One of the strongest demographic regularities is the observation that men marry later in life than women. In a study conducted by the United Nations,[22] the average age of marriage for males exceeded that for females in each of 90 countries and in every time period studied between 1950 and 1985. The age difference tends to be larger in traditional societies than in modern industrial countries and has diminished over time in most industrial countries.

Bergstrom and Bagnoli (1993) proposed an explanation for this difference. They suggest that, at least in traditional societies, women are valued as marriage partners for their ability to bear children and manage a household, while men are valued for their ability to make money. Information about how well a male will perform economically

[22] *Patterns of First Marriage: Timing and Prevalence* (1990).

– whether he is diligent and sober – becomes available at a later age than the relevant information about how well a female would perform her household roles. This leads to an "intertemporal lemons model", in which males who expect to do poorly in later life will seek to marry at a relatively young age and males who expect to prosper will postpone marriage until their success becomes evident to potential marriage partners. Females, on the other hand, marry relatively early, with more desirable females marrying the successful, older males who postponed marriage and the less desirable females marrying the young males who want to marry young. In equilibrium, a young male who attempts to marry is signaling a lack of confidence in his future economic prospects. While the most desirable females would not accept such males, the less desirable females have no better alternatives in the marriage market and hence are willing to marry young males.

This theory implies not only that males tend to marry later in life than females, but also that males who marry young will tend to be less prosperous in later life than males who postpone marriage. Bergstrom and Schoeni (1992) investigate the empirical relationship between age-at-first-marriage and lifetime income, for males and for females. Using 1980 US Census data, they plot wage income of males in later life as a function of the age at which they married. Income is highest for those who marry in their late 20s. Men who marry at age 28 or 29 have average earnings about 20% higher than men who marry at 18.

4.4. Alternative household structures

Most of the work by economists on the theory of the household has concerned either single-person households or monogamous couples, with or without children. There is, however, considerable evidence that nonmonogamous modes of household organization are too significant to ignore.

4.4.1. Polygyny in marriage markets

Becker (1981) devotes a chapter of his *Treatise on the Family* to "Polygamy and Monogamy in marriage markets". Becker's analysis of polygamy is more than a clever curiosum; it extends methods of economic analysis to a major social institution that has received all too little attention from economists. Although overt polygamy is rare in our own society, it is a very common mode of family organization around the world. Polygyny (men having multiple wives) is prevalent in 850 of the 1170 societies recorded in *Murdock's Ethnographic Atlas* (Murdock, 1967), while official polyandry (women having multiple husbands) is prevalent in only a handful of societies (Hartung, 1982).

One of the first economic issues that must be confronted by a polygynous society is the question of how are wives allocated. Not surprisingly (to economists at least), the price system usually comes into play. Becker suggests that theory would predict higher incomes for women under polygyny than under monogamy. He reasons that relaxing the constraint that a man can have only one wife would shift the demand schedule for wives upward, leading to higher bride prices with polygyny than with monogamy. The argument that polygyny leads to higher bride prices is theoretically compelling and appears to be supported by anthropological evidence.[23] It does not, however, follow that higher bride prices imply welfare gains for females. If "property rights" to an unmarried female lie with her family, it seems plausible that her family would use the proceeds from the sale of a bride to purchase a wife or an additional wife for one of her male siblings. This theoretical prediction appears to be strongly supported by anthropological field studies (see Goody, 1973).

4.4.2. *Unwed parents*

Economic theorists have done little work on extending bargaining models of sexual relationships and child support to noncohabiting, unmarried parents. This neglect might have been excusable 30 years ago on the grounds that the most children were born into households with two cohabiting adults. Recent statistics show that unwed parenthood is no longer rare. In the United States in 1960, only 5% of all births occurred out of wedlock. In 1990, more than 25% of births were to unwed parents. (About 30% of the unwed parents in 1990 were cohabiting couples.) The proportion of all children who live in single-parent, mother-only households has risen from 8% in 1960 to 23% in 1990. For Black Americans, the statistics are even more dramatic. In 1990, two-thirds of births were out of wedlock and more than half of all children live in single-parent households.[24]

Willis (1994) studies some of the interesting theoretical issues that arise in the analysis of unwed parenthood. Willis begins with an analysis of fertility decisions and child care expenditures for a single mother who is not able to identify the father(s) of her children. He then considers an equilibrium model of child support for noncohabiting parents. In this model, the father's identity is known and both parents care about the well-being of a child. Since they do not live together, it is difficult for them to monitor each other's behavior sufficiently to sustain efficient cooperative arrangements for child support. Willis examines a noncooperative Stackelberg equilibrium where the mother has custody of the child and the father can influence expenditure on the child only by transferring income to the mother. This equilibrium will not in general be efficient. Marriage, Willis argues, is likely to lead to more efficient, cooperative arrangements for child care between mother and father. The question arises: If it

[23] According to Gaulin and Boster (1992), about two-thirds of the societies found in Murdock's *Ethnographic Atlas* have positive bride prices, while in only about 3% of these societies is it the case that brides must pay a dowry to the husband. Moreover, according to Gaulin and Boster, almost all of the societies with dowries are monogamous.

[24] These statistics and many interesting related facts are reported by DaVonza and Rahman (1994).

is more efficient for the two parents of a child to live together than apart, why is unwed motherhood so common? Willis suggests some possible reasons. One force for unwed parenthood that leads to a particularly interesting analysis is imbalance between the number of marriageable women and the number of marriageable men. This explanation seems particularly compelling for the Black population. Wilson (1987) argues that women's search for partners will be confined primarily to a pool of "marriageable males" – males who would bring resources to a marriage. For statistical purposes, he identifies this pool with males who are currently employed. Wilson found that in 1980, the ratio of black marriageable males aged 20–44 to black females aged 20–44 was about 0.56 in the Northeast and North Central states of the US. (In 1960, this ratio was about 0.67.) The corresponding ratio of white marriageable males to females was about 0.85. Following Wilson's suggestion, Willis works out an equilibrium model in which men choose between monogamy and a polygynous life in which they father children by several women but marry none of them. Monogamous men are confined to a single mate. A polygynous life will have some advantages, because a man may father children by more women, and some disadvantages, including the inefficiencies in child care arrangements that arise when parents do not live together. In Willis's model, there is a threshold expected number of partners P such that men will be indifferent between monogamy and a polygynous life if their expected number of partners in polygyny is equal to P. Suppose that there are more women who want to have children than the number of marriageable males available to them, but suppose that there are not enough marriage females so that every male could have P partners. Then there would be an equilibrium in which some marriageable males (and an equal number of females) are monogamous and some marriageable males do not marry, but father children by P different women. As Willis shows, this model leads to an interesting algebra of a society with a mixture of monogamy and unofficial polygyny.

5. Interdependent preferences within families

5.1. *Benevolence and other forms of unselfishness*

5.1.1. *Preferences on allocations*

If household members love each other, copy each other, envy each other or annoy each other, then individuals care not only about their own consumptions, but also about the consumptions of other members. In the most general case, each member's utility would depend on the amount of each private good consumed by each member of the household as well as about the amount of each household public good.

It is often useful to consider a model of household interdependence that is intermediate between a fully general model of interdependence and the case where consumers care only about their own consumptions of private goods and the vector of household public goods. An interesting and much-studied assumption is that preferences on allocations are "weakly separable" between one's own consumption and that of others.[25] The assumption that consumer i's preferences are separable with respect to his own consumption means that i's preferences among alternative bundles (x_i, y) of private goods and household public goods are not changed by changes in the consumption bundles of others. In this case, a person may *care* about what other family members consume, but their consumption does not influence one's preferences about one's own consumption. In this case, each individual i has a well-defined "private utility function" $v_i(x_i, y)$ that represents i's preferences on private goods for himself given the vector y of public goods.

