Published in final edited form as:

J Econ Lit. 2017 September ; 55(3): 985–1045. doi:10.1257/jel.20150715.

Static and Intertemporal Household Decisions[†]

Pierre-Andre Chiappori and

Columbia University, Department of Economics

Maurizio Mazzocco

University of California at Los Angeles, Department of Economics

Abstract

We discuss the most popular static and dynamic models of household behavior. Our main objective is to explain which aspects of household decisions different models can account for. Using this insight, we describe testable implications, identification results, and estimation findings obtained in the literature. Particular attention is given to the ability of different models to answer various types of policy questions.

1. Introduction

Consider a social-welfare program like *Oportunidades* in Mexico or *Bolsa Familia* in Brazil, whereby poor households receive a cash transfer under the condition that their children attend school and receive some minimal health services. An important issue relates to the recipient of the transfer. Should the cash be given to the wife? To the husband? To both? Does it make a difference—and if so, in which respect and over which horizon?

Alternatively, consider income taxation for households. In some countries, the relevant fiscal unit is the household, not the individual. As a consequence, married people are always jointly taxed. In other countries, individuals are taxed independently. Finally, there is a third group of countries where households are allowed to choose between independent or joint taxation. How should an economist think about these different options? Are they equivalent? If not, can we predict their differential effect and make policy recommendations based on such an analysis?

Lastly, consider a change in the legislation governing divorce—for instance a reform that changes the distribution of wealth between spouses from separate ownership of assets to equal division, as occurred in England as a result of a landmark decision by the House of Lords in 2000. Such a reform clearly has an effect on couples who divorce and probably on couples who may divorce with a sufficiently high probability, since the perspective of singlehood is perceived in a different way before and after the decision. But what about couples who are unlikely to divorce? Could the change affect their decisions, for instance by modifying the spouses' respective bargaining powers? And could there be long-term

[†]Go to <https://doi.org/10.1257/jel.20150715> to visit the article page and view author disclosure statement(s).

consequences on intra-household allocation of resources and ultimately on household behavior?

Those three types of policies—and many others—share two common features. They cannot be analyzed without referring, explicitly or implicitly, to a specific model of household behavior; and the policy recommendation stemming from such an analysis will significantly depend on the model adopted.

Until recently, the standard approach to modeling household behavior was based on versions of the so-called “unitary” approach, which assumes that a household can be represented by a unique utility function that is independent of prices and income. In a framework of this type, the answers to the set of questions stemming from the cash-transfer policy are straightforward and clear-cut: the identity of the recipient *cannot* make a difference in terms of household behavior. What exclusively matters, as far as household decisions are concerned, is the total amount of resources at the household’s disposal. Whether resources are provided by the husband, the wife, or both is irrelevant. Essentially, incomes from different sources are pooled and only the aggregate amount has an effect. The answers to the questions arising from the tax-reform policy are complex even within the unitary framework, which predicts that the various tax regimes will have different effects on behavior. For instance, if the tax schedule is progressive, a change to joint taxation *de facto* increases the tax rate for one spouse, and possibly for both. Still, there exist potential consequences of the reform that unitary models are not equipped to consider. One example is a potential change in the spouses’ bargaining positions. A divorce reform is even more difficult to analyze in a unitary model, since it does not distinguish or it cannot identify the utilities that characterize *individual* household members: if the household is represented as a black box summarized by a unique utility function, predicting individual reactions to changes in divorce laws is all but impossible.

In the past three decades, economists have developed models that address some of the limitations of the unitary approach as a framework used to answer policy questions. Those models explicitly recognize that household members have their own preferences and therefore may sometimes disagree on the optimal decisions. Using the new models, researchers interested in evaluating cash-transfer programs can account for the recipient of the transfer and establish whether her or his identity has significant effects on individual welfare and decisions. Economists wishing to assess the effect of a tax reform or a modification in the existing divorce legislation can measure potential changes in bargaining positions and the consequent effects on decisions— including long-term aspects such as education choices, human-capital accumulation, and intra-household specialization.

The previous discussion suggests—and the remainder of this survey will argue more precisely—that the choice of a specific model of household behavior is never irrelevant, and almost never innocuous. Over the last thirty years, considerable progress has been accomplished in the development and assessment of household models that can be used to answer relevant policy questions. The aim of this article, which is divided into four parts, is to survey these advances.

In the first part, we review static models of household decisions. We consider two classes of static models: models that belong to the unitary framework and models that explicitly recognize that households are composed of more than one decision maker. With regard to the second class, we survey models that use noncooperative concepts to characterize household decisions, as well as collective models of the household, i.e., models that assume that household decisions are Pareto efficient. In the second part of the survey, we review intertemporal models of household behavior. The discussion focuses on three main dynamic models: the intertemporal unitary model; a model that extends the static collective model to an intertemporal environment in which household members cannot commit to future allocations of resources; and a similar model with commitment.

A crucial requirement for any model, static or dynamic, is empirical tractability. A model is empirically tractable, and therefore helpful in understanding behavior and answering policy questions, only if it fulfills two requirements. First, the model should be testable: it should generate a set of empirically falsifiable restrictions that fully characterize it, in the sense that a given behavior is compatible with the model if and only if these conditions are satisfied. Second, the model should be identifiable: it should be feasible, possibly under some assumptions, to recover the structure of the model—typically individual preferences and the decision process—from the observation of household behavior. In the third part of the survey, we evaluate whether the models considered in the first two parts of this article satisfy the double requirement by reviewing tests, identification, and estimation results that have been derived in the literature.

We conclude the survey by looking back at the three policies we started with and evaluating how the different models we have reviewed can account for their main effects. The discussion emphasizes two main points. To assess most policies that have an effect on individual welfare and decisions, researchers must rely on a particular model. Without a model, it is not possible to evaluate the effect of the tax reform on a spouse's bargaining position or the long-term effects of changes in divorce laws on intra-household specialization and risk sharing. Moreover, the choice of the model is crucial, since different frameworks have different abilities to assess the various effects of the policy under investigation.

2. Static Models of Household Decisions

In this section, we introduce static models that have been used to study household behavior. The section is divided into four parts. In the first subsection, we describe the setting. Next, we introduce the most commonly used static model: the unitary model. The third subsection discusses noncooperative models of household decisions. In the last subsection, we consider static collective models, which are the main alternative to the static unitary model. Throughout the section, we take households as given and do not discuss household formation and dissolution.¹

¹For a detailed discussion of household formation and dissolution, see Browning, Chiappori, and Weiss (2014).

2.1 The Setting

In the survey, we consider the decisions of a two-person household over the consumption of various commodities and the allocation of time to leisure, labor supply, and household production. We restrict our attention to a two-person household to simplify the discussion. All the results can be generalized to a household composed of I individuals. The commodities consumed by the household include private, as well as public, goods. Specifically, N commodities are publicly consumed within the household. We will denote with Q_k the consumption of public good k and with Q the N -vector of those quantities. Similarly, the household consumption of private good h is denoted with q_h and the n -vector containing private consumption with q . The private good h is consumed by both household members, with member i consuming q_h^i and the spouse consuming the remaining quantity $q_h^j = q_h - q_h^i$. The vectors of private goods that members 1 and 2 consume are denoted with q^1 and q^2 respectively. The associated market prices for public and private goods are given by the N -vector P and the n -vector p , respectively.

Members of the household are each endowed with their own preferences over consumption. In particular, a married person has preferences that are separate from those of the spouse. When modeling the preferences of a married individual, it is important to establish his or her degree of altruism toward the spouse. The most general version of individual preferences for a married individual allows for an unrestrictive form of altruism and can be represented using a utility function of the form $U^i(Q, q^1, q^2)$. In this case, spouse i is concerned directly with the spouse's consumption, and not only with the spouse's welfare.² Even this general setting generates strong testable restrictions on household behavior. Two examples are income pooling and "Symmetry plus rank 1" of the Slutsky matrix, both of which will be described later in the survey.

In many cases, however, tractability and especially the need to identify some aspects of the household decision process demand a more restrictive form of altruism. A standard assumption is that preferences are of the *caring* type. To provide a definition of caring preferences, it is helpful to introduce the concept of felicity function. It measures the part of the individual welfare that a married individual derives from his/her own (public and private) consumption. In an environment with altruism, it isolates the egotistic component of welfare from the component that is generated from caring for the spouse. Without altruism, the felicity function corresponds to the standard utility function. We will denote member i 's felicity function with $u^i(Q, q^i)$. Preferences of the caring type can then be written in the following form:

$$U^i(Q, q^1, q^2) = W^i(u^1(Q, q^1), u^2(Q, q^2)), \quad (1)$$

²Throughout the survey, we clearly distinguish between goods that are public and goods that are private. This modeling choice implicitly assumes that public consumption and households' private consumption, but not individual private consumption, are observed in the data. As an alternative, one could consider a setting in which only the sum of households' public and private consumption of each commodity is observed in the data (see for instance Browning and Chiappori 1998). We have opted for the first specification even if the second one is more general, because it makes clear the different ways in which public and private goods enter the individual preferences.

where W^i is an increasing function.

The assumption that preferences are of the caring type incorporates an important moral principle: i is indifferent between bundles (q^j, Q) that j consumes whenever j is indifferent between them. In this sense, caring is different from the paternalistic view implicit in the general altruistic form of individual preferences, where a spouse cares about the partner's choices and not only about her or his welfare. A consequence of assuming that preferences are caring is that direct externalities between members are ruled out, since i 's evaluation of her private consumption q^i does not depend directly on the private goods that j consumes; it only depends through the felicity function $u^i(Q, q^j)$.

A particular but widely used version of caring is *egotistic* preferences, whereby members only care about their own well-being. In this case, individual preferences can be represented by the felicity function $u^i(Q, q^j)$. Note that such egotistic preferences for consumption do not exclude noneconomic aspects, such as love and companionship. A person's utility may be affected by the presence of the spouse, but not by her consumption. Formally, the true preferences are of the form $F^i(u^i(Q, q^j), a)$, where a is a vector describing marital status and the spouse's characteristics. The function F^i and the vector a will typically play a crucial role in the decision to marry or divorce and in the choice of a partner. However, they are irrelevant for the characterization of the preferences of married individuals over consumption bundles.

Leisure can be introduced in a household model as one of the goods consumed by household members. To highlight time allocation decisions, however, in most of the survey we will make explicit the distinction between standard consumption goods and leisure. We will denote with ℓ^i person i 's leisure, with h^i her (market) labor supply, with w^i her hourly wage, with y^i her nonlabor income, and with T^i the total time available to her. The above discussion about altruism applies also to leisure. It can be assumed that leisure is only privately consumed or that it also enters the spouse's utility function.³

We will allow some commodities to be produced within the household. Following Becker's (1965) seminal contribution, we will assume that a subset of the commodities are the output of a household production function that has two types of inputs: good purchased in the market and hours spent on household production by each of the members. We will denote with X_k the vector of market goods used in the production of public good Q_k , with x_h the corresponding vector employed in the production of private good q_h , with $D_k = (D_k^1, D_k^2)$ the vectors of hours devoted by members 1 and 2 to the production of public good Q_k , and with $d_h = (d_h^1, d_h^2)$ the corresponding vectors spent in the production of private good q_h . The entire set of goods purchased in the market and hours devoted to household production will be summarized using the notation $(X, x) = (X_1, \dots, X_N, x_1, \dots, x_n)$ and $(D, d) = (D_1, \dots, D_N, d_1, \dots, d_n)$, respectively. Therefore, when the public good Q_k and the private good q_h are produced within the household, we have that

³As for private consumption, leisure may affect the spouse's utility through the felicity function or directly. For a detailed investigation of this topic, see for instance Fong and Zhang (2011).

$$Q_k = F_k(X_k, D_k)$$

and

$$q_h = q_h^1 + q_h^2 = f_h(x_h, d_h).$$

The goods purchased in the market, as well as the time allocated by the members to household production, affect the constraints faced by the household. The vector (X, x) enters the budget constraint as part of household expenditure and the vector (D, d) enters the time constraint that each person faces. Notice also that if some commodity a is bought in the market and directly consumed, we can include it in a framework with household production by setting $Q_a = x_a$ if the commodity is public and $q_a = x_a$ if the good is private.

The budget constraint with household production can therefore be written in the following general form:

$$p' \left(\sum_{k=1}^N X_k + \sum_{h=1}^n x_h \right) + \sum_{i=1}^2 w^i \left(l^i + \sum_{k=1}^N D_k^i + \sum_{h=1}^n d_h^i \right) = \sum_{i=1}^2 (y^i + w^i T^i) = Y, \quad (\text{BC})$$

where $Y = \sum_i (y^i + w^i T^i)$ is the household's total potential income and we have substituted out member i 's labor supply h^i using the time constraint $T^i = h^i + \sum_{k=1}^N D_k^i + \sum_{h=1}^n d_h^i$.

Finally, an important concept in models of household decisions, which will be frequently used in the survey, is the notion of *distribution factor*. A distribution factor z_k is any variable that (i) does not affect preferences or the budget constraint, but (ii) may influence the decision process by affecting the decision power of household members.⁴ The survey will make clear how the intra-household decision power can be modeled and how distribution factors can modify it.

For ease of exposition and to maintain consistency across sections, in the rest of the survey we will consider almost exclusively a setting with egotistical preferences and household production in which standard consumption goods are clearly separated from leisure.

2.2 The Static Unitary Model

Historically, the most commonly used model of household behavior has been the static unitary model. The main assumption implicit in this approach is that households behave as single decision makers, independently of the number of household members. This assumption is equivalent to postulating that the household's preferences can be represented using a unique utility function that does not depend on prices, income, and distribution factors, irrespective of the number of individuals that compose the household. This is a

⁴For a detailed discussion of distribution factors, see Browning, Chiappori, and Weiss (2014) and Bourguignon, Browning, and Chiappori (2009).

natural starting point for modeling household behavior, since it makes the model tractable, simple to test, and easy to estimate. Whether the unitary model is a good description of household behavior is, however, a different question altogether.

The most standard version of the unitary model assumes that there is a unique utility function $U^H(Q, q, l^1, l^2)$ that characterizes household's preferences. This function does not depend on individual private consumption, but on the household's aggregate private consumption. A household's behavior can then be described as the solution of the following problem:

$$\max_{(X, x, l^1, l^2, d^1, d^2, D^1, D^2)} U^H(Q, q, l^1, l^2) \quad (\text{UM})$$

subject to

$$p' \left(\sum_{k=1}^N X_k + \sum_{h=1}^n x_h \right) + \sum_{i=1}^2 w^i \left(l^i + \sum_{k=1}^N D_k^i + \sum_{h=1}^n d_h^i \right) = Y, \quad (\text{BC})$$

$$Q_k = F_k(X_k, D_k) \quad (\text{HP})$$

and

$$\sum_i q_h^i = f_h(x_h, d_h) \text{ for all } k \text{ and } h.$$

An immediate argument against the unitary model is that it violates a fundamental requirement of microeconomics, namely methodological individualism. Individuals, not groups, have preferences; the notion of “household preferences” is certainly acceptable, but only insofar as it can be derived from a model that explicitly includes individual preferences and some decision process, what Alderman et al. (1995) called “shifting the burden of proof.” The literature has devised methods that can be used to reconcile the unitary approach with the individualism principle by deriving the “household utility” from a well-specified model of individual preferences. These are reviewed in the next subsection.

2.2.1 Justifying the Unitary Approach: Samuelson's Index, Becker's Rotten Kid, and Transferable Utility—There are three main alternative ways of reconciling the unitary model with an environment in which household members are endowed with their own utility functions: Samuelson's welfare index, Becker's rotten kid theorem, and transferable utility.

Samuelson's welfare index: The first approach assumes that the household's utility is characterized by a welfare index over the individual utility functions of the form $W(u^1(Q,$

$q^1, l^1, u^2(Q, q^2, l^2)$). Then, according to the unitary approach, the household's behavior can be described as the solution of the following problem:

$$\max_{(X, x, l^1, l^2, d^1, D^1, d^2, D^2)} \quad (\text{UMW})$$

$$W\left(u^1(Q, q^1, l^1), u^2(Q, q^2, l^2)\right)$$

subject to

$$p' \left(\sum_{k=1}^N X_k + \sum_{h=1}^n x_h \right) + \sum_{i=1}^2 w^i \left(l^i + \sum_{k=1}^N D_k^i + \sum_{h=1}^n d_h^i \right) = Y, \quad (\text{BC})$$

$$Q_k = F_k(X_k, D_k) \quad (\text{HP})$$

and

$$\sum_i q_h^i = f_h(x_h, d_h)$$

for all k and h .

It is important to note, however, that the two versions (UMW) and (UM) are not empirically distinguishable from each other. To get some insight for why that is the case, one may first remark that the individual utility functions $u^1(Q, q^1, l^1)$ and $u^2(Q, q^2, l^2)$ cannot be separately recovered from the welfare index, W . Indeed, define the function $\bar{U}^H$ by:

$$\bar{U}^H(Q, q^1, l^1, q^2, l^2) = W\left(u^1(Q, q^1, l^1), u^2(Q, q^2, l^2)\right). \quad (2)$$

Data on consumption, individual labor supplies, and prices allow us to identify $\bar{U}^H$, at best. However, for any given $\bar{U}^H$ there exists a continuum of different functions W , u^1 , and u^2 such that (2) is satisfied. As a consequence, u^1 , u^2 , and W cannot be separately identified. Intuitively, variations in the wage of member 1, and therefore in his leisure, provides only information on the value that member assigns to leisure interacted with the value that the household assigns to member 1's preferences for leisure. The two components of household preferences cannot be separated. The same argument applies to variation in the prices of private and public consumption goods.

A second remark is that even the function $\bar{U}^H$ cannot be generally identified from data. This is a consequence of Hicks's composite-good theorem.⁵ Since the prices paid by member 1 for private consumption are generally equal to the prices paid by member 2, one can define a household utility function U^H that depends on household private consumption $q = q^1 + q^2$ (and not on q^1 and q^2) by:

$$U^H(Q, q, l^1, l^2) = \max_{q^1, q^2} \bar{U}^H(Q, q^1, l^1, q^2, l^2)$$

subject to

$$q^1 + q^2 = q.$$

Without independent variation in the prices paid by members 1 and 2 for private consumption, only $U^H(Q, q, l^1, l^2)$ can be identified. And again, for each U^H there is a continuum of $\bar{U}^H$ that are consistent with it. As a result, $\bar{U}^H(Q, q^1, l^1, q^2, l^2)$ cannot be recovered.

