

Family Decision-Making

Shelly Lundberg
University of Washington

Robert A. Pollak
Washington University in St. Louis

I. Introduction

Economic models of consumer demand and labor supply begin with an individual economic agent choosing actions that maximize his or her utility function subject to a budget constraint. How can we reconcile this individualistic theory of the consumer with the reality that people tend to live, eat, work, and play in families? Economists have dealt with a possible multiplicity of decision makers in the family in two ways. The first approach, in ascendancy from the 1950s until the 1980s, was the unitary approach--treating the family as though it were a single decision-making agent, with a single pooled budget constraint and a single utility function that includes the consumption and leisure time of every family member. The second approach, pioneered by Manser-Brown and McElroy-Horney in the early 1980s, was to model family demands as the solution to a cooperative bargaining game. Other non-unitary approaches have been developed, including the "collective" model of Chiappori and various non-cooperative models.

Most models of family behavior allow two decision-makers--the husband and the wife. Children are customarily excluded from the set of decision-making agents in the family, though they may be recognized as consumers of goods chosen and provided by loving or dutiful parents. Bargaining models can also be used to analyze interactions between elderly parents and adult children, interactions that typically involve family members living in different households. The empirical implications of non-unitary models of marriage depend upon their assumptions about the form of the game or other joint decision-making process but, in general, these models widen the range of "rational" family behavior.

II. Unitary models

Two models provide the theoretical underpinning of the unitary, or common preference, approach to family behavior: Samuelson's (1956) consensus model and Becker's (1974,1981) altruist model. The consensus model was introduced by Samuelson to exhibit the conditions under which family behavior can be rationalized as the outcome of maximizing a single utility function. Consider a two-member family consisting of a husband and a wife. Each has an individual utility function that depends on his or her private consumption of goods but, by consensus, they agree to maximize a social welfare function of their individual utilities, subject to a joint budget constraint that pools the income received by the two family members. Then we can analyze their

aggregate expenditure pattern as though the family were a single agent maximizing a utility function. That is, the household maximizes $U(c^h, c^w)$, where c^h and c^w are the private consumptions of husband (h) and wife (w), subject to the budget constraint $p(c^h + c^w) = y = y^h + y^w$ which pools the individual incomes of husband and wife. This problem generates demand functions $c^i = f^i(p, y)$ that depend only on prices and total family income and that have standard properties provided the utility functions are well-behaved. Thus, the comparative statics of traditional consumer demand theory apply directly to family behavior under the consensus model. Samuelson did not, however, purport to explain how the family achieves a consensus regarding the joint welfare function, or how this consensus is maintained.

Becker's altruist model (1974,1981) addresses these questions, and also provides an account of how resources are distributed within the family. In Becker's model, the family consists of a group of purely selfish but rational "kids" and one altruistic parent whose utility function reflects his concern for the well-being of other family members. Becker argues that the presence of an altruistic parent who makes positive transfers to each member of the family is sufficient to induce the selfish kids to act in an apparently unselfish way. The altruistic parent will adjust transfers so that each "rotten kid" finds it in his interest to choose actions that maximize family income. The resulting distribution is the one that maximizes the altruist's utility function subject to the family's resource constraint, so the implications of the altruist model for family demands coincide with those of the consensus model.

The unitary framework is a simple, powerful mechanism for generating demand functions and establishing their comparative statics for use in applied problems. In recent years, however, these models have been criticized on both empirical and theoretical grounds. Empirical evidence inconsistent with the implications of the common preference approach has been accumulating. Dissatisfaction with unitary models on theoretical grounds has been the product of serious study, by economists, of marriage and divorce. Models of marriage and divorce require a theoretical framework in which agents compare their expected utilities inside marriage with their expected utilities outside marriage, but the individual utilities of husband and wife cannot be recovered from the social welfare function that generates consumption, labor supply, fertility, and other behavior within marriage. If the analysis of marriage and divorce is awkward, the analysis of marital decisions in the shadow of divorce is even more so. If unilateral divorce is possible, individual rationality implies that marital decisions cannot leave either husband or wife worse off than they would be outside the marriage. This individual rationality requirement, however, alters the comparative statics of the model, and destroys the correspondence between the behavior of a single rational agent and the behavior of a family.

III. Non-unitary models

Cooperative Bargaining Models

A viable alternative to unitary models of the family must recognize, in a non-trivial fashion, the involvement of two or more agents with distinct preferences in determining family consumption. Bargaining models from cooperative game theory satisfy these conditions. A typical cooperative bargaining model of marriage begins with a family that consists of only two members: a husband and a wife. Each has a utility function that depends on his or her consumption of private goods ($U^h(c^h)$ for the husband and $U^w(c^w)$ for the wife). If agreement is not reached, then the payoff received is represented by the “threat point,” $(T^h(Z), T^w(Z))$ --the utilities associated with a default outcome of divorce or, alternatively, a noncooperative equilibrium within the marriage. The threat point depends, in turn, upon a set of exogenous distribution factors Z .