In a model with private goods only, Winter (1969) and Bergstrom (1971a,b) define preferences of consumer i to be benevolent (nonmalevolent) if there is weak separability and every family member favors (does not object to) a change in another family member's consumption that ranks higher in that person's private preferences. If there is benevolence (nonmalevolence), then preferences of every person i can be represented by a utility function of the Bergson–Samuelson form, $U_i(v_1(x_1, y), \ldots, v_h(x_h, y))$. where U_i is an increasing (nondecreasing) function of v_j. Archibald and Donaldson (1976) define preferences that can be represented by utility functions of the form $U_i(v_1(x_1, y), \ldots, v_h(x_h, y))$ where U_i is not necessarily monotone increasing in its arguments to be *nonpaternalistic preferences*. They point out that nonpaternalistic preferences permit not only nonmalevolence and benevolence, but also malevolence as well as preferences for equity such that U_i may not be monotonic.

5.2. *Interdependent utility functions*

When family members love (or hate or envy) each other, their interlinked joys and sorrows may feed on each other in curious ways. No matter how these feelings are entwined, economists concerned with resource allocation are likely to be more interested in derived "reduced form" preferences over allocations of goods than in a tangle of interrelated preferences about the happiness of others. Therefore, although preferences over household allocations may be founded on interrelated preferences, economists are likely to want to disentangle the interrelated utilities of family members and find the corresponding derived preferences on allocations. This problem has been addressed by several economists, including Bergstrom (1971b, 1988, 1989b), Barro (1974), Becker (1974), Pearce (1983), Kimball (1987) and Bernheim and Stark (1988).

Bergstrom (1989, 1990) studies models in which there is a group of consumers whose happiness depends on their own consumption and on their perceptions of the happiness of other members of the group. Then the happiness of each person can only

[25] A thorough treatment of a variety of separability assumptions is found in Blackorby et al. (1978).

be determined if one knows the happiness of each of the others. One resolution of this paradoxical simultaneity is to suppose that each individual's happiness is observable by others, but with a lag. Each person's current happiness depends on his or her own current consumption and on her observation of the happiness of all other family members in the previous period. With this structure, the time path of happiness for each person is determined as a system of difference equations.[26]

As a concrete example, consider a family with h members. Let $c_i(t)$ be family member i's consumption bundle at time t and let $U_i(t)$ be i's utility at time t. Suppose that utility interdependence takes the additive form

$$U_i(t) = u_i(c_i(t)) + \sum_{j \neq i} a_{ij} U_j(t-1),$$

where the constant a_{ij} represents the marginal effect of person j's happiness in the previous period on person i's current happiness. This system of difference equations can be written as a matrix equation $U(t) = u(c(t)) + AU(t-1)$, where $c(t) = (c_1(t), \ldots, c_h(t))$, $u(c(t)) = (u_1(c_1), \ldots, u_h(c_h))$, $U(t) = (U_1(t), \ldots, U_h(t))$ and A is the matrix with zeroes on the diagonal and with $A_{ij} = a_{ij}$ for $i \neq j$.

Let us evaluate the path of utilities in the case where each family member receives a constant consumption over time so that $c(t) = c$ in every period. Suppose that in period 0, family members start with an arbitrary distribution of utilities $(U_1(0), \ldots, U_h(0))$. If the eigenvalues of the matrix A all have absolute values less than unity, the distribution of utilities will converge to a constant vector that we will define to be $U(c)$. This equilibrium distribution of utilities must satisfy the equation $U(c) = (I - A)^{-1}u(c)$. As Pearce (1983) and Bergstrom (1988) observe, when utility interdependence is nonmalevolent, the matrix A_{ij} is nonnegative and the formal structure of the model is the same as that of Leontief input–output matrices. The theory of productive Leontief matrices[27] can be borrowed to good effect. A nonnegative matrix A is said to be *productive* if there exists some positive vector x such that $(I - A)x$ is a strictly positive vector. Gale proves the following properties of productive matrices:

(i) If A is a nonnegative, productive matrix, the matrix $(I - A)^{-1}$ exists and is nonnegative in every element.

(ii) A nonnegative matrix A is productive if and only if all eigenvalues of A are smaller than one in absolute value.

From property (i), it follows that if A is a productive matrix, then where the allocation c of consumption over time is constant, there must be a unique limiting distribution of

[26] In a paper which proposes several interesting models of interdependent utility, Pollak (1976) introduces the idea of using lagged, rather than simultaneous independence.

[27] For an elegant treatment of productive matrices, see Gale (1960). An equivalent condition is known as the Hawkins–Simon condition. Yet another equivalent condition is that the matrix $I - A$ be "dominant diagonal". See McKenzie (1960).

utility $U(c)$ such that $U(c) = u(c) + Au(c)$. Thus $U(c) = (I - A)^{-1}u(c)$. Writing out in full, the implied utility functions on allocations, we have $U_i(c_1, \ldots, c_h) = \sum_{j=1}^{h} a_{ij}b_{ij}u_j(c_j)$ where $b_{ij} \geq 0$ is the ijth element of the matrix $(I - A)^{-1}$.

The requirement that A be a productive matrix limits the strength of benevolent interdependence. For example, in a two-person family, A will be a productive matrix if and only if $a_{12}a_{21} < 1$. Bergstrom (1971b, 1989b) shows that for two persons, a system of *superbenevolent* interdependent utilities in which $a_{12}a_{21} > 1$ has the property that at all Pareto optimal allocations, disagreements between the two persons take the form of each wanting the other to have the better part.[28] In the case where there are more than two persons, the matrix A will be productive if $\sum_{j=1}^{h} a_{ij} < 1$ for all i. If there is nonmalevolence and the matrix A is not productive, then the dynamical system implied by the equation $U(t) = u(c) + AU(t-1)$ is not stable. This would imply that starting from certain configurations of utility, although consumption of each consumer is constant, the interrelated happinesses would feed on each other and diverge. The dynamics of unstable utility interactions have not yet been studied by economists.

5.3. *Intergenerational utility interdependence*

Utility interdependence in families does not begin and end with a single nuclear family. Everyone's parents were children of parents who were children of parents and so on.[29] Samuelson (1958) pioneered formal modeling of an "overlapping generations" economy, in which a new generation appears in every time period, and each generation ages and dies. In Samuelson's model, there is no benevolence between parents and offspring. Each newborn enters the world, not as a helpless baby, but as a rational decision-maker aware that she has a specific pattern of endowments of labor to sell over the course of her life. Her encounters with preceding and subsequent generations are entirely commercial – borrowing or lending to smooth her lifetime consumption.[30]

Strotz (1955) argues that individual preferences need not be time-consistent in the sense that if one makes an optimal lifetime consumption plan from the viewpoint of the present, one's "future self" may choose not to abide by this plan. In the absence of time-consistency, Strotz suggests two possible theories of consumer behavior. These theories, which are clarified and refined by Pollak (1968) and by Blackorby et al. (1978), are known as theories of "naive" strategies and of "sophisticated" strategies. A person with a naive strategy takes the first step of the intertemporal consumption plan

[28] Matzkin and Streufert (1991) present an interesting example in which *supermalevolence* leads to paradoxes similar to those induced by superbenevolence.

[29] It is tempting to say that every child will be a parent of children who will be parents, but of course not everyone has children. Most economic models of overlapping generations do not, however, take this fact into account.

[30] Diamond (1965a,b) extends this model to allow accumulation of capital and to study the effects of national debt.

that is optimal given his current preferences while making the (incorrect) assumption that in the future he will stick to this plan. In an equilibrium of sophisticated strategies, a person with intertemporally inconsistent plans chooses his current consumption, knowing that in the next period, his preferences over the future will not be consistent with his current preferences. If he knows what these preferences will be, then in equilibrium, each period's choice will be optimal for that period based on what he knows will be chosen in future periods.