All this implies that there is no gain from using the alternative formulation of the unitary model with individual utility functions (UMW) over the standard formulation (UM).⁶ This is an important point, which stresses the intrinsic limits faced by a unitary representation when considering issues related to intra-household inequality or resource allocation.

Becker's rotten kid theorem: An alternative way of reconciling the unitary model with an individualistic environment has been proposed by Becker (1974) in his famous rotten kid theorem. Becker starts with a criticism of Samuelson's approach for simply postulating the index W , instead of deriving it from a more structural model of household behavior. He then proceeds to provide an explicit model of household decision processes from which the unitary setting can be derived.

The main argument underlying the rotten kid theorem can be described as follows. Consider a household with n children and an altruistic parent. The children have preferences over the consumption c of one commodity and a vector a of actions taken by them. These preferences can be represented using the utility function $U^i(c^i, a)$, with $i = 1, \dots, n$. The altruistic parent has preferences over the welfare of the n children, which can be characterized using the utility function $W(u^1(c^1, a), \dots, U^n(c^n, a))$. The actions taken by the children affect the income available to the household: $Y = Y(a)$. Two examples of actions are children's labor supply and the children's contribution to the purchase of a public good.

The children and the parent play a two-stage game. In the first stage, each child independently chooses the action that maximizes his or her own welfare. In the second stage,

⁵For a discussion of Hicks's composite-good theorem, see chapter 5 in Deaton and Muellbauer (1980).

⁶A delicate issue is related to the separability property implied by (2), which could in principle help identification. However, the unitary model is one of the rare cases in which identification does not obtain even with separability, as shown in Chiappori and Ekeland (2009).

given the children's actions, and hence the amount of income available to the household, the parent chooses the allocation of income among the n children that maximizes her or his own utility. Formally, in the first stage, child i chooses the action a_i that maximizes own utility subject to the constraint that own consumption must equal the transfer the child will receive from the parent, i.e., the child solves

$$\max_{a_i} U^i(c_i, a)$$

subject to

$$c_i = t_i(a_1, \dots, a_n).$$

In the second stage, the parent decides the transfers $t_1, \dots, t_n$ that maximize own utility:

$$\max_{t_1, \dots, t_n} W(U^1(t_1, a), \dots, U^n(t_n, a)) \quad (3)$$

subject to

$$\sum_{i=1}^n t_i = Y(a_1, \dots, a_n).$$

The question, now, is whether the outcome of this game is optimal, in the sense that the actions chosen by the kids coincide with what the parent would have chosen if she could freely decide the kids' actions. Technically: does the solution to the two-stage game also solve the parent's program

$$\max_{\{t_i, a_i\}_{i=1, \dots, n}} W(U^1(t_1, a), \dots, U^n(t_n, a)) \quad (4)$$

subject to

$$\sum_{i=1}^n t_i = Y(a_1, \dots, a_n)?$$

A key remark is that *commitment* issues play a central role in that formulation. If the parent could commit over transfers, the solution would be straightforward: she would simply announce that transfers will be nil unless the kids exactly choose the actions she wants to implement. However, Becker implicitly recognizes that such a commitment would not be credible. Once the children have chosen some possibly suboptimal action, the parent will choose the transfers that maximize her own utility taking the children's action as given, which is exactly what program (3) states. Technically, we have a Stakelberg equilibrium with

children as first movers. It is precisely this problem—inducing the right actions in the absence of commitment mechanisms—that the rotten kid theorem aims at solving.

Becker claims that the problem is solved, in the sense that the solution to program (4) is an equilibrium outcome of the two-stage game, under three conditions: (i) there is only one consumption good, (ii) each child's welfare is a normal good in the parent's utility, and (iii) the parent makes a positive transfer to all children. These conditions (and the corresponding conclusion) are commonly known as the rotten kid theorem. Becker's argument goes as follows. Consider the two-stage game. If at the solution the parent makes transfers to all children, the children's consumption is decided by the parent. If each child's welfare is a normal good in the parent's utility, it will increase with the income available to the parent. The children will therefore have the incentives to choose the actions that maximize family income and the two problems have the same solution.

The rotten kid theorem was never formally proved until Bergstrom (1989) derived necessary and sufficient conditions for the result to hold. Bergstrom shows that a necessary and sufficient condition for the rotten kid theorem to be satisfied is that the children's utilities are transferable conditional on the children's actions—i.e., that they take the following form:

$$u^i(c^i, a) = A(a)c^i + B_i(a).$$

A crucial implication of Bergstrom's result is thus that the conditions proposed by Becker are not sufficient for the rotten kid theorem to hold. Specifically, the theorem does *not* hold unless the children's utilities are linear (or affine) in consumption, and the coefficient $A(a)$ is, moreover, the same for all kids. The second restriction is quite strong: since a denotes the vector of actions taken by *all* kids, it must be the case that the action taken by kid i enters the utility of kid $j \neq i$ in exactly the same way as it enters i 's utility. For instance, if the children can choose leisure and their utilities only depend on their own leisure and consumption (but not on the siblings' leisure), the rotten kid theorem fails because the children will typically devote too much time to leisure—even though this reduces the amount of income received by the parent, therefore the sum transferred to them.

Transferable utility: A last and practically important approach to reconcile the unitary model with the existence of individual preferences relies on the transferable utility (TU) assumption. We say that preferences satisfy the TU property if there exists a cardinal representation $u^i(Q, q^i)$ of i 's preferences, $i = 1, 2$, such that, for all price–income bundles (p, P, Y) , the Pareto frontier takes the form

$$u^1 + u^2 = K(p, P, Y).$$

In words, for a well chosen cardinalization of preferences, the Pareto frontier is a straight line with slope equal to -1 for all values of prices and income. The practical translation is that whenever agents behave efficiently, then for a well-chosen cardinalization of preferences, they must maximize *the sum* of individual utilities, as opposed to a weighted sum. This is equivalent to saying that, for that particular cardinalization, the household

members *must* have the same Pareto weights. An important implication is that if preferences are TU, any household model that assumes efficient outcomes—a primary example being the collective model that will be introduced below—must boil down to a unitary framework.

In practice, what do we need to assume about preferences for the TU property to hold? Partial answers were given by Bergstrom and Varian (1985), who consider the case of purely private consumption, and Bergstrom and Cornes (1983), who analyze a model in which all commodities but one are publicly consumed. These works are generalized by Chiappori and Gugl (2015), who provide necessary and sufficient conditions. These authors refer to the notion of conditional indirect utility introduced by Blundell, Chiappori, and Meghir (2005), defined as the maximum utility level an individual can reach by choosing the optimal bundle of private consumption for given values of private prices, total private expenditures, and conditional on a given vector of public consumption. They introduce a specific property of individual preferences, the affine conditional indirect utility (ACIU), which states that for a well chosen cardinal representation the conditional indirect utility is affine in total expenditures; and they show that TU obtains if and only if (i) each individual preferences exhibit the ACIU property, and (ii) the coefficient on total expenditures, which can be a function of private prices and public consumption, is the same for all individuals.

2.2.2 Income Pooling—Whatever argument is used to justify the unitary model, one of its main strengths is that it generates strong testable restrictions on household behavior. The most popular testable implication is income pooling. In a unitary setting, households maximize a single utility under a budget constraint; it follows that individual nonlabor incomes y^1 and y^2 affect household decisions only through the budget constraint, and only through the sum $y = y^1 + y^2$. As a consequence, after controlling for total nonlabor income y , individual nonlabor incomes should not affect household behavior. An equivalent statement is that while total income enters the budget constraint, any additional variable describing the respective magnitude of individual contributions—say, the ratio $y_1/(y_1 + y_2)$ —is in principle a distribution factor; as such, it cannot matter in a unitary context. A second well-known implication of the unitary model is that the Slutsky matrix constructed from household demands should be symmetric and negative semidefinite.

The income-pooling property is generally easier to test than Slutsky symmetry, if only because it does not require price variation, which is notoriously difficult to obtain. A description of several such tests will be provided in section 4. Let us simply mention here that the income-pooling property is generally rejected: individual nonlabor incomes affect household behavior in ways that go beyond the effect of total income on the budget constraint. A possible reason for the rejection of income pooling is that the unitary model aggregates individual preferences in a way that is not consistent with the data. It is plausible that households make actual decisions by assigning higher weight to the preferences of members that are perceived to be more important or, equivalently, to have more power within the household. The power of a person in a group is generally influenced by her or his outside options, which in turn depend on a collection of variables such as individual income, wealth, wages, human capital, and a series of distribution factors. If this is the case, households aggregate preferences in a way that depends on all those variables. In the unitary model, this possibility is ruled out. As indicated in the previous subsection, individual

preferences can only be aggregated by using some household index that is independent of any additional variable. On the contrary, several recent approaches—and particularly the collective model, which is discussed below—emphasize issues related to the intra-household allocation of power as central determinants of household behavior; and distribution factors matter precisely because they can influence this allocation. In that sense, the discussion around income pooling epitomizes the basic difference between unitary and non-unitary approaches.

The recent methods that have attempted to extend the static unitary model to a framework that allows for a more general way of aggregating individual preferences can be divided into two groups: models assuming that household members do not cooperate when making decisions (and using tools from noncooperative game theory), and models treating households as a group of individuals who cooperate and make efficient decisions. The next two subsections discuss these two literatures.⁷

2.3 Static Noncooperative Models

In noncooperative models, the key concept is a noncooperative Nash equilibrium, in which each spouse maximizes her or his own utility taking the partner's decisions as given. Several papers have used a non-cooperative approach to model household decisions, starting with the seminal papers by Leuthold (1968); Ashworth and Ulph (1981); and Bergstrom, Blume, and Varian (1986); and followed more recently by Chen and Woolley (2001); Browning (2000); Browning, Chiappori, and Lechene (2010); Lechene and Preston (2011); Cherchye, Demuynck, and De Rock (2011); Del Boca and Flinn (2012); d'Aspremont and Dos Santos Ferreira (2014); Boone et al. (2014); Del Boca and Flinn (2014); and Doepke and Tertilt (2014).

There are two main reasons for using noncooperative models to characterize household behavior. First, they may be directly relevant because it may be the case that at least some households behave in a noncooperative way (think, for instance, of households that are on the verge of a conflictual divorce). Second, as argued in Browning, Chiappori, and Lechene (2010), some cooperative models use the noncooperative outcome as a threat point. It is therefore important to study household decisions in a noncooperative setting to understand household behavior in a cooperative environment.

We now introduce the noncooperative model based on a Nash equilibrium. In most of the subsection, we will abstract from household production because almost all papers using noncooperative models have not considered this aspect of household behavior. We will consider a two-member household in which member i 's preferences depend on own private consumption, the spouse's private consumption, public consumption, own leisure, and the spouse's leisure. These preferences can therefore be represented using the following utility function: $u^i(Q, q^1, q^2, l^1, l^2)$. In a Nash noncooperative model, each household member makes independent decisions, taking as given the choices of the spouse. Given that the

⁷Some papers use a combination of cooperative and noncooperative methods. For instance, Del Boca and Flinn (2012) allow households to operate in a cooperative or noncooperative way and propose an estimator that enables the econometrician to evaluate which fraction of families are in each of the two regimes.

spouses make independent choices and some goods are public, an important concept in noncooperative models is the notion of an *individual contribution to public consumption*, which is defined as the part of public consumption that is provided by one of the household members. We will denote this variable with Q^i , with $q^1 + Q^2 = Q$.

Given this definition, the noncooperative model can be formally described as follows. Conditional on the spouse's choices, members 1 and 2 each choose their own private consumption, their individual contribution to public consumption, and their own leisure as a solution to the following problems:

$$\max_{Q^1, q^1, l^1} u^1(Q^1 + Q^2, q^1, q^2, l^1, l^2) \quad (5)$$

subject to

$$P'Q^1 + p'q^1 = Y^1,$$

and

$$\max_{Q^2, q^2, l^2} u^2(Q^1 + Q^2, q^1, q^2, l^1, l^2) \quad (6)$$

subject to

$$P'Q^2 + p'q^2 = Y^2.$$

The outcome of the household decision process is then assumed to be a Nash equilibrium, which is defined as the quantities q^{1*} , q^{2*} , l^{1*} , l^{2*} , Q^{1*} , and Q^{2*} that simultaneously solve problems (5) and (6).

The outcome of the Nash noncooperative model is generally inefficient. The reason for the inefficiency is intuitive. When household members decide on their individual contribution to public consumption and their own private consumption and leisure, they do not internalize the benefits their spouses derive from their choices. There are, however, two different sets of conditions under which the Nash noncooperative model generates outcomes that are efficient, and hence equivalent, to one of the solutions of a cooperative model. The first set is straightforward. If the household does not consume public goods and if each individual's utility does not depend on the spouse's private consumption and leisure, there is no conflict between the two spouses when making decisions and the outcome is efficient. The second set of conditions considers the case in which the household consumes exclusively public goods. Browning (2000) has shown that, in this environment, the solution of the noncooperative model is always efficient. The intuition behind this result is that, without private goods and leisure, household members have no incentive to underprovide the public good.

Economists have derived several implications that can be used to test the Nash noncooperative model. Lechene and Preston (2011) have shown that the Slutsky matrix derived from a noncooperative model does not satisfy the standard symmetry and negativity conditions. Instead, it is the sum of a symmetric matrix and an additional matrix whose rank is greater than one. We will see that a similar property applies to cooperative models, but the rank of the additional matrix is higher in a noncooperative environment.⁸ An additional testable implication, which is related to the income-pooling hypothesis, is derived in Bergstrom, Blume, and Varian (1986). The paper considers an environment in which two spouses use their individual income to privately provide a single public good and to purchase a private good. In that framework, the authors establish that there exist ranges of individual incomes for which both members contribute to the public good. They then show that, over those ranges, income is fully pooled, in the sense that a redistribution of income from one spouse to the other does not affect the household's choice of either public or private consumption. Browning, Chiappori, and Lechene (2010) extend the model proposed in Bergstrom, Blume, and Varian (1986) to the private provision of many public goods. They show that, with several public goods, there is at most one public commodity to which both spouses contribute; all other public goods are exclusively funded by one member or the other. Moreover, when both spouses contribute to a public good, the income-pooling result of Bergstrom, Blume, and Varian (1986) remains valid. Boone et al. (2014) generalize the noncooperative model considered in Browning, Chiappori, and Lechene (2010) to the case of endogenous income and derive results that are related to the findings of Bergstrom, Blume, and Varian (1986) and Browning, Chiappori, and Lechene (2010). They show that, in a noncooperative model with endogenous income, households can be in one of the following three regimes: a first regime in which the wife is a dictator, in the sense that the household's demand for public goods reflects exclusively the wife's preferences; a second regime in which the wife's as well as the husband's preferences are reflected in the household's demand for public consumption; and a third regime in which the husband is the dictator. Doepke and Tertilt (2014) generalize the results obtained in Browning, Chiappori, and Lechene (2010) and Boone et al. (2014) to an environment in which public goods are produced within the household. They show that noncooperation implies a narrow gender specialization in domestic chores, with each spouse specializing in the exclusive production of some goods. They also show that, in their context, specialization is exclusively driven by the spouses' respective wages. This result differs from the suggestion made by Lundberg and Pollak (1993), who propose a model in which social norms are the main driver of specialization.

Lastly, Chiappori and Naidoo (2015) consider an alternative, and more general, model in which agents first share aggregate household income and then privately provide the public goods according to a Nash equilibrium. In practice, thus, in programs (5) and (6), individual incomes Y_1 and Y_2 are replaced with general functions of the form $\rho(Y_1, Y_2)$ and $Y_1 + Y_2 - \rho(Y_1, Y_2)$; the goal being to investigate testable predictions stemming from the private provision of public goods only, not from the assumption that individuals can only use their

⁸d'Aspremont and Dos Santos Ferreira (2014) have studied cases that are intermediate between cooperation and noncooperation and showed that the rank of the asymmetric component may be even higher than in the noncooperative case.

own income (and cannot transfer resources between them). Using that framework, they derive implications that can be tested using cross-sectional data, in which no price variation is observed. They show that distribution factor proportionality, a property that will be discussed in detail in section 4, is satisfied in their noncooperative context when agents contribute to different public goods.⁹ Moreover, they derive additional restrictions and show that they are necessary and sufficient for household decisions to be the outcome of their model.

Very few papers have taken to the data some of the implications of noncooperative models. Boone et al. (2014) use the CEX and expenditure on children to test the noncooperative model against the unitary model and estimate the fraction of families that are in one of the three regimes predicted by their noncooperative framework. They reject the unitary model in favor of the noncooperative one for couples with two and three children, but for couples with only one child the unitary model cannot be rejected. Interestingly, they find that, under their assumption that the public good corresponds to expenditure on children, the majority of households are in a regime in which the wife is the dictator. We are not aware, however, of a paper that has tested the main hypothesis generated by the noncooperative model according to which all chores, except maybe one, are performed exclusively by one member. This hypothesis appears counterintuitive. More plausibly, some tasks, and possibly most tasks, are performed by both spouses, either jointly or alternatively. It is also not intuitive that, if there is an exclusive allocation of tasks, it is entirely driven by relative wages. But these are empirical questions and more research attempting to test noncooperative models with formally derived implications is required.

We conclude this subsection with one last remark. So far, no general identification result has been derived for noncooperative models. We believe that showing which part of the structure of noncooperative models can be recovered and the corresponding data requirements is a project worth pursuing.

2.4 Static Collective Models

In this subsection, we discuss static models that rely on cooperative outcomes to characterize household decisions. We will first outline the main assumption on which those models are based. We will then provide a mathematical formulation and discuss the concept of individual decision power, which is an important component of cooperative models of the household. We will conclude the subsection by introducing a two-stage formulation of those models, which is convenient to derive testable implications and identification results.