The Nash bargaining model provides the leading solution concept in bargaining models of marriage. The couple maximizes the Nash product function $N = [U^h(c^h) - T^h(Z)][U^w(c^w) - T^w(Z)]$ subject to a pooled budget constraint, and this results in demand functions of the form $c^i = f^i(p, y, Z)$. Thus demands and individual utilities depend upon the distribution factors Z , which may include individual incomes y^h and y^w . This solution can be illustrated by a diagram in utility space, where AB is the utility-possibility frontier. Nash (1950) shows that a set of four axioms, including Pareto optimality--which ensures that the solution lies on the utility-possibility frontier--uniquely characterizes the Nash bargaining solution.




Figure 1 The Nash Bargaining Solution

The utility received by husband or wife in the Nash bargaining solution depends upon the threat point; the higher one's utility at the threat point, the higher one's utility in the Nash bargaining solution. This dependence is the critical empirical implication of Nash bargaining models: family demands depend, not only on prices and total family income, but also on determinants of the threat point.

In divorce-threat bargaining models, the threat point is the maximal level of utility attainable outside the marriage. If divorcing partners maintain ownership of income received separately within marriage, the demands emerging from marital bargaining will depend not on total family income, but on the income received by the husband and the income received by the wife. The divorce threat point is also likely to depend on environmental factors (extrahousehold environmental parameters, or EEP's, in McElroy's (1990) terminology) that do not directly affect marital utility, such as conditions in the remarriage market and the income available to divorced men and women. The family demands that result from divorce-threat marital bargaining will therefore depend upon these parameters as well.

In the separate spheres bargaining model of Lundberg and Pollak (1993), the threat point is internal to the marriage, not external as in divorce-threat bargaining models. The husband and wife settle their differences by Nash bargaining, but the alternative to agreement is an inefficient noncooperative equilibrium within marriage. In a noncooperative equilibrium, each spouse voluntarily provides household public goods, choosing actions that are utility-maximizing, given the actions of their partner. Divorce may be the ultimate threat available to marital partners in disagreement, but a noncooperative marriage in which the spouses receive some benefits due to joint consumption of public goods may be a more plausible threat in day-to-day marital bargaining.

The introduction of this internal threat point has important implications, because separate spheres bargaining generates family demands that, under some circumstances, depend not on who receives income after divorce, but on who receives (or controls) income within the marriage. Lundberg and Pollak assume gender specialization in the noncooperative provision of household public goods, with the husband providing one good out of his own resources, and the wife providing a separate good from her individual resources. These decisions lead to a pair of reaction functions that determine a Cournot-Nash equilibrium in which the public goods contributions are inefficiently low, and depend upon the distribution of individual incomes within the family.

As the divorce-threat and separate spheres models show, cooperative bargaining does not necessarily imply income pooling, i.e. the property that demands depend only on total household income, rather than its separate components. Bargained outcomes depend upon the threat point, and the income controlled by husband and wife will affect family behavior (and the relative well-being of men and women within marriage) if this

control influences the threat point. This dependence implies that public policy (e.g., taxes and transfers) need not be neutral in their effects on distribution within the family. Also, the absence of pooling and the presence of extrahousehold parameters in family demands yield a model that can be tested against the unitary alternative. For example, changes in the welfare payments available to divorced mothers, or in laws defining marital property and regulating its division upon divorce, should affect distribution between men and women in two-parent families through their effect on the threat point.

The "Collective" Approach

Most models of the family either assume or conclude that family behavior is Pareto optimal. Unitary models ensure Pareto optimality by assuming a family social welfare function that is an increasing function of the utilities of all family members: no member can be made better off without making another worse off. Cooperative bargaining models characterize the equilibrium distribution by means of a set of axioms, one of which is Pareto optimality. Pareto optimality is the defining property of the "collective model" of Chiappori (1988, 1992). Rather than applying a particular cooperative or noncooperative bargaining model to the household allocation process, Chiappori assumes only that equilibrium allocations are Pareto optimal. He demonstrates that, given a set of assumptions including weak separability of public goods and the private consumption of each family member, Pareto optimality implies, and is implied by, the existence of a "sharing rule." Under a sharing rule, the family acts as though decisions were made in two stages, with total family income first divided between public goods and the private expenditures of each individual, and then each individual allocating his or her share among private goods. The collective framework thus imposes a set of testable restrictions on the observed demands of the household. In essence, the ratio of the marginal propensities to consume any two goods must be the same for all sources of income, because the independent incomes of husband and wife affect consumption only through the sharing rule. A series of empirical tests have found consumption expenditures in households has been found to be generally consistent in this sense with Pareto optimality (eg. Bourguignon, Browning, Chiappori, and Lechene, 1993; Browning, Bourguignon, Chiappori, and Lechene, 1994).