As Phelps and Pollak (1968) and Blackorby et al. (1978) suggest, the Strotz model is a natural starting point for a theory of interaction between benevolent parents and their descendants. Let c_t be the consumption vector of generation t and ${}_tc$ be the vector $(c_t, c_{t+1},\ldots)$ specifying the consumption of generation t and each subsequent generation. Then a person in generation t has preferences represented by a utility function of the form $U_t({}_tc)$. The Strotz model would allow a member of any generation t to choose its own c_t and to leave an inheritance to its successor generation. The next generation in turn is allowed to choose its consumption and the inheritance it leaves to its successor. A mother who follows a naive strategy chooses consumption and saving based on the (generally incorrect) assumption that her descendants will dispose of her inheritance in the same way she would wish them to. A mother who follows a sophisticated strategy chooses her preferred amount of saving in the knowledge that her daughter will spend her inheritance in a way that is optimal from the daughter's point of view.

Koopmans (1960) studied conditions on utility functions that guarantee time-consistency. Where ${}_tc$ is the vector $c_t, \ldots, c_m$ of consumption in time periods from t until the end of the decision-maker's life, Koopmans showed that if preferences are stationary, additively separable between time periods and time consistent, then (subject to some technical conditions) it must be that preferences of the individual in time t are representable by a "time-discounted" utility function of the form $U_t = \sum_{\tau=t}^{m} \alpha^{\tau} u(c_{\tau})$. If weak separability rather than additive separability is assumed and if the time horizon is infinite, then time consistent utility functions must take the recursive form $U_t({}_tc) = V(c_t, U({}_{t-1}c))$.

Naive application of single-person intertemporal models to family dynasties lack one important feature of modern economic life: the illegality of slavery. It is natural in a single-consumer model to allow the consumer to borrow on future income, even if he is not able to commit his future selves to a particular course of action. In the intergenerational interpretation, people are allowed to leave positive inheritances, but they are not allowed to sell the future labor services of their descendants and thus enhance their current consumption. Laitner (1979a,b, 1988, 1990) published a series of papers that explore bequests, saving and debt in models where parents cannot extract wealth from their descendants, under various assumptions about mating patterns.

One of the most influential applications of recursive intergenerational utility is Barro's (1974) paper, "Are Government Bonds Net Wealth?". Barro argues that if utility functions take the form $U_t({}_tc) = V(c_t, U({}_{t-1}c))$ and if each generation voluntarily leaves an inheritance to its successor, then government programs which impose intergenerational transfers (for example subsidized education, social security, and government debt) will be offset by corresponding changes in inheritance.

Barro finds neutrality in a model where reproduction is asexual or the only mating is between siblings. There are no marriages and no connections between family lines. Bernheim and Bagwell (1988) suggest that if Barro's model is to be taken seriously, then it must also apply in a model with intermarriage between families. Bernheim and Bagwell argue that if a daughter from one family marries a son from another family and if both parental families leave inheritance to the bride and groom, then a small income transfer from one parental family to another would be undone by offsetting changes in the inheritances of the two sides of the family. But this is only the beginning. If the bride and groom each have a sibling who marries someone else, then the two families that were directly linked by marriage will be indirectly linked to a third and fourth family, which in turn will be linked to other families. Since transfers between directly linked families are offset by changes in gifts, income transfers between indirectly linked families will likewise be fully offset, through a chain reaction of changes in gifts along the path of marriages relating these families.

Bernheim and Bagwell apply simulations and offer corroborating arguments from random graph theory to show that with reasonable models of mate selection, there is a very high probability that any two families in large finite populations will be indirectly linked by marriage where the links connecting people span only two generations. If it were the case that for all marriages, both sets of parents-in-law left inheritance to their offspring, then with very high probability, almost any small governmental income redistribution would be undone by offsetting private actions. Bernheim and Bagwell find this implausible and suggest that it is likely that there are large numbers of breaks in the chain, that is instances where one or both sets of parents-in-law do not leave estates to their children. Where there are many breaks in the intergenerational chain of giving, Barro's neutrality result cannot be expected to apply.

Laitner (1991) proposes a model in which marriage is not random but strongly assortative on income so that persons who expect large inheritances will marry others who expect similarly large inheritances. In Laitner's model, the cross-sectional neutrality found by Bernheim and Bagwell is absent because marriages between children from families of significantly different income levels are rare and when they do occur, typically the less wealthy parental family will leave no estate to the young couple.

An economically and mathematically interesting structure arises when each generation cares not only about its own consumption and the utility of its successor but also about the utility of its parent generation. Kimball (1987), models "two-sided altruism" by assuming that preferences of generation t take the additively separable form

$$U_t = u_t(c_t) + aU_{t-1} + bU_{t+1},$$

where a and b are positive constants. Kimball was the first to solve this system of interdependent utility functions for the equivalent set of utilities defined over allocations. Hori and Kanaye (1989) and Hori (1990) study extensions of the two-sided altruism model to cases where the interaction are of the nonadditively separable form $U_t = V(U_{t-1}, c_t, U_{t+1})$.

Bergstrom (1988) examines Kimball's model of two-sided altruism within the more general class of interdependent utilities that are expressed by the matrix equation $U = u + AU$ where A is a nonnegative matrix. In the overlapping generations model, there are infinitely many future generations. This fact threatens to pose formidable mathematical problems. While many of the fundamental results of finite dimensional linear algebra carry over to denumerable matrices and vectors, there are some nasty surprises. Among these surprises are the fact that matrix multiplication is not, in general, associative and the fact that a matrix may have more than one inverse (for a good exposition of this theory, see Kemeny et al., 1966). Fortunately, it turns out that denumerable productive matrices are much better behaved than denumerable matrices in general and in fact share all of the desirable properties of finite productive matrices (see Bergstrom, 1988).

In Kimball's case, the matrix A has values a everywhere on the first subdiagonal, b everywhere on the first superdiagonal and zeros everywhere else. It turns out that the matrix $I - A$ is dominant diagonal if and only if $a + b < 1$. In this case, the interdependent utility functions can be untangled by matrix inversion to yield simple, but very interesting utility functions defined over allocations of consumption. Kimball and Bergstrom both find that the generation t's utility for an infinite consumption stream over the past and future is given by

$$U_t = \sum_{j=1}^{\infty} \alpha^j u_{t-j}(c_{t-j}) + u_t(c_t) + \sum_{j=1}^{\infty} \beta^j u_{t+j}(c_{t+j}),$$

where α and β are constants, such that $0 < \alpha < 1$, $0 < \beta < 1$ and $\alpha/\beta = a/b$.

In this formulation, a person born in period t cares not only about her own consumption and the consumption of her dependents, but also about the consumption of her ancestors. While it may be true that she can do nothing to change the consumption of her ancestors, it could be that her preferences about her own consumption and that of her descendants would be shaped by what had happened to her ancestors. As it happens, preferences on allocations that are derived from the two-sided altruism model are additively separable between the consumption of one's ancestors, one's own consumption and that of one's descendants. Hence for this case, one can study preferences over future generations without investigating family history. This observation illustrates the usefulness of disentangling preferences on allocations from preferences on utilities. When one simply looks at the structure of two-sided altruistic preferences over utilities, it is not obvious without the mathematics that preferences on allocations will be additively separable across generations. One might also want to ask whether it is realistic to assume a preference structure that implies additive separability between one's preferences over the consumption pattern of ancestors and the consumptions of one's descendants. For example, in some families it is important not to leave a smaller estate to one's children than has been the norm for previous generations.