Similarly to noncooperative models, collective models of the household explicitly recognize that households generally consist of several individuals who may have distinct utilities. Differently from noncooperative formulations, collective models assume that household decisions are efficient in the sense that they are always on the Pareto frontier. A distinctive feature of collective models is their axiomatic nature. They do not rely on specific assumptions of the way household members achieve an efficient outcome, such as Nash

⁹As discussed earlier, agents can jointly contribute to, at most, one public good. If that is the case, the demand of that public good satisfies income pooling.

bargaining. They simply assume Pareto efficiency, which is satisfied if for any decision the household makes, there is no alternative choice that would have been preferred by all household members.

While the assumption of Pareto efficiency is undoubtedly restrictive, collective models are sufficiently general to include as special cases most of the static models used to study household behavior. One example is the unitary model, whose solution is clearly efficient as long as the household index W introduced in UMW is strictly increasing in the felicity functions $u^1(Q, q^1, l^1)$ and $u^2(Q, q^2, l^2)$. Models based on cooperative game theory are also particular cases of collective models. For instance, Nash bargaining models of household behavior, pioneered by Manser and Brown (1980) and McElroy and Horney (1981), generate an efficient outcome and are therefore part of the collective family. A last group of models that are special cases of the collective framework are models based on a market equilibrium, as proposed by Grossbard-Shechtman (1984), Gersbach and Haller (2001), and Edlund and Korn (2002).

The efficiency assumption is standard in many economic contexts and has often been applied to household behavior. Still, it needs careful justification. Within a static context, this assumption amounts to the requirement that married partners can find a way to take advantage of opportunities that make both of them better off. Because of proximity and durability of the relation, both partners are, in general, aware of the preferences and actions of each other. They should therefore be able to act cooperatively by reaching some binding agreement. Enforcement of such agreements can be achieved through mutual care and trust, by social norms, or by formal legal contracts. Alternatively, the agreement can be supported by repeated interactions with the possibility of punishment. A large literature in game theory, based on several “folk theorems,” suggests that in such situations efficiency should prevail.¹⁰ Even if one is not convinced by these arguments, at the very least, in a static environment, efficiency can be considered as a natural benchmark. In a dynamic framework, however, full efficiency becomes more debatable, because it requires commitment abilities that, in practice, may not be available to the spouses. The next section discusses how such restrictions to commitment can be introduced in the collective framework considered here.

We can now provide a formal characterization of the collective model. Pareto efficiency has a simple translation: the household behaves as if it were maximizing a weighted sum of the members’ utilities subject to a budget constraint and household production constraints. In a collective model, household decisions can therefore be derived as the solution to a problem of the form:

$$\max_{(x, X, d^1, D^1, d^2, D^2, l^1, l^2)} \quad (\text{P})$$

$$\mu^1 U^1(Q, q^1, l^1) + \mu^2 U^2(Q, q^2, l^2)$$

¹⁰Note, however, that folk theorems essentially apply to infinitely repeated interactions.

subject to

$$p' \left(\sum_{k=1}^N X_k + \sum_{h=1}^n x_h \right) + \sum_{i=1}^2 w^i \left(l^i + \sum_{k=1}^N D_k^i + \sum_{h=1}^n d_h^i \right) = Y, \quad (\text{BC})$$

$$Q_k = F_k(X_k, D_k) \quad (\text{HP})$$

and

$$\sum_i q_h^i = f_h(x_h, d_h)$$

for all k and h .

A few aspects of the collective model are worth discussing. First, the Pareto weights μ^1 and μ^2 generally depend on prices, wages, income, and distribution factors. As a consequence, the household makes decisions by aggregating preferences in a way that depends on all those variables. The collective model is therefore consistent with the empirical evidence collected by testing the income-pooling hypothesis, which suggests that individual income affects household behavior even after controlling for total income. Second, the Pareto weights have a natural interpretation in terms of *relative decision power*. To see this, observe first that the solution of the collective model does not change if the objective function is divided by the sum of the Pareto weights. Hence, only the relative weights $\mu = \mu^1/(\mu^1 + \mu^2)$ and $1 - \mu = \mu^2/(\mu^1 + \mu^2)$ are relevant to understanding household behavior. If μ is zero, member 1 has no say on household decisions, whereas if μ is equal to one, member 1 has perfect control over the choices made by the household. More generally, an increase in μ results in a move along the Pareto frontier that gives more resources and higher utility to member 1. In this sense, if we restrict ourselves to economic considerations, the Pareto weight μ can be interpreted as the relative decision power of member 1. Note, however, that there is one situation in which the Pareto weights do not represent the individual decision power—namely, the case of preferences satisfying the TU property. As indicated earlier, an important consequence of TU is that the two spouses always have equal Pareto weights or, equivalently, $\mu = \mu^1/(\mu^1 + \mu^2) = 1/2$. They cannot therefore represent the spouses' relative powers, at least for the cardinalization that generates the Pareto frontier with slope equal to -1 .

There exist situations under which the unitary and collective models generate the same set of household decisions. This is the case, for instance, if the relative decision power μ is constant and therefore does not depend on prices, wages, income, and distribution factors. In that case, the maximand in (P) is a standard utility function, and the choices generated can be rationalized using the unitary model. As pointed out in section 2.2, one example in which these conditions are fulfilled is provided by the TU model. We can therefore conclude that the unitary model is a good choice for modeling household behavior if one believes that the intra-household decision power is constant across households and over time.

Different conditions may also be derived in more specific situations. A standard example is risk sharing in an economy with one commodity and several states of the world. In such a context, Mazzocco (2007), using results from Gorman (1953) and Pollak (1971), shows that a group of individuals making efficient decisions behaves as a single agent if the individual utilities belong to the harmonic absolute risk aversion class with identical curvature parameter, which is known as the ISHARA condition. Under these restrictions, the assumption implicit in the unitary model that the household behaves as a single individual is therefore satisfied and the unitary and collective models generate the same outcomes. The ISHARA condition is fulfilled if, for example, all individual utilities exhibit constant absolute risk aversion or, alternatively, all utilities belong to the constant relative risk aversion class with identical risk aversion parameter.

In problem (P) we have only considered egotistic preferences. But the problem can easily be extended to preferences of the caring type by replacing $U^i(Q, q^i, P)$ with $W^i(u^1(Q, q^1, P), u^2(Q, q^2, P))$. It is important to point out, however, that the model with egotistical preferences (P) plays a special role. The reason for this is that the solution to the collective model with caring preferences must also be a solution of the collective model (P) with egotistical preferences.¹¹ The model with egotistical preferences provides, therefore, all the solutions generated by the model with caring preferences. The converse, however, is not true. A very unequal solution to (P) may fail to be Pareto efficient for caring preferences, since a transfer of resources from well-endowed but caring individuals to poorly endowed individuals may be Pareto improving in an environment with caring utilities.

The main feature that differentiates the collective model from the unitary framework is that household decisions depend on the intra-household decision power, in addition to prices and income. To formally make this point, we explicitly recognize the dependence of the intra-household decision power on prices, income, and distribution factors by setting $\mu = \mu(p, w, Y, z)$. The solution of collective models can then be written in the following form:

$$x(p, w, Y, \mu(p, w, Y, z)),$$

$$X(p, w, Y, \mu(p, w, Y, z)),$$

$$d(p, w, Y, \mu(p, w, Y, z)),$$

$$D(p, w, Y, \mu(p, w, Y, z)),$$

¹¹The argument underlying this result is straightforward. If a solution to the collective model with caring failed to be Pareto efficient with egotistical preferences, then there would exist an alternative allocation that would increase both agents' felicities. But such an allocation would also increase both agents' caring welfare, which contradicts the initial claim. See Chiappori (1992).

$$l^1(p, w, Y, \mu(p, w, Y, z)),$$

$$l^2(p, w, Y, \mu(p, w, Y, z)).$$

Making explicit the dependence of household decisions on the relative decision power is helpful because it highlights two aspects of collective models. First, distribution factors such as individual income or changes in divorce laws affect household decisions only through μ . This feature has been used to derive a set of testable implications for collective models of the household, which will be discussed in section 4. Second, the identification of individual preferences and production functions is complicated by the fact that the intra-household decision power is not observed and must be recovered from data. As a consequence, even if one were to observe the expenditure on commodities used in household production, the allocation of time to household and market production, prices, wages, income, and distribution factors, the only functions that are straightforward to recover are

$$\bar{d}(p, w, Y, z) = d(p, w, Y, \mu(p, w, Y, z)), \quad (7)$$

$$\bar{D}(p, w, Y, z) = D(p, w, Y, \mu(p, w, Y, z)),$$

$$\bar{l}^1(p, w, Y, z) = l^1(p, w, Y, \mu(p, w, Y, z)),$$

$$\bar{l}^2(p, w, Y, z) = l^2(p, w, Y, \mu(p, w, Y, z)),$$

$$\bar{x}(p, w, Y, z) = x(p, w, Y, \mu(p, w, Y, z)),$$

$$\bar{X}(p, w, Y, z) = X(p, w, Y, \mu(p, w, Y, z)).$$

which are a combination of demand functions and relative decision power.

In spite of this obstacle, results have been derived that enable a researcher to fully characterize static collective models. A first set of results provides necessary and sufficient conditions for a demand function to stem from a collective framework. A second set determines exclusion restrictions under which individual preferences and the intra-household decision power can be recovered from the sole observation of household behavior.

¹² The collective model is the only non-unitary model of the household for which similar results have been derived. They will be discussed in detail in section 4.

As we have argued along the way, the collective model only postulates that the household chooses an efficient outcome, but it does not specify which one. In most applications, there is a need to go one step further and predict the exact outcome of the decision process or, equivalently, the individual decision power. There are two possible ways to determine the actual efficient outcome. The first possibility is to specify a bargaining game played by the household members and the corresponding outside options available to each individual. An alternative path is to adopt a general-equilibrium approach and recognize that the members' outside options, and hence the efficient outcome selected by the household, are generally determined by the matching process through which the household is formed. For the sake of brevity, we only focus on the first approach.¹³

Bargaining models must make two choices. First, they have to select the bargaining concept to be used. Typically, bargaining models adopt an axiomatic approach by choosing a Nash-bargaining solution or, less frequently, a Kalai–Smorodinsky solution. Note that both solutions include Pareto efficiency as one of their axioms. The second choice that must be made is which outside option or threat point to use, i.e., the utility level a person could reach in the absence of an agreement that would generate the efficient outcome. This choice is crucial because, as argued by Chiappori, Donni, and Komunjer (2012), *any* Pareto-efficient decision process can be rationalized as the solution to a Nash-bargaining model for an adequate choice of the threat point. This means that no additional restriction can be introduced by the sole adoption of a Nash-bargaining framework; any new prediction must come from the definition of the threat points.

On this issue, the literature has mainly used two approaches. The first one relies on the idea that, with public goods, noncooperative behavior typically leads to inefficient outcomes that can be used as threats by the household members. Specifically, in the absence of an agreement, both members resort to a private provision of the public goods. This approach captures the idea that the person who would suffer more from this lack of cooperation—typically the person who has the higher valuation for public consumption—is likely to be more willing to compromise in order to reach an agreement. A variant of this idea is proposed by Lundberg and Pollak (1993) and is based on the notion of “separate spheres.” In their paper, each partner is assigned to a set of public goods to which she or he alone can contribute. This is defined as her or his sphere of responsibility or expertise which, according to Lundberg and Pollak, is determined by social norms. The threat point can then be defined as the value of being in a marriage in which the spouses act noncooperatively and privately provide the public goods in their sphere. In both cases, the cooperative bargaining model inherits some properties of the noncooperative framework used to construct the threat points. In particular, for any set of price and income such that the noncooperative model generates income pooling, so does the cooperative version.

¹²See Chiappori and Ekeland (2006, 2009) for a detailed discussion.

¹³The interested reader is referred to Browning, Chiappori, and Weiss (2014), chapter 7, and Chiappori (2017) for a general discussion.

The second approach to modeling the threat point is to select divorce as the no-agreement situation. The threat point can then be defined as the maximum utility a person could reach after divorce. The utilization of this approach, however, is not straightforward, since it requires information on the utility of divorcees. Using data on consumption and labor supply, one can recover an ordinal representation of preferences at best, whereas Nash-bargaining solutions require knowledge of the cardinal representation of preferences. Moreover, divorcees' utilities depend not only on their welfare when single, but also on their remarriage probability and their utility level in case of remarriage. The latter aspect, in turn, should be the outcome of a Nash-bargaining game of the same nature as the initial one.

All this indicates that when analyzing the bargaining situation of a couple, the threat points should not be considered as exogenous and, when divorce is considered, particular attention should be devoted to taking into account the remarriage probabilities and future utilities. These issues are of particular importance when analyzing the impact of large-scale reforms like those discussed in the introduction: assuming exogenous threat points in such contexts will generally generate misleading implications. Papers that structurally estimate collective models, such as, for instance, Bronson (2015); Gemici and Laufer (2012); Mazzocco, Ruiz, and Yamaguchi (2014b); and Voena (2015) deal with these issues by first solving the model by backward induction for each possible period and state of nature and then by simulating forward the path followed by a given individual. An alternative method for dealing with these issues is to use an equilibrium approach. Such an approach, which can be based on frictionless matching or search, explicitly recognizes the simultaneous nature of the problem.¹⁴

The previous discussion leads us to the following question: should a researcher choose a threat point based on noncooperation or divorce? We believe that the choice should be made using empirical evidence, since it is difficult to employ theoretical arguments in favor of one of the two choices. Bergstrom (1996) is one attempt to justify the adoption of a noncooperative threat point using theoretical arguments. His claims, however, are based on the assumption that the environment in which households make decisions is stationary. The birth of children, job promotions, job changes, health shocks, and simply aging are likely to invalidate this assumption, making Bergstrom's conclusion less persuasive.

We conclude this subsection by introducing an alternative formulation of collective models that relies on two separate stages. This specification is helpful to derive testable implications and identification results, and will be the basis of some of the discussion in section 4. We will proceed in three steps. We will first introduce the two-stage formulation for a simple model with no public goods and no household production. We will then generalize it by introducing public goods. Finally, we will consider the general case with household production.

Without public goods and household production and with egotistic preferences, any efficient decision can be described using the following two-stage process and any solution of the following two-stage process is efficient.¹⁵ In the first stage, the spouses choose jointly how

¹⁴See Browning, Chiappori, and Weiss (2014) and Chiappori (2017) for detailed analyses.

to allocate total household income Y between them. Denote with ρ^i the amount allocated to member i ; in the language of the collective model, (ρ^1, ρ^2) defines the *sharing rule*. Then, in the second stage, each spouse optimally chooses private consumption and leisure given ρ^i . The intuition behind this result is based on the second welfare theorem. Without altruism and public goods, the household can be considered as a small economy without externalities. As a consequence, from the second welfare theorem, any Pareto-efficient decision can be decentralized by choosing the correct transfer to the two spouses.

To formally describe the two-stage approach, it is useful to start from the second stage. Conditional on an arbitrary amount of resources allocated by the household to member i , ρ^i , in the second stage this spouse chooses private consumption and leisure as a solution to the following simple single-agent problem:

$$V_i(p, w^i, \rho^i) = \max_{q^i, l^i} U^i(q^i, l^i) \quad (8)$$

subject to

$$pq^i + w^i l^i = \rho^i,$$

where $V^i(p, w^i, \rho^i)$ is the indirect utility function that measures the welfare of member i , given prices and the hourly wage if she or he is endowed with ρ^i . Then, in the first stage, the household uses the indirect-utility functions V^1 and V^2 and the intrahousehold decision power μ to optimally allocate the household full income between the two spouses:

$$\max_{\rho^1, \rho^2} \mu V^1(p, w^1, \rho^1) + (1 - \mu) V^2(p, w^2, \rho^2)$$

subject to

$$\rho^1 + \rho^2 = \sum_{i=1}^2 y^i + \sum_{i=1}^2 w^i T^i = Y.$$

Note that the sharing rule (ρ^1, ρ^2) is simply the solution to this first-stage problem.¹⁶

The two-stage framework just introduced relies on the strong assumption that all commodities are privately consumed. Relaxing this assumption is important because the existence of public consumption is one of the motives of household formation. There are two approaches that can be used to construct the two-stage formulation for the case in which households consume public goods. The first approach relies on the notion of the *conditional*

¹⁵For a proof, see Chiappori (1992).

¹⁶The outcome of a collective model can be derived using the described two-stage formulation also for caring preferences. As argued above, any efficient decision with caring preferences can be obtained using the felicity functions that define the caring preferences as the egotistical utility functions. As a consequence, the solution of the collective model with caring preferences can be derived using the two-stage formulation, but the converse is not true.

sharing rule, which was introduced by Blundell, Chiappori, and Meghir (2005). The second approach is based on Lindahl prices. Since the two methods are similar to the two-stage formulation with only private goods, we only provide a short description.

The first approach is based on the following simple idea. The existence of public goods introduces externalities in the household decision process. To obtain an efficient decision, these externalities must be managed at the household level, since individual members are unable to deal with them on their own. But once the household has solved the externality problem, the two spouses can optimally choose their own private consumption and leisure. This insight can be implemented using the following two-stage method. In stage one, the household manages the externalities by choosing the consumption of public goods and the distribution of remaining income between members. The first stage is therefore identical to the first stage with only private goods except that public consumption is also selected and only income net of public consumption is allocated. Given the level of public consumption selected in the first stage, the second stage is identical to the one described for the private-goods case: the spouses spend their allotted amount of resources on private consumption and leisure so as to maximize their individual utility, *conditionally on* the level of public expenditures selected in the first stage—hence the term *conditional sharing rule*. Similarly to the situation with only private goods, with public consumption any efficient decision can be represented as stemming from the two-stage process just described. However, unlike the private-goods case, it is generally not the case that, for given public consumptions, any conditional sharing rule generates an efficient outcome; indeed, only sharing rules that satisfy the standard Bowen–Lindahl–Samuelson conditions generate efficient outcomes.

The second approach relies on an old result in public economics stating that, in the presence of public goods, any efficient allocation can be decentralized using personal prices for the public goods. The personal prices are also known as Lindahl prices and capture the willingness to pay of the household members for the public goods. This result establishes a symmetry between private and public goods. For private goods, the spouses face identical prices and purchase different quantities, whereas for public goods, household members buy the same quantity at individual-specific prices.¹⁷ Using the Lindahl prices, a collective household behaves as if it were using a two-stage process. In stage one, the household chooses the individual prices for the public goods (which must add up to market prices) and an allocation of total income between members. In stage two, the spouses choose optimally private consumption, leisure, and public consumption under a budget constraint constructed using their Lindahl prices.