Noncooperative Bargaining Models

Cooperative game theory motivates the assumption of Pareto optimality by assuming that information is relatively good (or at least not asymmetric) and that the players can make binding, costlessly-enforceable agreements. Since legal institutions do not provide for external enforcement of contracts regarding consumption, labor supply, and allocation within marriage, the binding-agreement assumption is unappealing. Noncooperative game theory, in contrast, does not assume that binding agreements enforce intrahousehold allocations, but focuses instead on self-enforcing equilibria.

Without binding agreements, much of the motivation for assuming Pareto optimality vanishes. It is possible, however, for noncooperative bargaining to yield Pareto optimal outcomes under certain conditions. In general, repeated noncooperative games have multiple equilibria, and Pareto optimal equilibria can often be sustained by the threat of punishment. One of the benefits of modeling distribution within marriage as

a noncooperative game is the opportunity to treat efficiency as endogenous, potentially dependent upon the institutions and social context of marriage in a particular society and upon the characteristics of the marital partners.

The prevalence of destructive or wasteful phenomena such as domestic violence and child abuse as well as the demand for marriage counseling and family therapy, suggests that we consider the possibility that family behavior is sometimes inefficient. Other researchers have pointed to gender segmentation in the management of businesses or agricultural plots in many countries as evidence of an essentially noncooperative, and possibly inefficient, family environment. One piece of evidence is provided by Udry (1995), who finds that the household allocation of resources to male- and female-controlled agricultural plots in Burkina Faso is inefficient.

Intertemporal Models

In dynamic bargaining models with investment, decisions made in one period can alter the relative bargaining power of individual family members in future periods. Several papers have shown that limited commitment in this situation can lead to inefficient allocations of household resources. Lundberg and Pollak [2001] use a two-stage model of a married couple's location decision to show that marital decisions that affect future bargaining power need not be efficient unless the husband and wife can make binding agreements regarding their future actions. Aura [2001] examines a married couple's consumption and savings choices when they are unable to commit to not renegotiate their decisions in the future, and the efficiency implications of different divorce asset division regimes.

*** This needs work ***

IV. Empirical evidence

Recent empirical evidence suggests that the restrictions imposed on demand functions by common preference models are not well-supported. Rejections of the family income pooling assumption have been most influential in weakening economists' attachment to unitary models. The fraction of income received or controlled by one family member should not influence demands, conditional on total family income. A large number of recent empirical studies have rejected pooling, finding that earned and unearned income received by the husband or wife significantly affect demand patterns when total income or expenditure is held constant. Some studies find that children appear to do better when their mothers control a larger fraction of family resources (Thomas, 1990; Hoddinott and Haddad, 1991).

A test of the pooling hypothesis requires a measure of husband's and wife's relative control over resources. Relative earnings would seem to be an attractive candidate for this measure, since labor income is by far the largest component of family income, and earnings data are readily available and reliably measured. Also, the earnings of wives relative to husbands have increased dramatically in the United States and many other countries, and we would like to assess the distributional consequences, if any, of this change. The difficulty with this approach is that earnings are clearly endogenous

with respect to the household's time allocation decisions, so that households with different ratios of wife's earnings to husband's earnings are likely to face different prices and may have different preferences.

One might try to avoid these problems by testing the pooling of unearned income rather than earnings. Unearned income is not contaminated by price effects, but most unearned income sources are not entirely exogenous with respect to past or present household behavior. Furthermore, variations in unearned income over a cross-section are likely to be correlated with other (possibly unobservable) determinants of consumption. For example, property income reflects, to a considerable extent, accumulated savings and is therefore correlated with past labor supply and, if those who worked a lot in the past continue to do so, current labor supply. Public and private transfers may be responsive to household distress due to unemployment or bad health, and may be related to expenditures through the events that prompted them (Schultz, 1990). *Unexpected* transfers such as lottery winnings, unexpected gifts, or unexpected bequests will affect resources controlled by individuals without affecting prices, but are likely to be sporadic and unimportant for most families. *** Resources at marriage ***

The ideal test of the pooling hypothesis would be based on an experiment in which some husbands and some wives were randomly selected to receive an income transfer. A less-than-ideal test could be based on a "natural experiment" in which some husbands or some wives received an exogenous income change. Lundberg, Pollak, and Wales (1995) examine the effects of such a natural experiment--the policy change in the United Kingdom that transferred a substantial child allowance from husbands to wives in the late 1970s. They find strong evidence that a shift towards relatively greater expenditures on women's goods and children's goods coincided with this income redistribution, and interpret this as a rejection of the pooling hypothesis.

*** Rubalcava & Thomas, PROGRESA studies ***