The utility function over allocations that is derived from two-sided altruism implies a time-consistency property which is an interesting generalization of the Strotz–Koopmans property. Consider two generations in the same family line, t and t' where $t < t'$. Generation t and t' will have identical preferences about the allocation of consumption among generations that come after t' and about generations that come before t. They will, however, in general disagree about income transfers among generations in the interval between t and t'.

Laitner (1988) studies gift and bequest behavior in a model of two-sided altruism where bequests must be positive and voluntary and where there are random differences in wealth between generations. As Laitner points out, in reasonable models of intergenerational preferences, there will be gifts from parents to children if the parents are much richer than their children and gifts from children to parents if the children are much richer than their parents, and over some (quite possibly large) intermediate range of relative incomes, there will not be gifts in either direction. Thus there is a positive (and possibly high) probability that in any generation, the chain of voluntary gift-giving necessary to sustain neutrality as in the Barro model or the Bernheim–Bagwell model will be broken.

5.4. Pareto optimality of competitive equilibrium in households with utility interdependence: the First Welfare Theorem

It is reasonable to ask what kind of decentralized allocation mechanisms can achieve Pareto efficient allocation in a household. A competitive equilibrium allocation within the household should certainly be included in any roundup of the usual suspects. But if there are benevolent consumers, there is in general no reason to expect that competitive equilibrium is Pareto optimal. If we define competitive equilibrium so as to exclude the possibility of gifts, then even in a two-person family a competitive equilibrium can fail to be Pareto optimal. The reason is simply that with benevolence it may be possible for both donor and recipient to benefit from a gift.[31]

The problem in the previous example could be fixed by extending the notion of competitive equilibrium to allow for voluntary "gifts". This approach is taken in Bergstrom (1971b), who shows that for two-person families, a competitive gift equi-

[31] For example, consider a family with two persons and one good. Utility functions $U_1(x_1, x_2) = x_1^2 x_2$ and $U_2(x_1, x_2) = x_1 x_2^2$, where x_i is the amount of good consumed by person i. Person 1 has an initial endowment of five units of the private good and person 2 has one unit. The initial endowment (5, 1) is a competitive equilibrium, but it is not Pareto optimal, since both persons would prefer the allocation where person 1 gets four units and person 2 gets two units.

librium is Pareto optimal. But for households in which more than one person cares about the consumption of others, a competitive gift equilibrium in which individuals decide independently how much to give each other is not Pareto optimal. The reason is that the well-being of someone who is loved by more than one person becomes a "public good". Purely bilateral gift arrangements will not result in a Pareto optimal allocation. In such an environment, Pareto efficiency requires multilateral coordination among those who are benevolent toward the same individual. Bergstrom (1971a) explores a Lindahl equilibrium in which those who are benevolent toward an individual each pay some share of the cost of that individual's consumption and all agree on the quantities, given their cost shares.

In Becker's Rotten Kid model (1974), competitive equilibrium with gifts leads to a Pareto optimal allocation in the household. However, this optimality is purchased with a very strong assumption. In particular it is assumed that there is one benevolent family member who makes voluntary gifts to each of the other family members, while no other family members choose to make gifts. Since by assumption, the head of the family is making gifts to all other family members, the allocation that results is the household head's favorite allocation among all allocations which cost no more than total family income.[32]

In an overlapping generations model where each generation has a property right to its own labor, the assumption that a current household head is willing to make positive gifts to all future generations is not attractive. But for families in which preferences are characterized by the recursive structure $U_t = U_t(c_t, U_{t+1})$, competitive equilibrium with voluntary inheritance turns out to be Pareto optimal even if some generations choose to leave nothing to their successors. There seems to be neither a statement nor a proof of this proposition in the literature, but for a family with a finite horizon, proving this proposition is a fairly easy exercise in backward induction.[33] One uses the recursive structure of preferences to show that if an allocation is at least as good for all family members and preferred by some family members to a competitive equilibrium, then the total cost of the proposed allocation to the family dynasty exceeds the total cost of the family's competitive allocation. The remainder of the proof mimics the Arrow–Debreu proof of the Pareto optimality of competitive equilibrium.

5.4.1. The efficiency of competitive equilibrium with nonbenevolence

According to the First Welfare Theorem, under very weak assumptions, a competitive equilibrium is Pareto optimal for selfish consumers. It seems plausible that this result

[32] Although in the simplest version of the Rotten Kid theorem, family members other than the head are assumed to be selfish, the optimality of competitive equilibrium would extend to the case where more than one family member is benevolent if it is assumed that the utility of the head depends positively on the overall utility of each family member.

[33] The proof extends to an infinite horizon if there is sufficient "impatience" so that the present value of resources to appear in the distant future converges to zero.

would extend to the case of malevolent (or nonbenevolent) preferences. Parks (1991) demonstrates that this conjecture holds for a broad class but not for all nonbenevolent preferences. Where all family members have preferences of the Bergson form

$$U_i(v_1(x_1), \ldots, v_n(x_n)),$$

Parks defines the n by n matrix $G(v_1, \ldots, v_n)$ to be the Jacobean matrix whose ijth element is $\partial U_i(v_1, \ldots, v_n)/\partial v_j$. He shows that a competitive equilibrium will necessarily be a local Pareto optimum if the matrix G^{-1} is a nonnegative matrix.[34] As Parks observes, in the case of nonbenevolence the off-diagonal elements of G are nonpositive and the diagonal elements are positive. In addition, as this matrix has the dominant diagonal property (McKenzie, 1960), it will be true that G^{-1} is a nonnegative matrix.

The matrix G will fail to be dominant diagonal if malevolence is too intense. In this case, a competitive equilibrium is not necessarily Pareto optimal. Consider for example a pure exchange economy with two consumers and one private good and suppose that free disposal is possible. Each consumer has an initial endowment of two units of the good. Consumer 1 has utility $U_1(x_1, x_2) = x_1 - x_2^2$ and consumer 2 has utility $U_2(x_1, x_2) = x_2 - x_1^2$, where x_i is consumption by consumer i. The no-trade outcome where $x_1 = x_2 = 2$ is a competitive equilibrium. In this case, each consumer has a utility of -2. But this outcome is evidently not Pareto optimal. For example, if $x_1 = x_2 = 1/2$, each consumer will have a utility of 1/4. In this example, the conditions of Parks' theorem fail since the matrix G turns out not to be dominant diagonal when $x_1 = x_2 = 2$.

5.5. *Sustainability of Pareto optimality as competitive equilibria: the Second Welfare Theorem*

Winter (1969) observed that the Second Welfare Theorem (with convex preferences, every Pareto optimum can be sustained as a competitive equilibrium) extends without modification to the case of nonmalevolent preferences. This result has an interesting application to the theory of family consumption because it suggests that in families where nonmalevolence reigns, consumption decisions can be efficiently decentralized by giving each family member an allowance to spend on personal consumption.

Winter's result, however, is not quite as powerful as it might first appear. Competitive equilibrium as defined by Winter requires that each family member spend his income only on himself. A more useful theorem for decentralization in a benevolent

[34] Parks' proof is as follows. A local Pareto improvement is possible only if $G\,dv > 0$ for some vector dv. If G^{-1} is nonnegative, then $G\,dv > 0$ implies $dv > 0$. But as in the proof of the First Welfare Theorem without externalities, it must be that starting from a competitive equilibrium, there is no feasible change in allocation for which $dv > 0$.

family would state that in a "competitive equilibrium with gifts", where people are allowed to choose their best combination of personal consumption and money transfers to others, every Pareto optimum is a competitive equilibrium. But this result is not true without some qualification. For example, consider the case where $U_1(x_1, x_2) = x_1 x_2^2$ and $U_2(x_1, x_2) = x_1^2 x_2$ and consider the allocation (3, 3), which is Pareto optimal. If no gifts are permitted, then this is a competitive equilibrium, but if gifts are allowed, person 1 would want to give one unit to person 2 and accept no gifts from her. If person 2 were allowed to choose, she would give one unit to person 1 and accept nothing from him. Thus there will be no equilibrium in which each is allowed to determine his or her net gift to the other. We could rescue the situation by defining a gift equilibrium to be one in which nobody wants to make a gift which other persons are willing to accept, or alternatively by assuming that persons are "selfish enough" so that it never happens that one person wants to make a gift that the other will not accept. More complicated versions of this problem arise when several generations have interconnected utility functions (see Bergstrom, 1971b; Pearce, 1983).