Lastly, household production can readily be introduced in this framework. When all produced commodities are privately consumed, the sharing-rule approach can immediately be extended, still invoking the second welfare theorem.¹⁸ If some goods produced within the household are public, it is important to distinguish two cases, depending on whether these goods can be freely sold and purchased on the market. If the public commodities are

¹⁷See Chiappori and Ekeland (2009) for a general presentation. For applications, see for instance Donni (2009) and Cherchye, De Rock, and Vermeulen (2007).

¹⁸Note, however, that the second welfare theorem now requires convexity of the production sets.

marketable, one can use the two-stage method based on the conditional sharing rule or the alternative method based on Lindahl prices. If the produced commodities cannot be purchased and sold on the market, a two-stage formulation of the collective model can still be constructed using the conditional sharing rule. However, the time (and possibly the other inputs) allocated to the production of the public goods create externalities within the household; they must therefore be chosen in the first stage jointly with the allocation of full income between spouses. Then, in the second stage, conditional on the amount of public goods produced and the amount of time left, the two household members choose leisure and the amount of hours and commodities to employ in the production of the private goods.

3. Intertemporal Models of Household Decisions

The first part of the survey has focused on static models of household decisions, which are clearly the right starting point. They are easier to develop and test, and they provide researchers with a relatively simple framework that can be used to study identification and perform estimation. They have, however, two main limitations. First, they cannot be used to answer questions and evaluate policies that have an intertemporal dimension; a serious weakness indeed, since most policies are characterized by some intertemporal aspects. The choice between a joint or individual taxation system for married couples outlined at the beginning of the survey is a good example of a policy whose effects are difficult to evaluate using static models. A second and related limitation is that static models cannot be employed to analyze the evolution of intra-household decision powers. Since the latter typically depend on over-time variation in the individual outside options, it can only be studied using models that can account for intertemporal aspects of household behavior.

The second part of this survey addresses these two limitations by reviewing intertemporal models of household decisions. It is divided into four subsections. The first subsection introduces an intertemporal model whose main advantage is to have as special cases the most popular static and intertemporal frameworks used to study household decisions. An important feature of the proposed model is that the individual outside options and their evolution play a critical role by affecting the intra-household decision power. In the second subsection, we propose a three-stage formulation of the household decision process and argue that its solution is equivalent to the solution obtained from the intertemporal model described in the first subsection. In the third subsection, we use the three-stage formulation to establish that the most commonly used static and intertemporal models are a special case of the general framework, and to describe the features of household behavior on which different models focus. Finally, we discuss the effect that decisions made at an earlier stage of the relationship can have on the spouses' outside options.

3.1 A General Framework: The Limited-Commitment Model

We start by introducing an intertemporal model that generalizes the static collective model to a dynamic environment. To accomplish this, two issues must be addressed. The first one is whether household members can commit to future plans. The second issue is whether household members have full or partial information about variables that characterize the spouse, such as income, savings, and preferences. To understand why these two issues arise

in a dynamic setting, remember that the main assumption of the static collective model is that household decisions are Pareto efficient. For this assumption to be satisfied in an intertemporal environment, household members must be able to commit to future allocations of resources and have full information about the spouse. Under these two assumptions, household decisions will be efficient in the sense that they are on the ex ante or “first-best” Pareto frontier, i.e., the intertemporal Pareto frontier at the time of household formation.

In what follows, we assume complete information in the sense that household members have full knowledge of preferences and probability distributions over the variables of interest. There is little research on household decision models with partial information, especially regarding the issues of testability and identifiability. We are aware of only one paper, Dubois and Ligon (2011), that develops and tests a model of household decisions in which information is not complete, and one paper, Ashraf (2009), that provides suggestive evidence that it is important to account for incomplete information to understand household decisions.¹⁹ Clearly, this is an area in which future research is needed.

We will relax, however, the assumption of commitment, which is clearly strong. It would imply that, after ten years of marriage, spouses cannot change the plans they formulated at the time of household formation, irrespective of the events that might have occurred since then. While such an assumption may perhaps be found to be empirically relevant, imposing it as a prerequisite seems excessive. For this reason, in this subsection we will consider an intertemporal model in which household members cannot commit. A model with full commitment will be described in the next subsection; we will show that it is a particular case of our general, limited-commitment model.

The model considered in this subsection is the limited-commitment intertemporal collective (LIC) model. To simplify the discussion, we consider a simpler version of the model with two periods and two states of nature. This is clearly an important simplification, especially if one is interested in estimating and simulating the LIC model to answer policy questions. But it enables us to describe the main features of the LIC model without having to deal with the cumbersome notation that would be required by a model with an arbitrary number of periods and states of nature. For a general version of the LIC model, one can read, for instance, Mazzocco (2001); Ligon (2002); Aura (2005); Mazzocco (2007); Gemici and Laufer (2012); Mazzocco, Ruiz, and Yamaguchi (2014b); Voena (2015); or Bronson (2015).

¹⁹Dubois and Ligon (2011) build three models to understand how food is allocated across household members in the rural Philippines and derive implications to test them. The first model is a collective framework. The second model allows the productivity of household members to be affected by the amount of food they consume. In the third model, Dubois and Ligon allow for incomplete information by considering a case in which the labor effort of household members is not observed. Their test outcomes indicate that models in which individual productivities depend on food intake and effort is not observed are better characterizations of how food is allocated across household members in rural villages. Ashraf (2009) uses data generated by a field experiment to study the role of information and communication within a household on financial decisions of married people in the Philippines. In the experiment, individuals were given a sum of money that corresponds to about a day's wage and were asked to deposit the money to their bank account, to their spouse's bank account, or to choose consumption in the form of committed consumption or cash. The main finding is that on average, in households in which the wife controls savings decisions, when choices are not observed by the wife, men choose to put the money in their personal account. In the same households, when choices are observable by the wife but she cannot communicate her preferences, men choose to commit the money to their own consumption. Finally, when choices are observable and wives can communicate their preferences to their husband, men choose to put the money in the wife's account. Wives in households in which the husband controls savings decisions make similar choices. These findings suggest that restrictions on information available to household members affect family decisions.

Although we concentrate here on the LIC model, it is worth mentioning that alternative intertemporal models of the collective type have been proposed. For instance, Basu (2006) develops a game-theoretical version of a collective model without commitment. The main difference between the LIC model considered in the survey and Basu's model is that the latter produces allocations that are generally inefficient, even from an ex post perspective. Lundberg and Pollak (2003) develop a two-period model in which couples make noncooperative migration and divorce decisions in the first period and share efficiently their joint resources in the second period, in case they stay married. Other authors have advocated the use of an intertemporal generalization of Nash bargaining in which the household solves a sequence of static Nash-bargaining problems, one for each time period. See for instance Greenwood, Guner, and Knowles (2003) and Lundberg and Pollak (2009) for a discussion of such a model. A limitation of the repeated Nash-bargaining model is that a household renegotiates the allocation of resources in each period even if the changes in the variables affecting the outside options are small, such as a one-dollar increase in yearly earnings. We have decided to focus on the LIC model instead of using one of the alternative frameworks because it includes as special cases most of the models used to characterize household decisions and it is a natural generalization of the static collective model discussed in the first part of the survey.

In the LIC model we consider, households live for two periods. In the first period, there is no uncertainty, whereas in the second period one of two possible states of nature can occur. We will denote with t a time period, with $\omega = \{\omega_L, \omega_H\}$ a state of nature, and with $P(\omega)$ its probability. The rest of the setting is identical to the one used for the static collective model with two additional features: household members can save using a risk-free asset with gross return R and in the second period, individual preferences are discounted using a discount factor β , which is assumed to be identical across spouses.²⁰

The main feature of the LIC model is that household decisions are efficient subject to the constraint that in each period and state of nature, both spouses can choose to leave the household and take the best available outside option if the level of welfare it provides is higher than the welfare provided by staying in the household at the current allocation of resources. This does not imply that household members will actually leave in that event. In many cases, the threat to leave will trigger a renegotiation that modifies the intrahousehold allocation plans to restore individual rationality, i.e., to make the new plans at least as good as the outside options for each member. In what follows, the utility provided by the best outside option of member i will be denoted with $\bar{u}^i(X)$, where X is a set of variables that can affect it.

In such a context, the definition of the outside option is crucial. Here the discussion closely mimics the description of the threat point provided for static models. One possibility is to define the outside option as the value of being divorced in the current period and making optimal decisions from next period onward. This definition is the most commonly used in

²⁰We assume equal discount factors to simplify the notation. The discussion that follows is not affected by this assumption. See Adams et al. (2014) for the implications of heterogeneity in discount factors for consumer intertemporal behavior and time inconsistency.

papers that study dynamic household behavior. For instance, it is used in models that test and estimate intertemporal collective models, such as Aura (2005); Mazzocco (2007); Gemici (2011); Mazzocco, Ruiz, and Yamaguchi (2014b); Bronson (2015); and Voena (2015), as well as in models based on search, such as Jacquemet and Robin (2013) and Gousse (2014). An alternative definition is considered in Lundberg and Pollak (1993). The outside option is defined as the individual welfare if the two spouses choose to stay married and not cooperate. Most of what follows is compatible with both interpretations and other alternatives. We will come back to this issue later on.

We will now introduce the LIC model by adding two sets of features to the static collective model considered in the first part of the survey. First, the static model must be augmented to capture the intertemporal aspects of household decisions and the existence of uncertainty. To do that, we add to the objective function of the static collective model the discounted expected value generated by the household in the second period, which takes the following form:

$$\beta \left[\mu^1(\Theta) u^1(Q_{\omega_L}, q_{\omega_L}^1, l_{\omega_L}^1) + \mu^2(\Theta) u^2(Q_{\omega_L}, q_{\omega_L}^2, l_{\omega_L}^2) \right] P(\omega_L) \\ + \beta \left[\mu^1(\Theta) u^1(Q_{\omega_H}, q_{\omega_H}^1, l_{\omega_H}^1) + \mu^2(\Theta) u^2(Q_{\omega_H}, q_{\omega_H}^2, l_{\omega_H}^2) \right] P(\omega_H),$$

where Θ includes all the variables that affect the relative decision power. Moreover, we include two second-period budget constraints, one for each state of the world:

$$p'_{\omega} \left(\sum_{k=1}^N X_{k,\omega} + \sum_{h=1}^n x_{h,\omega} \right) + \sum_{i=1}^2 \omega_{\omega}^i \left(l_{\omega}^i + \sum_{k=1}^N D_{k,\omega}^i + \sum_{h=1}^n d_{h,\omega}^i \right) = Y_{\omega} + Rb, \text{ for } \omega = \omega_L, \omega_H.$$

The second addition to the static collective models is necessary to introduce the restriction that the two household members cannot commit to future plans. This feature can be incorporated by including the following participation constraints for the two household members:

$$u^1(Q_{\omega}, q_{\omega}^1, l_{\omega}^1) \geq \bar{u}_{\omega}^1(X) \text{ for } \omega = \omega_L, \omega_H,$$

$$u^2(Q_{\omega}, q_{\omega}^2, l_{\omega}^2) \geq \bar{u}_{\omega}^2(X) \text{ for } \omega = \omega_L, \omega_H.$$

The idea behind the participation constraints is straightforward: for a set of decisions to be a solution of the LIC model, it must be that, for both household members, welfare in the second period is greater than the best outside option in both states of nature.

Now that we have described the new features of the LIC model, we can formally present it. In an LIC model, households choose the expenditure on commodities employed in household production, the allocation of time to household production and leisure, and savings as a solution to the following problem:

$$\begin{aligned} \max \quad & \mu^1(\Theta)u^1(Q, q^1, l^1) + \mu^2(\Theta)u^2(Q, q^2, l^2) \\ & + \beta \left[\sum_{i=1}^2 \mu^i(\Theta)u^i(Q_{\omega_L}, q_{\omega_L}^i, l_{\omega_L}^i) P(\omega_L) + \sum_{i=1}^2 \mu^i(\Theta)u^i(Q_{\omega_H}, q_{\omega_H}^i, l_{\omega_H}^i) P(\omega_H) \right] \end{aligned} \quad (\text{LIC})$$

subject to

$$\begin{aligned} & p' \left(\sum_{k=1}^N X_k + \sum_{h=1}^n x_h \right) \\ & + \sum_{i=1}^2 \omega^i \left(l^i + \sum_{k=1}^N D_k^i + \sum_{h=1}^n d_h^i \right) \\ & + b = Y, \quad p'_{\omega} \left(\sum_{k=1}^N X_{k,\omega} + \sum_{h=1}^n x_{h,\omega} \right) \\ & + \sum_{i=1}^2 \omega_{\omega}^i \left(l_{\omega}^i + \sum_{k=1}^N D_{k,\omega}^i + \sum_{h=1}^n d_{h,\omega}^i \right) = Y_{\omega} + Rb, \end{aligned}$$

for $\omega = \omega_L, \omega_H$;

$$u^1(Q_{\omega}, q_{\omega}^1, l_{\omega}^1) \geq \bar{u}_{\omega}^1(X)$$

and

$$u^2(Q_{\omega}, q_{\omega}^2, l_{\omega}^2) \geq \bar{u}_{\omega}^2(X),$$

for $\omega = \omega_L, \omega_H$;

$$Q_k = F_k(X_k, D_k)$$

and

$$\sum_i q_h^i = f_h(x_h, d_h)$$

for all k and h ;

$$Q_{k,\omega} = F_k(X_{k,\omega}, D_{k,\omega})$$

and

$$\sum_i q_{h,\omega}^i = f_h(x_{h,\omega}, d_{h,\omega})$$

for all k, h , and $\omega = \omega_L, \omega_H$

Just as in the static collective model, the function $\mu^i(\Theta)$ is the Pareto weight of member i that describes spouse i 's decision power or, equivalently, the weight of her preferences in the decision process. As argued in the first part of the survey, the intra-household decision power generally depends on a set of variables Θ , which includes the “distribution factors.” An interesting feature of the LIC model that will be further discussed later in this section is that these variables are fully defined *at the beginning* of the relationship and cannot change during the marriage. The only variables that can change over time are the ones that enter the members' outside options and are included in X . They are, therefore, the only variables that may change the household decision process between the first and second periods, as we shall see below.²¹

To better understand how the LIC model operates, it is helpful to incorporate the participation constraints in the objective function of (LIC) using a standard Lagrangian multiplier method. Let $\hat{\lambda}_{\omega}^i(W)$ be the Kuhn-Tucker multiplier associated with member i 's participation constraint in state ω ; notice that it depends on all the exogenous variables W that enter the household's problem. The variables that affect the initial decision power, Θ , and the best outside options, X , are therefore included in W . One can then form the Lagrangian that corresponds to the LIC model by adding to the objective function the participation constraints times the multipliers. The household's objective function takes therefore the following form:

$$\begin{aligned} & \mu^1(\Theta)u^1(Q, q^1, l^1) + \mu^2(\Theta)u^2(Q, q^2, l^2) \\ & + \beta \sum_{i=1}^2 \left[\mu^i(\Theta)u^i(Q_{\omega_L}, q_{\omega_L}^i, l_{\omega_L}^i) + \frac{\hat{\lambda}_{\omega_L}^i(W)}{\beta P_{(\omega_L)}} \left[u^i(Q_{\omega_L}, q_{\omega_L}^i, l_{\omega_L}^i) - \bar{u}_{\omega_L}^i(X) \right] \right] P_{(\omega_L)} \\ & + \beta \sum_{i=1}^2 \left[\mu^i(\Theta)u^i(Q_{\omega_H}, q_{\omega_H}^i, l_{\omega_H}^i) + \frac{\hat{\lambda}_{\omega_H}^i(W)}{\beta P_{(\omega_H)}} \left[u^i(Q_{\omega_H}, q_{\omega_H}^i, l_{\omega_H}^i) - \bar{u}_{\omega_H}^i(X) \right] \right] P_{(\omega_H)}. \end{aligned} \quad (9)$$

There are several features of the household's objective function (9) that are worth discussing. First, a well-known property of the Kuhn-Tucker multipliers is that they are

²¹Of course, the past realization of a variable may have influenced the decision power of a spouse and hence belong to the vector Θ , while the subsequent realization of the same variable may have an effect on the outside option and hence belong to X .

equal to zero if the participation constraints do not bind, and positive otherwise. This implies that the multipliers become relevant in the household decision process only if, at the current plan, one of the household members would strictly prefer the best outside option. This feature of the model will enable us to understand and interpret the intertemporal changes in decision power. Secondly, the second-period decisions, in any state of the world, maximize a weighted sum of utilities; as such, they are Pareto efficient in the static, “ex post” sense.

Thirdly, generally the outside options $\bar{u}_{\omega}^i(X)$ depend on variables chosen by the household, such as savings. Therefore, when making decisions, households take into account their effect on $\bar{u}_{\omega}^i(X)$. To simplify the notation, in the rest of the subsection we will aggregate all the parts in (9) that depend on $\bar{u}_{\omega}^i(X)$ under the function $g(X, W)$. This third feature highlights that, in an intertemporal collective model, when a household makes saving choices, it takes into account not only the standard consumption smoothing and precautionary motives, but also the effect that its decisions have on the members’ outside options. The last subsection provides further discussion of the relationship between household’s choices and $\bar{u}_{\omega}^i(X)$.