Archibald and Donaldson (1976) show that with certain restrictions, the Second Welfare Theorem extends to nonpaternalistic preferences which are not monotonic increasing in all of the v_j's. Their argument is based on the observation that the standard first-order conditions for Pareto optimality in an economy with nonpaternalistic preferences require that individuals all have the same marginal rates of substitution between goods. Given sufficient convexity, and given that the constrained optimality problem determining a Pareto optimum satisfies the appropriate constraint qualifications so that the standard first-order conditions are necessary for Pareto optimality, the Archibald–Donaldson conclusion follows.

5.5.1. Public goods and benefit–cost analysis in benevolent families

If family members want each other to be happy and if they share some household public goods, how do we determine a Pareto efficient expenditure on these public goods? For example, consider a married couple without children who are deciding whether to get a new car. The price of a new car is $\$P$. Suppose that the husband is willing to pay $\$H_1$ for the enjoyment he would get from using the car and $\$H_2$ for the enjoyment his wife would get from using the car. The wife is willing to pay $\$W_1$ for the enjoyment she would get from using the car and $\$W_2$ for the enjoyment the husband would get from using the car. How much should the couple be willing to pay in total for the car?

In the presence of "pure" nonmalevolence, there is a very simple and perhaps surprising answer to this question. Even though each person is willing to pay something for the other's enjoyment of the car, they should buy the car if and only if the sum $H_1 + W_1 \geq P$. This result is an instance of a very general result that also applies to multiperson families and to cases where the public goods are supplied continuously rather than discretely.

Consider a family with n members where the utility of each household member i can be expressed as $U_i(u_1(c_1, y), \ldots, u_n(c_n, y))$ where c_i is the vector of private consumption goods consumed by i and where y is the vector of household public goods consumed by the family. The assumption of nonmalevolence means that U_i is an increasing function of u_i and a nondecreasing function of u_j for all $j \neq i$. Therefore if an allocation is Pareto optimal (in terms of the U_i's), it must be that this allocation would also be Pareto optimal for an economy of selfish people in which each i has a utility function $u_i(x_i, y)$. But this means that any conditions which are necessary conditions for optimality in this selfish family are also necessary conditions for optimality in the actual benevolently related family.

For our example of the husband, wife, and car, it is easy to see that if $P < H_1 + W_1$, they can achieve a Pareto improvement by buying the car and dividing the costs so that the husband gives up less than H_1 dollars worth of private goods and the wife gives up less than W_1 dollars worth of private goods. Now suppose that $P > H_1 + W_1$. Imagine for the moment that husband and wife are selfish with private utilities $u_i(c_i, y)$. Then buying the car would be inefficient in the following sense. For any household allocation that they could afford if they buy the car, there will be another household allocation in which they do not buy the car and both of their private utilities will be higher. But since their preferences are benevolent, the fact that they can improve both of their private utilities by not buying the car implies that they can both increase the utilities that represent their benevolent preferences by not buying the car. If the couple were to use a decision rule such as "Buy the car if $P < H_1 + W_1 + H_2 + W_2$" they would act inefficiently whenever $P < H_1 + W_1 + H_2 + W_2$ but $H_1 + W_1 < P$.

Where the quantity of public goods is a continuous variable and consumers are selfish, the fundamental benefit–cost result for efficient supply of public goods in an economy is the Samuelson first-order condition (Samuelson, 1954) which requires that the sum of all individuals' marginal rates of substitutions between the public good and their own private consumption equals the marginal cost of public goods in terms of private goods. Since the Samuelson condition is a necessary condition for Pareto optimality in the selfish family where individual preferences are $u_i(c_i, y)$, and since Pareto efficiency in this selfish family is necessary for Pareto efficiency in the corresponding benevolent family, the Samuelson conditions measured from the selfish utility functions must be satisfied in order for there to be efficiency in the benevolent family.

Although the problem of benefit–cost analysis of household public goods in benevolent families seems interesting and important, it does not seem to have received much attention in the literature. The issue does, however, arise fairly frequently in discussions in the public policy literature about how to value persons' lives. If family members love each other, then the survival of each is a household public good. Jones-Lee (1991, 1992) has recently organized and clarified this discussion. According to Jones-Lee, the traditional prescription for evaluating a public project that saves "statistical lives" is that the evaluation should include not only people's willingness to pay for their own safety, but the sum of the amounts people would be willing to pay for improvements in the safety of others.[35] As Jones-Lee points out, Bergstrom (1982) claims this prescription is inappropriate if altruism takes the form of pure concern for other people's utility. Bergstrom's (1982) argument is essentially the same as the argument made above for household public goods, but was specialized to the analysis of risks to life.

Jones-Lee discusses the alternative case of "safety-focused altruism" in which people's only concern with the well-being of others is with their survival probabilities. In this case, he shows that it is appropriate in benefit–cost analysis to add people's willingness to pay for other people's survival probabilities to their willingness to pay for their own. Jones-Lee (1992) also suggests a model of interdependent preferences, which he calls paternalistic preferences, in which each person is "benevolent" towards others, but instead of accepting the other person's relative valuation of survival probability and wealth, the paternalistic individual wishes to impose his *own* relative values on the recipient.

5.6. Evolutionary models of benevolence with the family

In recent years, evolutionary biologists have developed a body of formal theory of the amount of altruism that can be expected to emerge among relatives in sexually reproducing species. Haldane (1955) remarked that according to evolutionary theory, one should be prepared to rescue a sibling from drowning if the likelihood of saving the sibling's life is at least twice the risk to one's own. To induce one to take the same risk for a first cousin, the likelihood of saving the cousin's life must be at least eight times the risk of drowning oneself.

Hamilton's remarkable papers (1964a,b) were the first to work out a formal justification for Haldane's calculus of altruism. Hamilton's main result has come to be known as "Hamilton's Rule". Hamilton states his rule as follows:

> The social behavior of a species evolves in such a way that in each distinct behavior-evoking situation the individual will seem to value his neighbors' fitness against his own according to the coefficients of relationship appropriate to that situation. (1964b: p. 19)

According to Hamilton's rule, natural selection will favor genes that lead a creature to be willing to exchange its own expected number of offspring for those of a relative so long as $c/b < r$, where c is the cost of the action in terms of ones own expected off-

[35] This prescription is advanced in Mishan's (1971) classic paper on the evaluation of human life and safety as well as in papers by Needleman (1976), Jones-Lee (1976), and Viscusi et al. (1988).

spring, b is the gain to the relative and r is the "coefficient of relatedness" between the individual and his relative. For diploid, sexually reproducing species with random mating, r is 1/2 for offspring and full siblings, 1/4 for grandchildren and half-siblings, 1/8 for great grandchildren and first cousins, and so on.

Dawkins' book *The Selfish Gene* (1976) popularized Hamilton's theory in a way that many economists have found accessible and stimulating. Dawkins advocates the viewpoint that the replicating agent in evolution is the *gene* rather than the animal. If a gene carried by one animal is likely to appear in its relatives, then a gene for helping one's relatives, at least when it is cheap to do so, will prosper relative to genes for totally selfish behavior. Trivers' book *Social Evolution* (1985) explores numerous applications of the theory of the evolution of altruism and conflict between relatives. This book is a pleasure to read, with a fascinating mixture of theories and applications of the theories throughout the animal kingdom.

Hamilton's rule is intriguing because it not only predicts a limited degree of altruism toward relatives, but makes explicit predictions of the degree of altruism as a function of the degree of relationship. Since the environments that shaped our genes are hidden in the distant past, most economists are skeptical about the usefulness of evolutionary hypotheses for explaining human preferences. Still, such fundamental features of family life as mating, child-rearing, and sibling relations are remarkably similar across existing cultures[36] and are likely not to have changed drastically over the millennia. This suggests that evolutionary theory can be expected to enrich the economics of the family.[37]

Trivers (1985) applies the Hamilton theory to parent–offspring conflict and to sibling rivalry and sibling conflict. According to Hamilton's theory, in a sexually reproducing diploid species, full siblings (who on average have half of their genes in common) will tend to value each other's survival probability half as much as they value their own. Parents, on the other hand, will value the survival probabilities of each offspring equally. Trivers illustrates these theoretical problems with field observations of feeding conflicts between mother and offspring and among siblings in several species of birds and mammals.

Hamilton proves his propositions only for environments where costs and benefits are purely additive. That is, each individual's survival probability can be expressed as a sum of "gifts" given to or received from relatives. Bergstrom (1995) extends the Hamilton model of altruism between siblings in order to allow general interactions in which benefits and costs from helping others may be nonlinear and nonseparably interactive.

Bergstrom and Stark (1993) offer a series of models in which altruism between

[36] For an anthropologist's view of the near-universality of much family structure, see for example, Stephens (1963).

[37] This view seems to be shared by Becker (1976) and Hirshleifer (1977, 1978). Hirshleifer's 1978 paper contains an engaging manifesto on behalf of an evolutionary theory of preference formation.

siblings and neighbors persists under evolutionary pressures. In these models, inheritance may be either genetic or "cultural".

5.7. *Conscious choice of altruism*

5.7.1. *When is more love not a good thing?*

In evolutionary models, the degree of altruism is selected endogenously by forces of natural selection. Bernheim and Stark (1988) consider some issues that arise if people are able to make conscious choices about how much to love others. Choices of this kind are especially pertinent to courtship and marriage. The metaphor "falling in love" suggests a certain lack of control of the process, but even here, one has some choice in choosing which precipices to approach.

Bernheim and Stark find interesting examples in which an increase in love by one individual may be bad for both the lover and the beloved. They parameterize love as a particular kind of interdependence of utility functions and show that altering the amount of love in a relationship can have surprising effects. For example, suppose that one member of a couple is naturally unhappy. If his partner were to increase her love for him, she would share his unhappiness and become visibly less happy herself. To make matters worse, her unhappy beloved would become even more miserable when he observes her reduced happiness.

Bernheim and Stark find further paradoxes in the application of noncooperative game theory to people who love each other. Consider two players in a nonzero-sum game. Suppose that an increase in one person's love is defined as making that person's payoff a convex combination of his own and his partner's payoff with an increased weight on the partner's payoff. In the Nash equilibrium for the resulting game, an increase in love may turn out to decrease rather than increase the payoffs of one or both partners. Bernheim and Stark also show that increased love may make both parties worse off in multi-stage games, where an increase in love may eliminate certain punishment strategies as credible threats and hence result in a Pareto inferior equilibrium.

5.7.2. *Maximizers and imitators*

Cox and Stark (1992) suggest that selfish people may choose to be kind to their aged parents because with some probability this behavior will be "imprinted" on their own children, who when the time comes will treat their own parents as they saw their parents treat their grandparents. Parents would then find it in their self-interest to treat their parents as they would like to be treated themselves when they are old.

Bergstrom and Stark (1993) remark that there is an incongruity in assuming that each generation rationally selects its behavior towards its parents, but believes that its children will copy their parents rather than make their own rational choices. Bergstrom and Stark suggest a model in which some fraction of children turn out to be imitators, while the others are maximizers. They assume that the environment is stationary, so that maximizers in any generation have the same utility functions and face the same probabilities that their children will be maximizers. A maximizer, although she may be entirely selfish, realizes that if her child is an imitator, then the help that she gives to her aged parents will be rewarded when she, herself, is old, by her child who has learned to treat aged parents generously. If, however, the child is a maximizer, the child's best action is independent of the way her mother acted. Bergstrom and Stark describe the optimizing conditions for maximizers in this situation and show that the more likely children are to be imitators, the better people will treat their aged parents. Since the behavior of imitators is ultimately copied from an ancestor who is a maximizer, the same analysis predicts the behavior of imitators.

References

Archibald, G.C. and D. Donaldson (1976), "Non-paternalism and the basic theorems of welfare economics", Canadian Journal of Economics 9: 492–507.

Axelrod, R. (1984), The evolution of cooperation (Basic Books, New York).

Barro, R. (1974), "Are government bonds net wealth?", Journal of Political Economy 82: 1095–1117.

Becker, G. (1965), "A theory of the allocation of time", Economic Journal 75: 493–517.

Becker, G. (1974), "A theory of social interactions", Journal of Political Economy 82: 1063–1094.

Becker, G. (1976), "Altruism, egoism, and fitness: economics and sociobiology", Journal of Economic Literature 14: 817–826.

Becker, G. (1981), A treatise on the family (Harvard University Press, Cambridge, MA).

Bergson, A. (1938), "A reformulation of certain aspects of welfare economics", Quarterly Journal of Economics 52: 310–334.

Bergstrom, T. (1971a), "A 'Scandinavian consensus' solution for efficient allocation with non-malevolent preferences", Journal of Economic Theory 2: 383–398.

Bergstrom, T. (1971b), "Interrelated consumer preference and voluntary exchange", in: A. Zarley, ed., Papers in quantitative economics (Kansas University Press, Lawrence, KS) pp. 79–94.

Bergstrom, T. (1988), "Systems of benevolent utility interdependence", Working paper (University of Michigan, Ann Arbor, MI).

Bergstrom, T. (1989a), "A fresh look at the rotten kid theorem", Journal of Political Economy 97: 1138–1159.

Bergstrom, T. (1989b), "Love and spaghetti, the opportunity cost of virtue", Journal of Economic Perspectives 3: 165–173.

Bergstrom, T. (1995), "On the evolution of altruistic ethical rules for siblings", American Economic Review 85: 58–81.

Bergstrom, T. and M. Bagnoli (1993), "Courtship as a waiting game", Journal of Political Economy 101: 185–202.

Bergstrom, T. and R. Cornes (1981), "Gorman and Musgrave are dual – an antipodean theorem on public goods", Economics Letters 7: 371–378.

Bergstrom, T. and R. Cornes (1983), "Independence of allocative efficiency from distribution in the theory of public goods", Econometrica 51: 1753–1765.

Bergstrom, T. and D. Lam (1989a), "The effects of cohort size on marriage markets in twentieth century Sweden", in: The family, the market, and the state in industrialized countries (Oxford University Press, Oxford).