The final noteworthy feature of the objective function is that, inside the summation, for each individual we can collect the second period utility function and have just one weight in front of it. Using the last two features, the household’s objective function can be rewritten in the following simpler and more intuitive form:

$$\begin{aligned} & \mu^1(\Theta)\mu^1(Q, q^1, l^1) \\ & + \mu^2(\Theta)\mu^2(Q, q^2, l^2) \\ & + \beta \left[\sum_{i=1}^2 M_{\omega_L}^i(W) u^i(Q_{\omega_L}, q_{\omega_L}^i, l_{\omega_L}^i) \times P(\omega_L) + \sum_{i=1}^2 M_{\omega_H}^i(W) u^i(Q_{\omega_H}, q_{\omega_H}^i, l_{\omega_H}^i) \times P(\omega_H) \right] + g(X, W), \end{aligned} \quad (10)$$

where

$$M_{\omega}^i(W) = \mu^i(\Theta) + \frac{\hat{\lambda}_{\omega}^i(W)}{\beta P(\omega)}. \quad (11)$$

Using this alternative formulation of the objective function, the household’s problem can be written as follows:

$$\begin{aligned}
& \max \mu^1(\Theta) u^1(Q, q^1, l^1) \\
& + \mu^2(\Theta) u^2(Q, q^2, l^2) \\
& + \beta \left[\sum_{i=1}^2 M_{\omega_L}^i(W) u^i(Q_{\omega_L}, q_{\omega_L}^i, l_{\omega_L}^i) P(\omega_L) + \sum_{i=1}^2 M_{\omega_H}^i(W) u^i(Q_{\omega_H}, q_{\omega_H}^i, l_{\omega_H}^i) P(\omega_H) \right] + g(X, W)
\end{aligned}
\tag{12}$$

subject to

$$\begin{aligned}
& p' \left(\sum_{k=1}^N X_k + \sum_{h=1}^n x_h \right) + \sum_{i=1}^2 \omega^i \left(l^i + \sum_{k=1}^N D_k^i + \sum_{h=1}^n d_h^i \right) + b = Y, \\
& p'_{\omega} \left(\sum_{k=1}^N X_{k,\omega} + \sum_{h=1}^n x_{h,\omega} \right) + \sum_{i=1}^2 \omega_{\omega}^i \left(l_{\omega}^i + \sum_{k=1}^N D_{k,\omega}^i + \sum_{h=1}^n d_{h,\omega}^i \right) = Y_{\omega} + Rb,
\end{aligned}$$

for $\Omega = \Omega_L, \Omega_H$,

$$Q_k = F_k(X_k, D_k)$$

and

$$\sum_i q_h^i = f_h(x_h, d_h)$$

for all k and h ,

$$Q_{k,\omega} = F_k(X_{k,\omega}, D_{k,\omega})$$

and

$$\sum_i q_{h,\omega}^i = f_h(x_{h,\omega}, d_{h,\omega})$$

for all k, h , and $\omega = \omega_L, \omega_H$.

We now use this alternative formulation to provide some insight on how the model accounts for the main aspects of household decisions. As mentioned above, the outside options may correspond to the value of noncooperation while married or to the value of divorce. In the following discussion, and for the sake of brevity, we use the terms “stay married” or “remain

married” to indicate a household in which the spouses are married and cooperate, as opposed to a situation in which they are married and do not cooperate.

In the LIC model, at the time of household formation the two spouses choose a contingent plan for their future lives. The initial contingent plan is based on the initial intra-household decision power $\mu^i(\Theta)$, which may be determined by bargaining or any other decision process, and will in general depend on the best outside options available to the two spouses at the time of household formation. Since $\mu^i(\Theta)$ is the decision power at the time of household formation, Θ includes all the variables that have an effect on the spouses’ bargaining strength at the time they marry, such as their education and human capital, their wages and wage prospects, the number of children, and the sex ratio in the marriage market. It is worth noting that $\mu^i(\Theta)$ corresponds to the intra-household decision power that characterizes the static collective model.

After the two spouses marry, they start living their life and, in doing so, experience various events. These events can be divided into two categories. The first category includes all events after which the two household members, at the current Pareto weights, are both better off staying married than taking their best outside options. Formally, these are shocks after which neither of the spouses’ participation constraints bind. In such cases, the household members will stick to the current contingent plan and continue their married life. The effect of these types of shocks are illustrated graphically in the first panel of figure 1, which depicts in the vertical and horizontal axis the expected lifetime utility of members 1 and 2 computed in a given period. The figure illustrates that these shocks have two main effects. First, they shift the Pareto frontier outward or inward. In the figure, we consider a positive event, such as a wage or income increase, that shifts the frontier outward. The second effect is to move the value of the spouses’ best outside options. For clarity, we illustrate the case in which only the value of the second household member increases. The first panel of figure 1 shows that after a shock of this type, the household remains married and continues to use the current contingent plan. Note that two conditions must be satisfied to guarantee this conclusion. First, the two best outside options must cross inside the frontier, which implies that after the shock the household still generates positive gains from marriage. It is therefore optimal for the two spouses to remain married. Second, after the event, the point on the new Pareto frontier that corresponds to the initial plan must still be inside the orthant defined by the two outside options. This indicates that, at the current plan, the participation constraints do not bind. As a consequence, it is optimal for the two spouses to continue their marriage using the contingent allocation of resources chosen at the beginning of their marriage, since any change of plan reduces (ex ante) efficiency.

The second category of events is composed of shocks that would make one or both spouses better off with the best outside option, if there is no change in the way resources are allocated. Examples of such shocks are a job promotion for one of the spouses or an inheritance received by one of the household members. When such events occur, there are two possible outcomes for the household. If, after the event, the household still generates a positive amount of marital surplus, it is optimal for the household to change the original plan in a way that satisfies the binding participation constraints. In the model, the household achieves this goal by increasing the Pareto weight of the constrained spouse, and therefore

the amount of resources allocated to that spouse in the current and future periods;

technically, the initial weight, $\mu^i(\Theta)$, is increased by the amount $\frac{\hat{\lambda}_\omega^i(W)}{\beta P(\omega_L)}$. As shown in Kocherlakota (1996) and Ligon, Thomas, and Worrall (2002), this change in decision power is optimal because it increases the resources allocated to the constrained individual just enough to make her or him indifferent between staying married at the new plan and taking the outside option. Any additional increase in decision power and reallocation of resources would be suboptimal from an ex ante viewpoint, because it would make the deviation from the initial plan too large.²² The effect of this type of shock is illustrated in the second panel of figure 1. The figure illustrates that after a shock a household remains married, but with different intra-household decision power and contingent plan, if the following two conditions are fulfilled. First, after the event, the spouses' best outside options still cross inside the Pareto frontier, indicating that the household produces a positive surplus even after the shock. It is therefore optimal for the two spouses to stay married. Second, after the shock, the initial contingent plan demands that the household chooses a point on the Pareto frontier that is outside the orthant generated by the two outside options, thus violating the participation constraint of one spouse. In the illustrated case, the point is to the left of the best outside option of the second spouse, implying that member 2's participation constraint is not satisfied. As a consequence, at the initial plan, that spouse is not willing to remain married. The household must therefore change the allocation of resources to convince her to stay. This is achieved by increasing the decision power of the constrained spouse by an amount that brings her exactly at her new outside option.

Finally, it may be the case that, after the shock takes place, there is no point on the Pareto frontier that provides each spouse with at least their outside option utility. Then the household no longer produces a positive amount of marital surplus, which implies that there is no renegotiation that could make both spouses better off as married. They will therefore choose their best outside options. This case is depicted in the third panel of figure 1, where, after the event, the two best outside options intersect outside the Pareto frontier.

The previous discussion emphasizes an important feature of the LIC model, namely that it imposes some discipline to the way the intra-household decision power can change over time. It can change only if one of the two spouses can make a *credible* threat of taking the outside option. This requires an event significant enough to generate a substantial increase in the outside value. Assume, for instance, that the outside option is defined by the person's expected utility if divorced. If the magnitude of the change is trivial—say, a ten-dollar raise in yearly earnings, or a 1 percent raise in hourly wage—the model will reasonably predict that there will be no change in decision power. Even a one-time cash transfer that corresponds to 10, 20, or even 50 percent of yearly income is unlikely to trigger a change in intra-household decision power, except maybe in marginal households for which the marital surplus was already minimal. On the other hand, serious changes will typically impact divorce probabilities, and this may in principle occur whether they are positive or negative; it depends on how the effect of the shock on the household Pareto frontier compares to its

²²This result depends on the assumption that a renegotiation of the plan is costless.

effect on the outside options. On a more technical side, note that the random process followed by the Pareto weights is Markovian: its evolution after some date t only depends on the values of the weights at t and not on their past history. This property allows for empirical tests of the LIC model.

The discussion also clarifies which variables can trigger changes in intra-household decision power: all variables that can modify the spouses' best outside options. In practice, changes in earnings, fertility, welfare subsidies, alimonies, and cash transfers can lead to a modification of $M_{\omega_L}^i(W)$.

In this section, we have pointed out one of the strengths of the LIC model: it enables one to determine when and how the intrahousehold decision power should change and which variables should have an effect on it. But there is one caveat that is important to discuss. Just as in the static framework, the model does not pin down the decision power *at the time of marriage* $\mu^i(\Theta)$. Figure 1 illustrates why that is the case. Consider two individuals with outside options $\bar{u}^1$ and $\bar{u}^2$ who choose to get married. At the time of marriage, any allocation on the Pareto frontier that is inside the orthant generated by $\bar{u}^1$ and $\bar{u}^2$ makes both spouses better off being married than single. All these allocations and the corresponding Pareto weights are therefore consistent with the marriage choice. Fortunately, and again just as in the static framework, the initial decision power can be identified using data that provide information on the point initially chosen by the household. For instance, if one observes the amount of leisure consumed by the wife and husband, the ratio of those two variables can be used to infer $\mu^i(\Theta)$. Alternatively, one can construct an “upstream” model describing the determination of Pareto weights; such a model can be based on an explicit bargaining framework, or alternatively be based on equilibrium approaches involving matching or search components (as in Chiappori, Costa Dias, and Meghir forthcoming). Such upstream models are currently an area of very active research. As mentioned above, they are, however, outside the scope of this survey.

The final specification of the LIC model is helpful for one last reason. It can be used to show that intertemporal household decisions with limited commitment can be decomposed in three simple stages that capture different aspects of household behavior. This is the subject of the next subsection.

3.2 A Three-Stage Formulation

In this subsection, we introduce a three-stage formulation of household decisions that has the same solution as the LIC model. This alternative formulation will be used to argue that the most popular household models correspond to a subset of the three stages. They therefore account for different aspects of household behavior.

We start by providing some insight on the distinct aspects of household behavior captured by the different stages. The first stage deals with the allocation of the household's lifetime resources across periods and states of nature. It therefore focuses on the dynamic and contingent nature of household decisions. For this reason, we will denote this stage with the term *intertemporal stage*. The second stage considers a particular period and state of nature

and, given the amount of resources allocated in the first stage to that period and state, it determines the optimal purchase of commodities to use in household production and the optimal allocation of time across labor supply, leisure, and household production. We will denominate this stage as the *resource-allocation stage*. The third and final stage deals with the same period and state of nature considered in the second stage and, conditional on the decisions made in that stage, it investigates the optimal allocation of private goods between spouses. We will refer to this stage as the *intra-household allocation stage*.

We can now provide a formal description of the three-stage formulation, starting with the *intra-household allocation stage*. In this stage, the household takes as given an arbitrary amount of private and public goods produced within the household and an arbitrary number of hours allocated to leisure. As discussed above, these quantities are optimally chosen in the *resource-allocation stage*. We will denote with $\bar{Q}_{t,\omega}$, $\bar{q}_{t,\omega}$, and $\bar{l}_{t,\omega}^i$, respectively, the vector of public goods, the vector of private goods, and the amount of time assigned to leisure that the household takes as given. We will also denote with $M_{t,\omega}^1$ and $M_{t,\omega}^2$ the individual decision power of the two spouses in the LIC model (12), with $M_{1,\omega}^i = \mu^i$. Then, in the *intra-household allocation stage*, the household chooses the allocation of private goods between spouses by solving the following static problem:

$$U(\bar{Q}_{t,\omega}, \bar{q}_{t,\omega}, \bar{l}_{t,\omega}^1, \bar{l}_{t,\omega}^2, M^1, M^2) = \max_{q_{t,\omega}^1, q_{t,\omega}^2} M_{t,\omega}^1 u^1(\bar{Q}_{t,\omega}, q_{t,\omega}^1, \bar{l}_{t,\omega}^1) + M_{t,\omega}^2 u^2(\bar{Q}_{t,\omega}, q_{t,\omega}^2, \bar{l}_{t,\omega}^2)$$

(13)

subject to

$$q_{t,\omega}^1 + q_{t,\omega}^2 = \bar{q}_{t,\omega}.$$

The solution to this problem provides the household utility function

$U(\bar{Q}_{t,\omega}, \bar{q}_{t,\omega}, \bar{l}_{t,\omega}^1, \bar{l}_{t,\omega}^2, M^1, M^2)$, which measures the welfare generated by the household in period t and state ω if (i) it is endowed with the vector of public goods $\bar{Q}_{t,\omega}$, the vector of private goods $\bar{q}_{t,\omega}$, and leisure $\bar{l}_{t,\omega}^1$ and $\bar{l}_{t,\omega}^2$, and (ii) the two spouses have decision powers equal to $M_{t,\omega}^1$ and $M_{t,\omega}^2$. Note that U is not, in general, a standard utility, since it depends on the Pareto weights $M_{t,\omega}^1$ and $M_{t,\omega}^2$ and on all the variables that affect them.

We can now introduce the resource-allocation stage, which also deals with period t and state ω . The household enters this period and state with an arbitrary amount of resources $\xi_{t,\omega}$ and uses the indirect utility function derived in the last stage $U(\bar{Q}_{t,\omega}, \bar{q}_{t,\omega}, \bar{l}_{t,\omega}^1, \bar{l}_{t,\omega}^2, M^1, M^2)$ to determine the optimal choice of commodities and time to employ in the production of

private and public goods, and the optimal choice of leisure. Formally, the household solves the following static problem:

$$V(\xi_{t,\omega}, p_{t,\omega}, w_{t,\omega}, M_{t,\omega}^1, M_{t,\omega}^2) = \max_{\{x_{a,t,\omega}, d_{a,t,\omega}^i, l_{t,\omega}^i\}} \quad (14)$$

$$U(Q_{t,\omega}, q_{t,\omega}, l_{t,\omega}^1, l_{t,\omega}^2, M_{t,\omega}^1, M_{t,\omega}^2)$$

subject to

$$p'_{t,\omega} \left(\sum_{k=1}^N X_{k,t,\omega} + \sum_{h=1}^n x_{h,t,\omega} \right) + \sum_{i=1}^2 w_{t,\omega}^i \left(l_{t,\omega}^i + \sum_{k=1}^N D_{k,t,\omega}^i + \sum_{h=1}^n d_{h,t,\omega}^i \right) = \xi_{t,\omega},$$

$$Q_{k,t,\omega} = F_k(X_{k,t,\omega}, D_{k,t,\omega}),$$

and

$$q_{h,t,\omega} = f_h(x_{h,t,\omega}, d_{h,t,\omega}),$$

where $V(\xi_{t,\omega}, p_{t,\omega}, w_{t,\omega}, M_{t,\omega}^1, M_{t,\omega}^2)$ is the indirect utility function for this stage, measuring the household welfare if $\xi_{t,\omega}$ resources are allocated to this period and state, prices and wages are $p_{t,\omega}$ and $w_{t,\omega}$, and the individual decision powers are $M_{t,\omega}^1$ and $M_{t,\omega}^2$. Observe that even the *resource-allocation stage* does not boil down to a unitary setting unless the indirect utility function derived in the *intra-household allocation stage*,

$U(\bar{Q}_{t,\omega}, \bar{q}_{t,\omega}, \bar{l}_{t,\omega}^1, \bar{l}_{t,\omega}^2, M^1, M^2)$, is independent of the Pareto weights $M_{t,\omega}^1$ and $M_{t,\omega}^2$ or those weights are constant.

In the *intertemporal stage*, the household deals with the allocation of lifetime resources over time and across states, using the indirect utility function derived in the *resource-allocation stage*. In the absence of the individual participation constraints, the allocation would simply solve the following standard intertemporal problem:

$$\max_{\xi, \xi_{\omega_L}, \xi_{\omega_H}} \quad (15)$$

$$V(\xi, p, \omega, M_1^1, M_1^2) + \sum_{\omega=\omega_L, \omega_H} \beta V(\xi_{\omega}, p_{\omega}, w_{\omega}, M_1^1, M_1^2) \times P(\omega) + g(X, W)$$

subject to

$$\xi + b = Y \text{ and } \xi_{\omega} = Y_{\omega} + Rb$$

for $\omega = \omega_L, \omega_H$.

But the lack of commitment and the corresponding participation constraints imply that we must consider three different alternatives. First, if in each state of nature the solution to problem (15) does not violate the spouses' participation constraints, it is the solution to the general three-stage problem. Second, if for some state of nature the solution to problem (15) violates the participation constraints of both spouses, the household does not generate a positive surplus and it is optimal for the two partners to choose their best outside options. Finally, if for some state of nature, at the solution, the participation constraint of just one spouse is violated, say agent 1's, the household increases that spouse's decision power $M_{2,\omega}^1$ according to equation (11) and solves the following new problem:

$$\max_{\xi, \xi_{\omega_L}, \xi_{\omega_H}} \quad (16)$$

$$V(\xi, p, \omega, M_1^1, M_1^2) + \sum_{\omega = \omega_L, \omega_H} \beta V(\xi_{\omega}, p_{\omega}, w_{\omega}, M_{2,\omega}^1, M_{1,\omega}^2) \times P(\omega) + g(X, W)$$

subject to

$$\xi + b = Y \text{ and } \xi_{\omega} = Y_{\omega} + Rb$$

for $\omega = \omega_L, \omega_H$.

One has then to consider the three cases just described at the new solution.

3.3 Existing Models from a Three-Stage Perspective

One can readily see, from this three-stage formulation, that several standard models of household decisions are special cases of the LIC model. Each class of models implicitly concentrates on one or several stages, and by so doing it accounts for some aspects of household behavior while missing others.