Bergstrom, T. and D. Lam (1989b), "The two-sex problem and the marriage squeeze in an equilibrium model of marriage markets", Working paper (University of Michigan, Ann Arbor, MI).
Bergstrom, T. and R. Schoeni (1992), "Income prospects and age of marriage", Working paper (Economics Department, University of Michigan, Ann Arbor, MI).
Bergstrom, T. and O. Stark (1993), "How altruism can prevail in an evolutionary environment", American Economic Review (Papers and Proceedings) 83: 149–155.
Bergstrom, T. and H. Varian (1985), "When do market games have transferable utility", Journal of Economic Theory 35: 222–233.
Bergstrom, T., L. Blume and H. Varian (1986), "On the private supply of public goods", Journal of Public Economics 29: 25–49.
Bernheim, B.D. and K. Bagwell (1988), "Is everything neutral?", Journal of Political Economy 96: 308–338.
Bernheim, B.D. and O. Stark (1988), "Altruism within the family reconsidered: do nice guys finish last?", American Economic Review 78: 1034–1045.
Binmore, K.G. (1985), "Bargaining and coalitions", in: A. Roth, ed., Game-theoretic models of bargaining (Cambridge University Press, Cambridge) pp. 259–304.
Binmore, K.G. (1987), "Perfect equilibrium in bargaining models", in: K. Binmore and P. Dasgupta, eds., The economics of bargaining (Blackwell, Oxford).
Binmore, K., A. Shaked and J. Sutton (1989), "An outside option experiment", Quarterly Journal of Economics 753–770.
Blackorby, C., D. Primont and R.R. Russell (1978), Duality, separability and functional structure: theory and economic applications (North-Holland, New York).
Brinig, M. (1990), "Rings and promises", Journal of Law, Economics, and Organization 6: 203–215.
Browning, M. (1992), "Children and household economic behavior", Journal of Economic Literature 30: 1434–1475.
Browning, M., F. Bourguignon, P.-A. Chiappori and V. Lechene (1992), "A structural empirical model of within household allocation", Working paper (McMaster University, Hamilton, Ontario).
Bruce, N. and M. Waldman (1990), "The rotten-kid theorem meets the Samaritan's dilemma", Quarterly Journal of Economics 105: 155–165.
Buchanan, J. (1975), "The Samaritan's dilemma", in: E. Phelps, ed., Altruism, morality, and economic theory (Russell Sage Foundation, New York) pp. 71–85.
Campbell, D. and M. Truchon (1988), "Boundary optima and the theory of public goods supply", Journal of Public Economics 35: 241–250.
Carmichael, H.L. and W.B. MacLeod (1993), "Gift giving and the evolution of cooperation", Working paper (Economics Department, Queens University, Kingston, Ontario).
Chiappori, P.-A. (1988), "Rational household labor supply", Econometrica 56: 63–90.
Chiappori, P.-A. (1992), "Collective labor supply and welfare", Journal of Political Economy 100: 437–467.
Chiappori, P.-A., L. Haddad, J. Hoddinot and R. Kanbur (1992), "Unitary versus collective models of the household: time to shift the burden of proof", Paper presented at the 1993 American Economics Meetings.
Chipman, J. (1974), "Homothetic preferences and aggregation", Journal of Economic Theory 8: 26–28.
Coleman, J. (1990), Foundations of social theory (Harvard University Press, Cambridge, MA).
Cox, D. and O. Stark (1992), "Intergenerational transfers and the demonstration effect", Mimeo. (Boston College, Chestnut Hill, MA).
Crawford, V. and A. Kelso (1982), "Job matching, coalition formation, and gross substitutes", Econometrica 50: 1483–1504.
Crawford, V. and E. Knoer (1981), "Job matching with heterogeneous firms and workers", Econometrica 49: 437–451.
Danzig, G. (1951), "Application of the simplex method to the transportation problem", in: T.C. Koopmans, ed., Activity analysis of production and allocation (Wiley, New York).

Das Gupta, P. (1973), "Growth of the U.S. population, 1940–1971, in the light of an interactive two-sex model", Demography 10: 543–565.

DaVonza, J. and M.O. Rahman (1994), "American families: trends and correlates", Population Index 59: 350–386.

Dawkins, R. (1976), The selfish gene (Oxford University Press, New York).

Deaton, A. and J. Muellbauer (1980), Economics and consumer behavior (Cambridge University Press, Cambridge).

Demange, G. and D. Gale (1985), "The strategy structure of two-sided matching models", Econometrica 53: 873–888.

Diamond, P. (1965), "National debt in a neoclassical growth model 2", American Economic Review 55: 1126–1149.

Diamond, P. (1965), "The evaluation of infinite utility streams", Econometrica 33: 170–17.

Diewert, E. (1977), "Generalized Slutsky conditions, for aggregate consumer demand functions", Journal of Economic Theory 15: 333–336.

Eisenberg, E. (1961), "Aggregation of utilities", Management Science 7: 337–350.

Gale, D. (1960), The theory of linear economic models (McGraw-Hill, New York).

Gale, D. and L. Shapley (1962), "College admissions and the stability of marriage", American Mathematical Monthly 69: 9–15.

Gaulin, S. and J. Boster (1992), "Dowry as female competition", American Anthropologist 92: 994–1005.

Goody, J. (1973), "Bridewealth and dowry in Africa and Eurasia", in: J. Goody and S. Tambiah, eds., Bridewealth and dowry (Cambridge University Press, Cambridge) pp. 1–57.

Gorman, W.M. (1953), "Community preference fields", Econometrica 21: 63–80.

Haddad, L. and J. Hoddinott (1991), "Household expenditures, child anthropomorphic status and the intra-household division of income: evidence from the Cote d'Ivoire", Discussion paper 155 (Research Program in Development Studies, Princeton University, Princeton, NJ).

Haddad, L. and J. Hoddinott (1994), "Womens' income and boy–girl anthropometric status in the Cote d'Ivoire", World Development 22: 543–553.

Haldane, J.B.S. (1955), "Population genetics", New Biology 18: 34–51.

Hamilton, W.D. (1964a), "The genetical evolution of social behavior. I", Journal of Theoretical Biology 7: 1–16.

Hamilton, W.D. (1964b), "The genetical evolution of social behavior. II", Journal of Theoretical Biology 7: 17–52.

Hartung, J. (1982), "Polygyny and inheritance of wealth", Current Anthropology 23: 1–12.

Hicks, J.R. (1956), A revision of demand theory (Clarendon Press, Oxford).

Hirshleifer, J. (1977), "Economics from a biological viewpoint", Journal of Law and Economics 1–52.

Hirshleifer, J. (1978), "Natural economy versus political economy", Journal of Social and Biological Structures 1: 319–337.

Hoddinott, J. and L. Haddad (1995), "Does female income share influence household expenditure? Evidence from Cote d'Ivoire", Oxford Bulletin of Economics and Statistics 57: 77–96.

Hori, H. (1990), "Utility functionals with nonpaternalistic intergenerational altruism: the case where altruism extends to many generations", Working paper.

Hori, H. and S. Kanayo (1989), "Utility functionals with nonpaternalistic intergenerational altruism", Journal of Economic Theory 49: 241–255.

Johnson, D.S. (1990), "Team behavior in the family: an analysis of the rotten kid theorem", Working paper (Bureau of Labor Statistics).

Jones-Lee, M. (1991), "Altruism and the value of other people's safety", Journal of Risk and Uncertainty 4: 213–219.

Jones-Lee, M. (1992), "Paternalistic altruism and the value of statistical life", The Economic Journal 102: 80–90, 102.