Static Unitary Models—As mentioned in the first part of the survey, historically static unitary models have been the most popular framework to study household decisions. Some example of papers using this type of model are Samuelson (1956); Becker (1962); Becker (1981); Deaton and Muellbauer (1980); Blundell, Pashardes, and Weber (1993); and Lewbel (2001). An interesting feature of the static unitary model is that it corresponds to a special case of the resource-allocation stage of the household decision process. To see this, remember the result discussed in the first part of the survey according to which the static unitary model is equivalent to the static collective model if the two functions characterizing

the intra-household decision power $M_{t,\omega}^1(W)$ and $M_{t,\omega}^2(W)$ do not depend on the vector of variables W and are therefore constant. Under this restriction, for some utility function U^H , the resource-allocation stage can be written in the following form:

$$\max_{\{x_{h,t,\omega}, X_{k,t,\omega}, d_{h,t,\omega}^i, D_{k,t,\omega}^i, l_{t,\omega}^i\}} \quad (17)$$

$$U^H(Q_{t,\omega}, q_{t,\omega}, l_{t,\omega}^1, l_{t,\omega}^2)$$

subject to

$$p'_{t,\omega} \left(\sum_{k=1}^N X_{k,t,\omega} + \sum_{h=1}^n x_{h,t,\omega} \right) + \sum_{i=1}^2 w_{t,\omega}^i \left(l_{t,\omega}^i + \sum_{k=1}^N D_{k,t,\omega}^i + \sum_{h=1}^n d_{h,t,\omega}^i \right) = \xi_{t,\omega},$$

$$Q_{k,t,\omega} = F_k(X_{k,t,\omega}, D_{k,t,\omega}),$$

and

$$q_{h,t,\omega} = f_h(x_{h,t,\omega}, d_{h,t,\omega}).$$

This problem corresponds to the standard unitary model of the household when leisure is considered and some of the consumption goods are produced within the household. As a consequence, under the restriction that the intra-household decision power does not vary, the resource-allocation stage and the static unitary model are equivalent.

Using the resource-allocation stage, we can therefore evaluate which aspects of household decisions the static unitary model can and cannot deal with. This popular model can clearly account for the allocation of resources across different types of consumption goods and for the allocation of time to labor supply, household production, and leisure. But it cannot deal with the aspects of household behavior that are implicit in the intra-household allocation and intertemporal stages. Specifically, the unitary model cannot be used to study the allocation of consumption across household members. Moreover, since it relies on the assumption that the relative decision power is constant, it cannot account for the effect of differences across households in decision power on labor supply, leisure, and household production choices.

Static Collective Models—The static collective models surveyed in the first part of the paper consider a wider set of aspects of household decisions, since they account for all dimensions of household behavior considered in the resource-allocation and intra-household allocation stages. Similarly to the static unitary model, they can be used to study the allocation of household resources across goods and the allocation of time to labor supply, household production, and leisure. Differently from the unitary model, they enable

researchers to examine the allocation of consumption across household members, as implied by the intra-household allocation stage, and to study the effect of differences in decision power across households on consumption, labor supply, household production, and leisure choices. The main limitation of these models is that, by construction, they cannot account for the intertemporal aspects of household behavior.

Intertemporal Unitary Models—The most commonly used intertemporal model to study household decisions is the intertemporal unitary model, as exemplified by papers such as Heathcote, Storesletten, and Violante (2010); Scholz, Seshadri, and Khitatrakun (2006); and Krueger and Perri (2006). Results in the first part of the survey imply that the intertemporal unitary model is equivalent to the intertemporal collective model only if household decisions are unaffected by differences in decision power across households, and by changes in decision power over time. As a consequence, this model can only deal with aspects of household behavior that characterize the resource-allocation and intertemporal stages of the household decision process. It can therefore account for the allocation of income across goods, the allocation of time to labor supply, household production, and leisure, the distribution of resources over periods, the accumulation of human capital, and the intertemporal changes in intra-household specialization. However, if any of these aspects of household decisions are affected by cross-sectional or time-series variation in decision power, the intertemporal unitary model will generally generate misleading implications. Moreover, it cannot be used to study the allocation of consumption across spouses and it is not well suited to the study of household formation or dissolution, particularly when external shocks may or may not result in divorce, depending on the intra-household reallocation that can be triggered.

Intertemporal Collective Models—Two main specifications of the intertemporal collective model have been used to analyze household behavior. The first one is the LIC model described at length in this part of the survey. The second model is the full-commitment intertemporal (FIC) model. Similarly to the LIC model, it accounts for all three stages of the household decision process. The only difference is that it assumes that at the time of household formation, household members can commit to future plans. Formally, the FIC model corresponds to the LIC model when the participation constraints are eliminated and takes the following form:

$$\begin{aligned} \max \mu^1(\Theta) u^1(Q, q^1, l^1) + \mu^2(\Theta) u^2(Q, q^2, l^2) \\ + \beta \left[\sum_{i=1}^2 \mu^i(\Theta) u^i(Q_{\omega_L}, q_{\omega_L}^i, l_{\omega_L}^i) P(\omega_L) + \sum_{i=1}^2 \mu^i(\Theta) u^i(Q_{\omega_H}, q_{\omega_H}^i, l_{\omega_H}^i) P(\omega_H) \right] \end{aligned} \quad (\text{FIC})$$

subject to

$$\begin{aligned}
& p' \left(\sum_{k=1}^N X_k + \sum_{h=1}^n x_h \right) \\
& + \sum_{i=1}^2 w^i \left(l^i + \sum_{k=1}^N D_k^i + \sum_{h=1}^n d_h^i \right) \\
& + b = Y, p'_{\omega} \left(\sum_{k=1}^N X_{k,\omega} + \sum_{h=1}^n x_{h,\omega} \right) \\
& + \sum_{i=1}^2 w_{\omega}^i \left(l_{\omega}^i + \sum_{k=1}^N D_{k,\omega}^i + \sum_{h=1}^n d_{h,\omega}^i \right) \\
& = Y_{\omega} + Rb, \text{ for } \omega = \omega_L, \omega_H;
\end{aligned}$$

$$Q_k = F_k(X_k, D_k)$$

and

$$\sum_i q_h^i = f_h(x_h, d_h)$$

for all k and h ,

$$Q_{k,\omega} = F_k(X_{k,\omega}, D_{k,\omega})$$

and

$$\sum_i q_{h,\omega}^i = f_h(x_{h,\omega}, d_{h,\omega})$$

for all k, h , and $\omega = \omega_L, \omega_H$.

The assumption that the two spouses can commit to an allocation of resources and the consequent absence of participation constraints has two main implications. The first is that the FIC model is a special case of the LIC model: the two models have the same solution in the particular case in which the participation constraints in the LIC model never bind and therefore the household never experiences the need to renegotiate its initial plan. The second implication is that, with commitment, only the intra-household decision power at the time of household formation can affect family behavior. Since household members can commit to an allocation of resources for the future, changes in the outside options are inconsequential. All this implies that only variation across households in decision power can affect consumption, leisure, household production, and saving decisions.

An important question is: which of the two intertemporal collective models is a better characterization of household behavior? The level of commitment required by the FIC is

very high. In that framework, household members cannot renegotiate their initial plan even if, after some years of marriage, a large shock hits one of the two spouses, such as the long-term loss of a well-paid job or the announcement that one of the spouses will receive a large inheritance. While such level of commitment is probably feasible (think, for instance, of religious beliefs prohibiting divorce), whether it can realistically be assumed in general is dubious. Even in places where religious beliefs prohibit divorce, household member can always revert to noncooperation. Therefore, without computational and data constraints, the LIC model should be preferred to the FIC model. If, in the data, renegotiations and intertemporal changes in decision power are irrelevant, the LIC model can easily account for this possibility by setting the intra-household decision power in each period and state of nature equal to the initial one.

Computational issues and data requirements are, however, important considerations that should be taken into account when an intertemporal collective model is chosen. The LIC model is computationally demanding and the data requirements for its identification and estimation are larger. The choice of one model over the other should therefore depend on the environment considered by the researcher. For instance, in the United States, where prenuptial agreements are not common and the cost of divorce is relatively low, the LIC model is probably a better model of household decisions, since renegotiations of plans and intertemporal changes in intra-household decision power should play an important role. Aura (2005), Mazzocco (2007), and Lise and Yamada (2014) provide evidence in favor of this hypothesis. In more traditional environments (such as rural societies in many developing countries), renegotiations may be less frequent, since the cost of divorce is relatively high, threats of ending a marriage are therefore less credible, and noncooperation is less appealing since households members are bound to spend a lifetime together. As a consequence, the FIC model may in some cases represent a good approximation of household behavior. In the end, these are empirical questions, and the crucial issue is whether the two versions can be empirically distinguished. We address this topic in the next section.

Table 1 summarizes the relationship between the three-stage formulation of household behavior and the models described in this subsection.

3.4 Endogenous Outside Options

Earlier in the section, we argued that the outside options $\bar{u}_{\omega}^i(X)$ are generally affected by past decisions. There are two main channels through which outside options can be endogenous. First, household members may choose, at the time of household formation or in later years, to make use of explicit or implicit contracts that will affect the spouses' outside options. There are two sets of contracts that are frequently used by spouses who believe that divorce is their best outside option. The first set is composed by prenuptial agreements, which, by specifying the transfers that should take place between spouses in case of divorce, have a direct effect on the outside options, as discussed in Bayot and Voena (2014). The second set includes financial contracts that improve the welfare of one or both spouses in case of divorce. For example, a family may choose to put the household's residence under the wife's name if the existing divorce law recognizes the household's choice in case of separation. The benefits of using these contracts is that households may be able to alleviate

the limitations to full commitment generated by the existence of the outside options. This is an old idea proposed by the incomplete-contracts literature: even when full commitment cannot be achieved and, hence, renegotiations are unavoidable, the full-commitment optimum can sometimes be achieved or approached by a sequence of short-term changes in bargaining powers; see for instance Rey and Salanie (1996). As mentioned above, existing laws play an important role in determining whether households are able to use these contracts to increase ex ante efficiency. For example, some countries enforce prenuptial agreements, whereas others do not recognize them.

The second channel through which outside options are endogenous is investment. Household members can invest in financial assets and human capital, with both types of investment affecting the individual outside options. To simplify the discussion, in this section we have considered a version of the LIC model in which households can only invest in a risk-free asset and cannot accumulate human capital. The extension to many financial assets, however, is straightforward. Moreover, human-capital investment can be readily incorporated in the LIC model. In that version of the model, spouses enter the relationship with a given stock of human capital and, in each period, choose whether and by how much to increase that stock through education, on-the-job training, and experience. This implies that the stock of human capital in a period generally depends on labor supply in the previous periods.²³

An interesting question is whether the possibility of affecting the outside options brings the household closer to or further away from the level of efficiency that can be achieved with full commitment. In the LIC model, this depends on whether the decisions that affect the outside options are made before or after household formation. If the choices are made before, such as education decisions, it cannot be determined whether the household will move closer to or further away from first best. For instance, if the secondary earner made education decisions taking into account their insurance value in case of divorce, as it is the case in Bronson (2015), these decisions have two effects on household welfare. First, for a constant level of marital surplus, they generally increase the number of renegotiations because the secondary earner is more likely to enter the labor market. Second, they generally increase the marital surplus because the secondary earner is more likely to pursue higher education. As a consequence, the number of renegotiations will decline. Therefore, whether the household will be closer to or further away from ex ante efficiency depends on which effect dominates. If decisions are taken after household formation, the LIC model always predicts that they will increase ex ante efficiency, since the objective of the household is to maximize its welfare. But this does not have to be the case in real life. It is possible that some of the decisions that affect the outside options are made noncooperatively. This is particularly plausible for households that are on the verge of divorce. This is a view advocated by Konrad and Lommerud (2000) among others, who posits that such decisions are taken in a noncooperative fashion.

²³If the stock of human capital depends on labor supply in the previous periods, the three-stage formulation has to be generalized. Specifically, the household must choose labor supply in the intertemporal stage because labor supply decisions have an intertemporal dimension through the accumulation of human capital.

One example that illustrates the difference between the two approaches is a household in which the secondary earner has to choose whether to participate in the labor market in an environment in which the outside options correspond to divorce. In the LIC model, this decision is made cooperatively and the household, after taking into account its effect on the outside options, will choose the alternative that maximizes ex ante efficiency. To further understand how the LIC model deals with this situation, consider a case in which the decision not to participate is detrimental to the secondary earner's welfare because it would excessively reduce her outside options. But the decision to participate would reduce the surplus generated by the family because intra-household specialization is optimal. In this case, the LIC model predicts that the household should use alternative ways of increasing the outside options of the secondary earner (for instance, by putting the household's residence under the secondary earner's name, or writing a new marital contract that favors the secondary earner in case of divorce) and ask the secondary earner not to participate in the labor market.²⁴ Noncooperative models will generally produce a different outcome. They predict that the secondary earner will, in most cases, choose to participate in the labor market even if it is detrimental to the welfare of the household as a whole. Once again, it is an empirical question which framework better captures the effect of this type of decisions on the individual outside options.²⁵

It is also an empirical issue whether endogenous outside options increase or reduce efficiency. It is therefore crucial to investigate the empirical content of the LIC or other models that can account for this aspect of household behavior. Are those models able to generate testable predictions? Do they allow one to empirically evaluate the relevance of their main assumptions? Is it possible to recover their underlying structures, namely preferences and the decision process, from observed behavior? These questions are considered in the next section.

4. Tests, Identification, and Estimation Results

In this section, we first survey the tests used in the literature to evaluate the ability of different classes of models to explain household decisions. We then discuss results that establish whether the structure of a particular model can be recovered using commonly available data. Lastly, we describe some of the papers that have estimated collective household models.

4.1 Tests

The derivation and implementation of testable implications of proposed models is an essential part of economic research, since it enables researchers to assess whether a particular framework represents a good characterization of the data. In this subsection, we will review the most popular tests of household models. For each test, we will first describe

²⁴Of course, the legislative environment may in some cases limit the couple's ability to alter outside options; for instance, divorce law (or courts) may disregard contractual agreements, or may not allow spouses to maintain the property rights of goods put under their name during marriage.

²⁵See Basu (2006) for additional discussions about the effect of spouses' decisions on outside options.

the implication on which it is based. We will then refer to the three-stage formulation to outline which aspect of household behavior is tested.

4.1.1 Income Pooling and Symmetry of the Slutsky Matrix—The first test we consider is the income-pooling hypothesis. Because of its simplicity, it is one of the most popular tests in the literature on household decisions, as shown by the large number of papers that have relied on it. As argued in the first part of the survey, it is based on the implication of the unitary model according to which, after controlling for household income, individual incomes should have no effect on household decisions. The three-stage formulation makes clear that this test is based on the resource-allocation stage and it evaluates whether the allocation of income to different consumption goods is consistent with the unitary model. Any feature of income allocation that is important in the data but missing from the unitary model will generate a rejection of income pooling. This implies that the income-pooling test provides no information on the ability of the static collective model to characterize the data, since it does not evaluate the intra-household allocation stage. Similarly, the income pooling test gives no information on which model is better equipped to characterize the intertemporal allocation of household resources, since it does not evaluate the intertemporal stage. Many papers have tested and rejected the income-pooling hypothesis. For instance Schultz (1990); Thomas (1990); Bourguignon et al. (1993); Phipps and Burton (1998); Fortin and Lacroix (1997); Lundberg, Pollak, and Wales (1997); Attanasio and Lechene (2002); and Ward-Batts (2008) and have provided evidence that the unitary model is not a good characterization of household decisions using this testable implication.

A second prediction of the unitary approach is that the Slutsky matrix derived from household demand functions should be negative semidefinite and symmetric. This property has been tested and repeatedly rejected; see for instance Lewbel (1995); Browning and Chiappori (1998); Dauphin and Fortin (2001); Haag, Hoderlein, and Pendakur (2009); Dauphin et al. (2011); and Kapan (2010). Tests of Slutsky symmetry should be considered with some caution, since they require a host of auxiliary assumptions (from exogeneity of price variations to various forms of separability) to which rejection may always be attributed. However, Browning and Chiappori (1998) and Kapan (2010) run the same test of Slutsky symmetry on three subsamples—single males, single females, and couples. Although the samples are of comparable size, the tests give strikingly different answers; in both papers, Slutsky symmetry is *not* rejected on the two subsamples of singles, but strongly for couples. This suggests that the rejections may not be due to auxiliary assumptions—which are the same for all samples—but may rather reflect systematic behavioral differences between couples and single individuals.

An interesting feature of the first testable implication discussed here, income pooling, is that the unitary model is not the only framework predicting it. As indicated earlier, noncooperative models with public goods typically lead to two possible regimes, one of which (when both partners simultaneously contribute to at least one public good) entails a strong form of income pooling. In that sense, the rejection of income pooling casts doubt on the empirical validity of noncooperative models as well. This claim should, however, be qualified because the actual test is (or should be) more complex than simple income pooling,

since the property only holds for a subset of prices and income bundles, and this subset is endogenously determined. Browning and Lechene (2001) provide a test of this kind, and fail to find evidence of (local) income pooling.

4.1.2 The Static Collective Model—The numerous rejections of the unitary model have induced many researchers who study household behavior to switch to the static collective model. Given its frequent use, it is important to understand whether this generalization of the unitary model is consistent with the available data. Several tests have been proposed. A first test is based on an implication of the collective model that is known as the *proportionality condition*. This restriction is based on the following idea. In a collective model, any distribution factor should affect household decisions only through the relative decision power, which is a one-dimensional function. This feature of the collective model is evident in the individual demand functions specified in equation (7). Using this idea, the proportionality condition can be derived following the next two steps. First, in equation (7), take the derivative of the leisure demand functions for members 1 and 2 with respect to the distribution factor z_j to obtain

$$\frac{\partial \bar{l}^1}{\partial z_j} = \frac{\partial l^1}{\partial \mu} \frac{\partial \mu}{\partial z_j} \text{ and } \frac{\partial \bar{l}^2}{\partial z_j} = \frac{\partial l^2}{\partial \mu} \frac{\partial \mu}{\partial z_j}.$$

Second, divide member 1's derivative by member 2's derivative to derive the following equality:

$$\frac{\partial \bar{l}^1}{\partial z_j} / \frac{\partial \bar{l}^2}{\partial z_j} = \frac{\partial l^1}{\partial \mu} / \frac{\partial l^2}{\partial \mu}. \quad (18)$$

An important characteristic of the ratio on the right-hand side is that it is independent of the distribution factor being considered. This characteristic is known as the proportionality condition and can be tested if one observes leisure and at least two distribution factors. If such variables are observed, one can estimate the ratio on the left-hand side of equation (18) for the two distribution factors and test that the two estimated ratios are statistically equal. The same testable implication can be derived using household private or public consumption $q = q^1 + q^2$, which is also observed. Bourguignon et al. (1993); Browning et al. (1994); Thomas, Contreras, and Frankenberg (1999); Browning and Chiappori (1998); Attanasio and Lechene (2002); Bobonis (2009); Bourguignon, Browning, and Chiappori (2009); and Attanasio and Lechene (2014) have tested the proportionality condition using different data sets without being able to reject it.

Three remarks related to the proportionality conditions are in order. First, the restriction that the distribution factors affect household decisions only through the one-dimensional function μ is fundamental for the derivation of the test. This result has an important consequence. Any model of a two-member household that makes the assumption of efficient decisions has as an implication that the distribution factors enter the decision process only through the relative Pareto weight, which is a one-dimensional function. As a result, the

proportionality condition is a testable restriction that applies to any model in that class. The three-stage formulation makes this clear. The proportionality condition tests the intra-household allocation stage of the household decision process. It therefore evaluates whether the allocation of resources across household members is consistent with a static definition of efficiency. Secondly, from a more pragmatic perspective, the proportionality test requires the distribution factors under consideration to be *exogenous* in the econometric sense. In particular, they must be uncorrelated with preferences, otherwise their effect may simply capture some unobservable heterogeneity in taste. In practice, the most convincing tests use either explicit randomizations, as in Attanasio and Lechene (2014), or aggregate variables that operate at the market level, such as sex ratios or laws governing divorce, as in Chiappori, Fortin, and Lacroix (2002). Lastly, observe that, similarly to the income-pooling test, the proportionality condition provides no information on the ability of different models to characterize the intertemporal allocation of resources observed in the data. Any intertemporal model that assumes within-period efficiency, such as the LIC or FIC model, is consistent with the proportionality condition.

A second test that has been used to evaluate collective models is the symmetry plus rank one of the Slutsky matrix. This testable implication is based on an idea that is similar to the one used to derive the proportionality condition. In a unitary model, prices and household income affect family decisions only through the budget constraint. A consequence of this is that the Slutsky matrix derived from a unitary model is symmetric. In a collective model, however, prices and income have an additional effect on household decisions through the intra-household decision power. As indicated earlier, in a two-member family the intra-household decision power is fully summarized by a one-dimensional function, namely the relative Pareto weight. Therefore, in a collective model, the Slutsky matrix is the sum of a symmetric matrix, which captures the effect of prices and income through the budget constraint, and a matrix of rank one, which accounts for the effect of prices and income on the relative decision power.²⁶ Browning and Chiappori (1998) and Kapan (2010) have used this restriction to test the collective model. Both papers find that symmetry is clearly rejected for couples, but symmetry plus rank one is not. Dauphin and Fortin (2001) and Dauphin et

²⁶This result can be proved mathematically using the collective demand functions derived in equation (7). To simplify the derivation, we only consider private consumption goods and we include wages in the vector of prices p . To derive the generic element of the Slutsky matrix for good i and price j , s_{ij} , in terms of variables that are observed, we consider the household consumption of good i , $q_i = q_i^1 + q_i^2$, and denote with q_i^* its optimal level. The generic element s_{ij} is equal to the derivative of the Marshallian demand function with respect to p_j plus the derivative with respect to income times the quantity demanded of good i . Taking derivatives with respect to prices and income of both sides of equation (7), the generic element of the Slutsky matrix can be written as follows:

$$s_{ij} = \frac{\partial q_i}{\partial p_j} + \frac{\partial q_i}{\partial \mu} \frac{\partial \mu}{\partial p_j} + \left(\frac{\partial q_i}{\partial Y} + \frac{\partial q_i}{\partial \mu} \frac{\partial \mu}{\partial Y} \right) q_i^*. \quad (19)$$

Rearranging terms, s_{ij} can be rewritten in the following form:

$$s_{ij} = \frac{\partial q_i}{\partial p_j} + \frac{\partial q_i}{\partial Y} q_i^* + \frac{\partial q_i}{\partial \mu} \left(\frac{\partial \mu}{\partial p_j} + \frac{\partial \mu}{\partial Y} \right) q_i^*.$$

al. (2011) present tests of the “Slutsky symmetry plus rank k ” property that characterizes households with $(k + 1)$ members, and find that the collective approach is not rejected.

The three-stage formulation of household behavior offers an alternative way to provide the intuition behind the symmetry plus rank one condition. It makes clear that the symmetry plus rank one condition relies on the resource-allocation stage and intra-household allocation stage of the household decision process. Without the intra-household allocation stage, the resource-allocation stage implies that the allocation of within-period resources to different consumption goods should generate a symmetric Slutsky matrix. The allocation of consumption across household members generated by the intra-household allocation stage adds a new matrix to the standard Slutsky matrix. This new matrix must be of rank one since, under efficiency, consumption is distributed to the two spouses according to their relative decision power. Similarly to the previous two tests, the symmetry plus rank one condition does not evaluate the ability of models to rationalize the inter-temporal allocation of resources observed in the data since it has no relationship with the intertemporal stage.

Cherchye, De Rock, and Vermeulen (2007, 2009, and 2011) have devised an alternative test to evaluate models of the households that assume static efficiency. The test is derived by extending the generalized axiom of revealed preference (GARP) to a collective framework. Similarly to the symmetry plus rank one test, the null hypothesis is that, in a given period, households make efficient decisions and it evaluates simultaneously the resource-allocation stage and intra-household allocation stage. It provides no information on intertemporal models. The main advantage of the GARP test is that it is nonparametric and therefore its outcome does not rely on specific functional forms chosen for its implementation. Using this nonparametric test and a variety of data sets, Cherchye, De Rock, and Vermeulen could not reject the null hypothesis of Pareto efficiency.

Udry (1996) has proposed an implication based on household production that enables one to test the static unitary model and the static collective model. The test evaluates whether decisions are consistent with the resource-allocation stage of the household decision process and is based on the following idea. If (i) a household engages simultaneously in consumption and production and (ii) household production is simultaneously carried out for different goods, the unitary and collective models imply that the inputs to household production should be allocated efficiently across the production of these goods. As a consequence, within the household, variation across inputs should only be a function of variation in the characteristics of the goods being produced. If other variables, such as individual incomes or education, affect the allocation of inputs, one or more important aspects of household behavior are missing from the unitary or collective model. Udry finds evidence against his implication, therefore rejecting the static unitary and collective models.

The first part $\frac{\partial q_i}{\partial p} + \frac{\partial q_i}{\partial Y} q_i^*$ is the generic element of the standard Slutsky matrix which is symmetric. The second part $\frac{\partial q_i}{\partial \mu} \left(\frac{\partial \mu}{\partial p} + \frac{\partial \mu}{\partial Y} \right) q_i^*$ is the addition generated by the collective model. It produces a matrix of rank one because $\frac{\partial q_i}{\partial \mu}$ as well as $\frac{\partial \mu}{\partial p} + \frac{\partial \mu}{\partial Y}$ form a vector and a matrix obtained from the product of two vectors has always rank one.

This test only focuses on static aspects of household decisions and does not provide information on the way households allocate resources intertemporally.

4.1.3 The Intertemporal Collective Models—We will now describe tests that evaluate the intertemporal allocation of resources. The most common and popular test is the excess sensitivity test, which is based on household Euler equations. The test focuses exclusively on the intertemporal stage of the household decision process. The null hypothesis is therefore that the intertemporal unitary model correctly describes the data. Hall (1978), Sargent (1978), Flavin (1981), Hall and Mishkin (1982), and Zeldes (1989) were among the first to implement this test, which is based on the following idea. Under the null hypothesis, the household should choose current and future consumption according to the Euler equations using all the information available at the time of the decision. As a consequence, the difference between the current marginal utility of consumption and next period expected marginal utility of consumption should be independent of variables that are known to the household at the time of the decision. Many papers have tested this implication of the intertemporal unitary model and regularly rejected it. Since this test does not evaluate the intra-household allocation stage, it provides no information on inter-temporal collective models of the household. There is, however, one paper whose results suggest that a possible explanation for the rejection of the excess sensitivity test is that the intertemporal allocation of resources is affected by differences across households and over time in intra-household decision power. Attanasio and Weber (1995) find that if one controls for a sufficient number of preference shocks, such as age of the household members and number of children, and for a sufficient number of variables that capture possible non-separabilities between consumption and leisure, such as the labor-force participation and wage of the wife, the household Euler equations are not rejected. This result leads us to ask the question, why do those variables appear in the household Euler equations? Is it because they correctly capture preferences shocks and potential non-separabilities, or because they proxy for the allocation of household resources across spouses? The test we discuss next attempts to answer this question.

Mazzocco (2007) proposes a test that evaluates the three intertemporal models of household behavior described in this survey: the intertemporal unitary model, the FIC model, and the LIC model. The test is an attempt to assess the importance of the intertemporal stage and of the intra-household allocation stage in explaining observed data and it is based on the following idea. As mentioned above, under the assumption that the unitary model is the correct characterization of the data, no variable known at the time of household decisions should affect the difference between the current marginal utility of consumption and next-period expected marginal utility of consumption. Under the alternative that one of the intertemporal collective models is correct, however, any variable that has an effect on the intra-household decision power should influence that difference. To understand why, consider two households in which the two husbands have identical risk aversion, the two wives have identical risk aversion, and the wives' risk aversion is greater than the husbands' risk aversion. The two households are identical in any other dimension, except that the wife in the first household has more decision power than the wife in the second household because she has a college degree instead of the high-school diploma of the other wife.

Generally, the first household will be more risk averse than the second because a higher weight is assigned to the wife's preferences. As a consequence, the first household assigns more value to consumption smoothing and has therefore a flatter consumption path. The difference in consumption paths will be detected in the estimation of the Euler equations and will be explained using the only variable that can rationalize it, education, and variables that are correlated with it, such as wages and labor-force participation. The reduced-form result is therefore that education, wages, and labor-force participation will have an effect on the household Euler equations even if they are known at the time that decisions are made. We can therefore conclude that if one of the variables that affect the intra-household decision power has a significant effect on the Euler equations, the intertemporal unitary model can be rejected. But this outcome provides no information about the FIC and LIC models.

To distinguish between the FIC and LIC models and therefore test the assumption of commitment, Mazzocco (2007) exploits the following difference between the two models. In the FIC model, only the intra-household decision power at the time of household formation can affect household decisions, and therefore the Euler equations. The LIC model instead predicts that the intra-household decision power in each period can influence household behavior, hence, Euler equations. To better understand the different implications of the two models, consider the previous example in which the two wives have different levels of education at the time of marriage. Suppose in addition that, after five years of marriage, because of a significant increase in college grants that was not expected at the time of marriage, the wife with lower education decides to attend college and subsequently graduates. In the FIC model, only the difference in education at the time of household formation can have an effect on the consumption path chosen by the two households. The change in education after five years cannot modify the consumption path because it was not predicted at the time of marriage, and therefore had no effect on the initial decision power. The LIC model has a different prediction. The difference in education at the time of marriage, as well as the unexpected change in education after five years, should affect the Euler equations. A consequence of this result is that in the FIC model, only variation across households in variables entering the decisions power can influence the Euler equations. In the LIC model, both variation across households and over time in those variables can have an effect on the Euler equations. Intuitively, this implication can be tested by verifying whether in the data, the consumption path of a given household changes over time in response to variation in variables that are known at the time the Euler equations were computed but were unexpected before that period. If it does, this provides evidence that intertemporal changes in decision power affect the household allocation of resources. The FIC model can therefore be rejected, which is the case using US data.

Lise and Yamada (2014) use an idea similar to the ones employed in Mazzocco (2007) to test the FIC against the LIC model. The paper estimates a LIC model with private consumption, public consumption, labor supply, and household production using data from the Japanese Panel Survey of Consumers (JPSC). The estimated model enables the authors to evaluate the amount of heterogeneity across households in intra-household decision power at the time of marriage and the amount of heterogeneity generated by the arrival of information that was not predicted at the time of marriage. The paper finds that differences across households in decision power at the time of marriage are quantitatively large and are

the most important factor in generating different outcomes across households. The arrival of news unexpected at the time of marriage has significant but smaller effects on intra-household decision power. Moreover, remarkably, those effects can be detected only if the news is sufficiently big. These results suggest that renegotiations take place, but only if the shocks experienced by households are sufficiently large. The authors therefore conclude that the FIC model can be rejected, whereas the LIC model is consistent with the data.

Table 2 summarizes the relationship between the three-stage formulation of the household decision process and the testable implications described in this section.

4.2 Identification Results

In this subsection, we describe results related to the identification of household models. We will consider first the static unitary and static collective models with and without household production. We will then discuss identification results for intertemporal models.

4.2.1 Static Models—We start with the unitary model. Suppose that one can observe the solution to program (UM). With this information, can household preferences and production functions be identified? The answer to the question is negative, and this, arguably, is a serious weakness of the unitary model. If some of the goods are produced within the household, household preferences cannot be separately identified from the household production functions—a point initially made by Pollak and Wachter (1975). To provide some insight on this result, it is instructive to substitute out private and public consumption from the utility function $U^H(Q, q, l^1, l^2)$ using the production function constraints $Q_k = F_k(X_k, D_k)$ and $q_h^i = f_h^i(x_h, d_h)$ and define the function V^H as follows

$$V^H(x, X, \delta^1, \delta^2, l^1, l^2) = \max_{D, d} U^H(F(X, D), f(x, d), l^1, l^2)$$

subject to

$$\sum_k D_k^i + \sum_h d_h^i = \delta^i, \quad i=1, 2,$$

where $F(X, D) = (F_1(X_1, D_1), \dots, F_N(X_N, D_N))$, $f(x, d) = (f_1(x_1, d_1), \dots, f_n(x_n, d_n))$, and δ^i de-notes the total time allocated by i to domestic production. In words, $V^H(x, X, \delta^1, \delta^2, l^1, l^2)$ denotes the household utility level reached by optimally allocating the domestic production time of each spouse δ^i between the various activities. Then, the household's decisions solve the following problem:

$$\max_{(x, X, \delta^1, \delta^2, l^1, l^2)} V^H(x, X, \delta^1, \delta^2, l^1, l^2) \quad (21)$$

subject to

$$p' \left(\sum_{k=1}^N X_k + \sum_{h=1}^n x_h \right) + \sum_{i=1}^2 \omega^i (l^i + \delta^i) = Y.$$

That is, the household uses the function V^H to choose the optimal quantity of goods to purchase in the market and the time to devote to leisure and household production. That problems (21) and (UM) have the same solution is a direct consequence of Hicks's composite-good theorem, given that the marginal price (namely the wage) is the same for all components of D^i and d^i ($i = 1, 2$).²⁷

We can now explain why household's preferences cannot be separately identified from the household production functions. Since in the data, the amount of consumption goods produced within the household is generally not observed, only the indirect utility function $V^H(X, x, \delta^1, \delta^2, l^1, l^2)$ can be recovered using available data on goods purchased in the market and time allocated to leisure and household production. Moreover, V^H contains joint information on preferences and production functions, which implies that household's preferences and production functions cannot be separately recovered. Technically, there exists a continuum of different preference and production functions that generate the same utility V^H ; these various combinations are empirically indistinguishable.²⁸

This weakness probably explains why, despite its theoretical appeal, the notion of intrahousehold production has exclusively been used, within a unitary framework, for situations in which the output of domestic production was directly observed. As we shall see later, this criticism does not apply to collective approaches, in which preferences and production technologies can typically be independently identified.

We now move to the collective framework. The main difficulty in the identification and estimation of collective models is that some household decisions are generally not observed. Most data sets contain information on total household consumption, leisure, prices, wages, and individual income, but they rarely collect data on the individual consumption of private goods. For this reason, in what follows, we assume that the following variables are observed: household consumption of all (public and private) commodities, leisure of each member, prices, wages, and income. We then consider the following question: when is this information sufficient to recover the underlying structure of a given model, i.e., individual preferences, production functions, and the intra-household decision power?

We first answer this question for a static collective model with egotistic preferences and without household production. In that context, the answer is provided in Chiappori and Ekeland (2009). They show that individual preferences and the intra-household decision power can be recovered if, for each household member, the following exclusion restriction is satisfied: there is one private good that the household member does not consume and, hence, that does not enter his or her egotistic utility function. For instance, this exclusion restriction

²⁷Becker's (1965) initial paper on the allocation of time was allowing for different "prices" of time, depending on the type of work performed. However, such a general framework raises serious identification issues. See, for instance, Chiappori and Lewbel (2015).

²⁸In a unitary model without household production, the household's preferences, but not the individual preferences, can be identified using standard results.

is satisfied if individual leisure is observed. In this case, with egotistic preferences, one private good for each member—the spouse’s leisure—does not enter her or his utility function.

Some remarks about this identification result are in order. First, the individual preferences are identified in the form of the indirect utility of each member, which can be obtained by replacing the demand functions (7) in the individual utility functions and can be written in the following form:

$$\bar{v}^i(p, w, Y, z) = v^i(p, w, Y, \mu(p, w, Y, z)). \quad (22)$$

From a welfare perspective, only the individual indirect utility functions are needed, since they fully characterize the utility each spouse can achieve given the exogenous variables. The identification of individual direct utilities is more problematic. The reason for this is that the identification of individual preferences relies on the identification of the sharing rule introduced in section 2.4. Chiappori and Ekeland (2009) have shown that the sharing rule is generally only identified up to some function of the prices of nonexclusive private goods. For each choice of this function, the direct-utility functions can be fully recovered. But without knowing that function, there is no hope of recovering the individual utilities. This is not a concern in the identification of the indirect utility functions, since they include the particular function of prices that characterizes the data and there is no need of separately identifying it. The second remark is that the identification result just described does not require the observation of distribution factors. If distribution factors are observed, Chiappori and Ekeland (2009) have shown that the identification result is stronger. In that case, identification of preferences and decision power can be achieved if one of the following two conditions is satisfied: the researcher observes one excluded good, instead of one for each spouse; or she observes a good that is assignable, where a good is assignable when it is consumed by both members and the consumption of each member is independently observed. As a third remark, observe that the identification result is only generic in the sense that it may fail in particular cases. The most interesting case in which it fails was discussed in section 2.2, where we argued that the unitary model does not allow the identification of individual preferences.