Juster, F.T. and F.P. Stafford (1991), "The allocation of time: empirical findings, behavioral models, and problems of measurement", Journal of Economic Literature 29: 471–522.
Kaneko, M. (1982), "The central assignment game and the assignment markets", Journal of Mathematical Economics 10: 205–232.
Kemeny, J., J. Snell and A. Knapp (1966), Denumerable Markov chains (Van Nostrand, Princeton, NJ).
Keyfitz, N. (1971), "The mathematics of sex and marriage", in: Proceedings of the sixth Berkeley symposium on mathematical statistics and probability, Vol. IV (University of California Press, Berkeley, CA).
Kimball, M. (1987), "Making sense of two-sided altruism", Journal of Monetary Economics 20: 301–326.
Koopmans, T. (1960), "Stationary ordinal utility and impatience", Econometrica 28: 287–309.
Koopmans, T. and M. Beckman (1957), "Assignment problems and the location of economic activities", Econometrica 25: 53–76.
Laitner, J. (1979a), "Household bequests, perfect expectations, and the national distribution of wealth", Econometrica 47: 1175–1194.
Laitner, J. (1979b), "Bequests, golden-age capital accumulation and government debt", Economica 46: 403–414.
Laitner, J. (1988), "Bequests, gifts, and social security", Review of Economic Studies 55: 275–299.
Laitner, J. (1990), "Intergenerational preference differences and optimal national saving", Journal of Economic Theory 22: 56–66.
Laitner, J. (1991), "Modeling marital connections among family lines", Journal of Political Economy 99: 1123–1141.
Lam, D. (1988), "Marriage markets and assortative mating with household public goods", Journal of Human Resources 23: 462–487.
Lindbeck, A. and J. Weibull (1988), "Altruism and efficiency, the economics of fait accompli", Journal of Political Economy 96: 1165–1182.
Lotka, A. (1922), "The stability of the normal age distribution", Proceedings of the National Academy of Sciences USA 8: 339.
Lundberg, S. and R. Pollak (1993), "Separate spheres bargaining and the marriage market", Journal of Political Economy 101: 988–1011.
Lundberg, S. and R. Pollak (1994), "Noncooperative bargaining models of marriage", American Economic Review 84: 132–137.
Manser, M. and M. Brown (1980), "Marriage and household decision theory – a bargaining analysis", International Economic Review 21: 21–34.
Marshack, J. and R. Radner (1972), Economic theory of teams (Yale University Press, New Haven, CT).
Matzkin, R. and P. Streufert (1991), "On interdependent utility", Working paper (University of Wisconsin, Madison, WI).
McElroy, M. (1990), "The empirical content of Nash-bargained household behavior", Journal of Human Resources 25: 559–583.
McElroy, M. and M. Horney (1981), "Nash-bargained decisions: toward a generalization of the theory of demand", International Economic Review 22: 333–349.
McFarland, D. (1972), "Comparison of alternative marriage models", in: T.E. Greville, ed., Population dynamics (Academic Press, New York) pp. 89–106.
McKenzie, L. (1960), "Matrices with dominant diagonals and economic theory", in: K. Arrow, K.S. Karlin and P. Suppes, eds., Mathematical methods in the social sciences (Stanford University Press, Stanford, CA).
Mishan, E.J. (1971), "Evaluation of life and limb: a theoretical approach", Journal of Political Economy 79: 687–705.
Moldevanu, B. (1990), "Stable bargained equilibria for assignment games without side-payments", International Journal of Game Theory 19: 171–190.
Murdock, G.P. (1967), Ethnographic atlas (University of Pittsburgh Press, Pittsburgh, PA).
Muth, R. (1966), "Household production and consumer demand functions", Econometrica 34: 699–708.

Nash, J. (1950), "The bargaining problem", Econometrica 18: 155–162.

Needleman, L. (1976), "Valuing other people's lives", Manchester School 44: 309–342.

Parks, R. (1991), "Pareto irrelevant externalities", Journal of Economic Theory 54: 165–179.

Pearce, D. (1983), "Nonpaternalistic sympathy and the inefficiency of consistent intertemporal plans", Working paper (Cowles Foundation, Yale University, New Haven, CT).

Phelps, E. and R. Pollak (1968), On second-best national saving and game-equilibrium growth, Review of Economic Studies 35: 185–200.

Pollak, R. (1968), "Consistent planning", Review of Economic Studies 35: 201–208.

Pollak, R. (1976), "Interdependent preferences", American Economic Review 66: 309–320.

Pollak, R. (1986), "A reformulation of the two-sex problem", Demography 23: 247–259.

Pollak, R. (1987), "The two-sex problem with persistent unions: a generalization of the birth matrix-mating rule model", Theoretical Population Biology 32: 176–187.

Pollak, R. (1990), "Two-sex demographic models", Journal of Political Economy 98: 399–420.

Pollak, R. and M. Wachter (1975), The relevance of the household production function and its implications for the allocation of time", Journal of Political Economy 68: 349–359.

Rader, T. (1964), "Edgeworth exchange and general economic equilibrium", Yale Economic Essays 4: 133–180.

Rochford, S. (1984), "Symmetrically pairwise-bargained allocations in an assignment market", Journal of Economic Theory 34: 262–281.

Roth, D. (1992), "A theory of partnership dynamics, learning, specific investment, and dissolution", Working paper (University of Michigan, Ann Arbor, MI).

Roth, A. and M. Sotomayor (1990), Two-sided matching (Cambridge University Press, Cambridge).

Rubinstein, A. (1982), "Perfect equilibrium in a bargaining model", Econometrica 50: 97–109.

Rubinstein, A. and A. Wolinsky (1985), "Equilibrium in a market with sequential bargaining", Econometrica 53: 1133–1150.

Samuelson, P. (1947), Foundations of economic analysis (Harvard University Press, Cambridge, MA).

Samuelson, P. (1950), "Evaluation of real national income", Oxford Economic Papers II: 1–20.

Samuelson, P. (1954), "The pure theory of public expenditures", Review of Economics and Statistics 36: 387–389.

Samuelson, P. (1956), "Social indifference curves", Quarterly Journal of Economics 70: 1–22.

Samuelson, P. (1958), "An exact consumption-loan model of interest with or without the social contrivance of money", Journal of Political Economy 66: 467–482.

Samuelson, P. (1961), "A new theorem on nonsubstitution", in: Money, growth and methodology, Published in honor of Johan Akerman (Lund Social Science Studies, Lund, Sweden) pp. 407–423.

Schultz, T.P. (1990), "Testing the neoclassical model of family labor supply and fertility", Journal of Human Resources 25: 599–634.

Shafer, W. (1977), "Revealed preference and aggregation", Econometrica 45: 1173–1182.

Shaked, A. and J. Sutton (1984), "Involuntary unemployment as a perfect equilibrium in a bargaining model", Econometrica 52: 1351–1364.

Shapiro, P. (1977), "Social utility functions", Journal of Economic Theory 16: 475–480.

Shapley, L. and M. Shubik (1972), "The assignment game I: the core, International Journal of Game Theory 1: 111–130.

Stephens, W.N. (1963), The family in cross-cultural perspective (Holt, Rinehart, Winston, New York).

Strotz, R. (1955), "Myopia and inconsistency in dynamic utility maximization", Review of Economic Studies 23: 165–180.

Sutton, J. (1986), "Noncooperative bargaining theory: an introduction", Review of Economic Studies 53: 709–724.

Thomas, D. (1990), "Intra-household resource allocation: an inferential approach", Journal of Human Resources 25: 635–696.

Trivers, R. (1985), Social evolution (Benjamin/Cummings, Menlo Park, CA).

Varian, H. (1983), "Nonparametric tests of consumer behavior", Review of Economic Studies 50: 99–110.
Varian, H. (1984), "Social indifference curves and aggregate demand", Quarterly Journal of Economics 99: 403–414.
Viscusi, W.K., W.A. Magat and A. Forest (1988), "Altruistic and private valuations of risk reduction", Journal of Policy Analysis and Management 7: 227–245.
Weiss, Y. and R. Willis (1985), "Children as collective goods in divorce settlements", Journal of Labor Economics 3: 268–292.
Weiss, Y. and R. Willis (1993), "Transfers among divorced couples: evidence and interpretation", Journal of Labor Economics 11: 629–679.
Willis, R. (1994), "A theory of out-of-wedlock childbearing", Working paper (University of Chicago, Chicago, IL).
Wilson, W.J. (1987), The truly disadvantaged (University of Chicago Press, Chicago, IL).
Winter, S. (1969), "A simple remark on the second optimality theorem of welfare economics", Journal of Economic Theory 1: 99–103.
Woolley, F. (1988), "A non-cooperative model of family decision making", Working paper TIDI/125 (London School of Economics, London).