When household production is added to the static collective model, a key issue is whether the goods that are domestically produced can be bought and sold in the market. Chiappori (1997) shows that, if the goods are marketable, the entire structure of the collective model can be identified, including the production functions. When produced commodities cannot be bought and sold in the market, however, identification is more problematic, because their price is now endogenous and household-specific. Apps and Rees (1997) and Chiappori (1997) show that, in that case, the structure is not identified in general; specifically, Chiappori (1997) demonstrates that if the production functions exhibit constant return to scale, then the sharing rule can be identified up to an additive function of individual wages.

These limitations to the “pure” identifiability properties of the collective models can, however, be overcome in different ways. First, they are specific to *local* identification; i.e.,

they consider the identification issue when the demand functions are only observed in the neighborhood of a given point. In particular, they ignore such global restrictions as nonnegativity, which may often provide more identification power. For instance, in the model without household production, while the sharing rule is locally identified up to some additive function of the prices of the nonexclusive private goods, imposing nonnegativity of individual expenditures for small values of total income pins down that function in general (see Browning, Chiappori, and Weiss 2014, chapter 5 for a detailed discussion). Along similar lines, approaches based on revealed preferences, which are global by nature, are typically able to compute narrow bounds on the sharing rule, even in very general frameworks, as shown for instance in Cherchye et al. (2014). A second way of improving the identification properties of the collective model is to assume that some aspects of individual preferences are unchanged by marriage. It is then possible to use data on singles to complement the “pure,” collective identification. For instance, a series of papers recently published in the *Review of Economics of the Household* (Bargain et al. 2006; Beninger et al. 2006; Myck et al. 2006; Vermeulen et al. 2006) consider a collective model of labor supply with public and private consumption, and assume that the marginal rate of substitution between leisure and private consumption is independent of marital status. Then the model is fully identified. Similarly, Browning, Chiappori, and Lewbel (2013) investigate a general framework in which agents, when they get married, keep the same preferences but can access a different, and generally more productive, consumption technology. The idea behind this assumption is that married households can exploit economies of scale in consumption and be more productive in the production of domestic goods because of intra-household specialization. In practice, while the basic rates of substitution between consumed commodities remains unaffected by marriage or cohabitation, the relationship between purchases and consumption is not. The demand structure of married individuals is therefore different from that of singles. Again, under these assumptions the entire model, including household production, can be identified.

Explicit restrictions on some functions can also allow full identification. Dunbar, Lewbel, and Pendakur (2013) explore the so-called “independence of scale” assumption, whereby the fraction of resources going to the wife does not depend on nonlabor income, although it may vary with prices. Under independence of scale and the additional assumption that the parents’ preferences regarding their own consumption do not depend on the number of children, they obtain full identification. In Dunbar, Lewbel, and Pendakur (2014), the same authors show an even stronger result. They find that, in the presence of a distribution factor, independence of scale guarantees full identification even without assumptions on preferences.

Finally, recent approaches have adopted an equilibrium perspective to deal with identification issues. Following an initial insight by Becker (1981), the objective is to model the genesis of individual decision power as stemming from equilibrium conditions in the marriage market and identify the corresponding framework. Various versions have been explored, based either on frictionless matching as in Choo and Siow (2006) and Chiappori, Salanie, and Weiss (2017) or on search models as in Jacquemet and Robin (2013) and Gousse (2014). Some of these papers consider exclusively matching patterns, whereas others analyze matching patterns and induced behavior as in Chiappori, Costa Dias, and Meghir

(forthcoming). The reader is referred to Browning, Chiappori, and Weiss (2014), chapters 7 and 8, for a general presentation.

4.2.2 Intertemporal Collective Models—We will now discuss identification of intertemporal models of the household. The three-stage formulation of the household decision process makes clear that the identification results derived for static collective models can be used to recover the part of the structure of intertemporal collective models that appears in the resource-allocation and intra-household allocation stages. Specifically, they can be used to identify the static part of the individual utility functions, which controls the individual preferences over allocations across a given set of goods. For instance, in a standard Cobb–Douglas utility function, $u(c, l) = (c^\sigma l^{1-\sigma})^{1-\gamma} (1-\gamma)$, the static part of individual preferences is represented by σ . The static results can also be used to identify the decision power at the time of household formation, how it varies across households, and the household production functions.

There are two elements of the intertemporal collective models that cannot be identified using results derived in a static framework. The first element is the intertemporal component of individual preferences, which captures the willingness to transfer resources across periods. In the Cobb–Douglas utility function introduced above, the intertemporal part is summarized by the parameter γ . The second element that cannot be identified using the static results is how the decision power changes over time with the arrival of new information. Casanova and Mazzocco (2014) study the identification of the first intertemporal element. They consider a FIC model and show that, in addition to the static features of household decisions, it is also possible to identify the intertemporal part of individual preferences. The idea used in the paper is to rely on individual Euler equations, one set of Euler equations for each household member, rather than the standard household Euler equations. Individual Euler equations provide the proper variation to identify the intertemporal part of individual preferences, since they contain information on how individual consumption varies in response to income or other types of shocks. The main limitation of the individual Euler equations is that they depend on individual consumption that is generally not observed. To overcome this problem, the paper employs the static relationship between the marginal utility of individual consumption and the marginal utility of leisure to derive individual consumption as a function of individual leisure and the individual wage rate, which are both observed. This function is then used jointly with the individual Euler equations to identify the intertemporal part of the individual utility functions. The paper studies identification in a FIC model of the household. As a consequence, it is not possible to analyze how the decision power changes over time. We are not aware of papers that have examined identification in the more general LIC model. This is an important topic that we believe should be considered in future research.

4.3 Estimation Results

We will now discuss some of the papers that have estimated collective models with the objective of describing in which circumstances a model that recognizes that households are generally composed of several decision makers is a better choice than a standard unitary model. To do that, for each paper we will highlight the gains of using a non-unitary

approach. To keep the survey of a reasonable size, we only review a small sample of the existing papers.

One of the first attempts to estimate a collective model of the household is the work by Browning et al. (1994). The main objective of the paper is to evaluate the effect of household wealth and individual incomes on the intra-household allocation of resources. A unitary model would not allow an economist to perform this analysis since, as argued earlier, it predicts that only household income should have an effect on family choices. The main finding of the estimation of the static collective model is that individual income has a significant effect on the amount of resources received by a spouse and that the effect depends on household wealth. For instance, in a poor household in which the wife's share of income is only 25 percent of the total household income, she receives 42 percent of total expenditure. At the other extreme, in a wealthy household in which she receives 75 percent of income, she has a 58 percent share in total expenditure.

Dunbar, Lewbel, and Pendakur (2013) use Malawi data to estimate the share of resources allocated to children and evaluate how this share depends on variables such as family structure, family size, and income. Earlier in this survey, we presented results indicating that, in a unitary model, the share of resources allocated to different household members cannot be identified. Dunbar, Lewbel, and Pendakur (2013), therefore, use a static collective model to perform their exercise. The paper documents two main findings. First, children receive a significant share of resources, approximately 20 percent for the first child. Second, the total share of expenditures going to children slightly increases with the number of children. As a consequence, the per child share decreases sharply. The authors also find that mothers are willing to sacrifice a fraction of their resources to their children and that the father's share does not respond to the number of children. They conclude that policies aimed at reducing extreme poverty cannot discard issues related to intra-household allocation. In particular, an analysis solely based on equivalence scales would be severely biased: it would grossly underestimate women's and children's poverty, especially in large households.

Gemici (2011) uses a collective model of the household to understand the main determinants of migration decisions in the United States. The paper accounts for the following two crucial aspects of migration choices. First, migration decisions are affected by job offers from other locations. Second, about two third of US households are couples, which implies that migration decisions are, for the most part, taken jointly by two spouses. The first part of the paper provides descriptive evidence that document the following facts. First, a couple has the lowest probability of moving when both spouses are working in their current location (1.63 percent) and the highest probability when both spouses are unemployed (5.26 percent).

Second, conditional on not working at t , married men are three times more likely than married women to become employed at $t + 1$ in periods when they do not move. If they move, however, males are approximately six times more likely to become employed, relative to their wives. Third, the divorce rate in periods that do not involve a change of location is 1.83 percent, whereas it is 4.54 percent in periods that involve a change of location. Lastly, married men who moved at least once have higher wages, whereas for married women migration is associated with lower wages. These facts make clear that in many instances,

migration decisions of couples benefit one spouse, but reduce the welfare and economic conditions of the partner. A collective model is better able to capture these features of the data, especially if one is interested in divorce decisions, since it allows for individual preferences. For this reason, in the second part of the paper Gemici (2011) develops, estimates, and simulates an intertemporal collective model that includes migration decisions. The simulation results indicate that it is important to account for the joint decisions of household members to rationalize the observed migration patterns. Specifically, Gemici finds that without family ties, 25 percent of men and 23 percent of women migrate, whereas only 18 percent of individuals move when married. Moreover, migration increases the average wages of men without family ties by 10 percent more than for married men. Similarly, the average wages of women increase by about 3 percent more than for married women after a move.

Voena (2015) studies the effect on labor-force participation and savings decisions of two types of divorce law changes: changes from mutual consent to unilateral divorce; changes from title based property, in which the property rights on wealth items are clearly assigned, to community property or equitable distribution, where wealth items are divided equally or distributed equitably at the time of divorce. The first part of the paper provides evidence that in states with unilateral divorce laws, a switch from a title-based regime to a community-property or equitable-distribution regime generates a statistically significant increase in household savings. She also finds that a change from mutual consent to unilateral divorce produces a decline in the wife's labor-force participation. Intertemporal collective models are well-suited to explaining the effects of the described changes in divorce laws for two reasons. First, they model household members using individual utility functions. It is therefore possible to follow individuals before and after a divorce and study when it is optimal to divorce. Second, the change from mutual consent to unilateral divorce can be understood as a change from an environment where households make decisions with commitment to an environment in which there is no commitment. The second part of the paper takes advantage of these two features of the intertemporal collective models to rationalize those findings. Specifically, Voena estimates an intertemporal collective model of the household in which married couples can commit to future plans in periods with mutual consent divorce, but they cannot commit after the divorce law has changed to unilateral divorce. The estimated model can match the patterns observed in the data. After the introduction of unilateral divorce, in the LIC model, the spouse with lower initial decision power, typically the wife, experiences an increase in decision power in all cases in which she is entitled to a larger fraction of wealth at divorce. As a consequence, in all those cases, this spouse reduces her labor supply and enjoys more leisure. The increase in savings in community-property states is generated in the model by the desire of husbands to self-insure against the potential loss of half of their assets to their spouses in case of divorce.

Bronson (2015) provides an answer to the following two questions. Why do women today attend college at higher rates than men, whereas in the 1970s men used to attend college at higher rates? And given the high college attendance rate of women today, why do they select disproportionately into lower-paying majors? The first part of the paper uses descriptive methods to document which factors are behind the reversal of the gender gap in college attendance and the persistence in gender differences in choices of majors. The main finding

is obtained using quasi-experimental variation generated by divorce law reforms, which indicates that changes in the probability of divorce contributed in a significant way to producing the observed patterns in graduation rates and major choices. The objective of the second part of the paper is to quantify the importance of this and other channels and to evaluate policies that have been proposed as ways of reducing the earnings gap between women and men. To quantify the effect of future probabilities of divorce on educational decisions when young, one needs a model with the following two features. First, it must be able to follow an individual over the following stages of life: singlehood, marriage, and divorce. Second, during marriage, spouses should make joint decisions to allow for intra-household specialization. The intertemporal unitary model does not possess these two features since, at the time of marriage, only one utility function is used for the entire household. It is therefore difficult to study divorce decisions and their effect on earlier choices. Bronson (2015) chooses to estimate and simulate the LIC model, which has the advantage over the FIC model of allowing her to evaluate the significance of changes in intra-household decision power on educational decisions. Using the estimated model and historical counterfactuals, Bronson (2015) finds that about half of the convergence in the gender gap in college graduation rates observed in the 1970s and 1980s was produced by the increase in the value of insurance that a college degree gives to women if they experience a divorce. The rest of the convergence is explained, for the most part, by changes in wages. The model is then used to evaluate the effect on education, labor supply, and marriage-related decisions of policies that change the financial return of different majors and family policies like paid maternity leave and part-time work entitlements. The results suggest that these policies, currently considered by policy makers and discussed in the media, have important unintended effects that can be measured only by estimating a model of the collective type.

Mazzocco, Ruiz, and Yamaguchi (2014b) examine the relationship between marital status, labor supply, and household production. The paper is motivated by evidence documented in Mazzocco, Ruiz, and Yamaguchi (2014a) using the PSID, which indicates that marital decisions affect labor supply and household production choices. As it may be expected, they find that households display large differences by marital status in labor-force participation, labor supply, and time devoted to household production. Married men work, on average, about 200 hours more than single men, who work 200 hours more than single women, whose labor hours exceed the labor hours of married women by about 200 hours. The same ranking applies to labor-force participation. The ranking for household production is reversed. Married women spend the highest number of hours in household production at 1,287 annual hours. They are followed by single women with 604 annual hours, single men with 372 hours, and married men with 366 hours. A more surprising finding is that the differences between married and single individuals in labor-force participation, labor supply, and household production do not arise suddenly at the time of marriage, but rather emerge gradually over time. To document this pattern, they study the evolution of these variables for women and men who experience a marriage or a divorce. They show that, before women choose to marry, they supply on average the same amount of labor hours as the average single woman. Starting from two years before marriage, however, they begin to gradually reduce their labor hours. At the time of marriage, their labor supply is about 200 hours lower

than the average single woman and about 300 hours higher than the average married woman. The labor supply of women who experience this transition into marriage continues to decline until after 4–5 years of marriage it reaches the level of the average married woman. The labor supply of men who enter marriage displays a similar transition, but with a positive trend. The transition of married couples into divorce is also characterized by significant and gradual changes in labor supply and time spent in household production. The objective of Mazzocco, Ruiz, and Yamaguchi's (2014b) paper is to understand why intra-household specialization changes gradually after marriage and before divorce, instead of being set at an optimal level at the time of marriage. Similarly to other papers described in this section, an intertemporal collective model is better suited to achieve that objective than a unitary model for two reasons. First, it enables one to follow an individual through his three main stages of life—singlehood, marriage, and potentially divorce—since that individual does not lose his identity/preferences when he chooses to join a household. Second, in that model, a married individual makes joint decisions with his spouse, which allows one to analyze intra-household specialization and how it evolves over time as a function of the variables that affect household choices. Mazzocco, Ruiz, and Yamaguchi (2014b) chose to estimate an LIC model instead of an FIC model because it enables them to evaluate which fraction of the gradual change in intra-household specialization can be explained by lack of commitment and the corresponding variation in individual decision power. Their results suggest that there are three main factors that determine the gradual change in intra-household specialization: the birth of a child, the accumulation of human capital, and the lack of commitment. Whereas the birth of a child and the accumulation of human capital have most of the effect after household formation, the lack of commitment plays a significant role after household formation, as well as before divorce. These findings suggest that it is important to account for lack of commitment among household members to evaluate policies aimed at influencing family decisions.

5. Back to Policies

At the end of this survey, we can go back to the policy issues evoked in the introduction and analyze the ability of different models to evaluate them.

Cash Transfers

The cash-transfer policy has two straightforward effects. The first effect is similar for married couples and singles: in both cases, it relaxes the budget constraint by increasing the amount of resources available to them. The second effect, however, is specific to married couples. If the benefit is paid to the wife, it may increase the value of her best outside option, especially if it would still be paid to her in case of separation. This, in turn, may impact her “power” within the relationship, ultimately affecting household's expenditures and investments. The static unitary model can only measure the effect that the cash transfer has on the budget constraint; it cannot account for the additional effect on the intra-household decision power. Since increasing the decision power of women is one of the stated objectives of *Progres-a*-type programs, one can expect that the unitary model is not well-equipped to evaluate cash-transfer policies. This intuition is largely confirmed by the empirical analysis. For instance, Attanasio and Lechene (2014) find that the additional income provided by

Progresa is not spent in the same way as income coming from another source. In particular, the general pattern that wealthier people spend a smaller fraction of their income on food has been regularly observed in Mexican data, as well as in many other countries. But it is violated in the case of “*Progresa* money.”

Static collective models, on the other hand, are perfectly designed to evaluate a one-time unconditional cash transfer. The effect of the transfer on the intra-household decision power can be accounted for by allowing the individual powers that enter the resource- allocation stage and intra-household allocation stage to change. In particular, the observation that *Progresa* money is spent in a different way can be explained by the fact that it is paid to the wife, which may have the effect of improving her weight in the decision process. Importantly, this interpretation generates testable predictions. According to the collective model, if changes in the Pareto weights are a driving force, any other distribution factor should have a similar effect on these weights. Similar, here, is taken in a very specific sense: behavior should satisfy the proportionality condition discussed above.²⁹ This property has been tested by Attanasio and Lechene (2014), who did not reject it.

The LIC model considers all three-stages of household decisions and allows for changes in intra-household decision power. It can therefore account for both effects of the cash-transfer program. It is clearly able to capture the first effect through a modification of the joint budget constraints. It is also well-equipped to account for and quantify the second effect. This can be understood by observing that in the LIC model, the cash transfer produces two adjustments: it increases the value of being married for both spouses and it improves the best outside option of the spouse who receives the transfer. If the increase in outside option is sufficiently large, relative to the increase in the individual value of being married, the participation constraint of this spouse will bind and her decision power will increase. This discussion also clarifies that not all cash transfers will modify the intra-household decision power. A transfer has this effect only if it is significant enough to make one of the spouses willing to take the outside option. The fact that the LIC model predicts that cash-transfer policies change the individual decision power only in some circumstances is probably a reasonable conclusion. A benefit that is received only for one year is unlikely to have a significant effect, except for marginal couples for whom the marital surplus is already small. A change in intra-household decision power requires, in most households, either extremely large transfers or some commitment to maintain the transfer for many years. Haushofer and Shapiro (2014) provide evidence in support of this view. They consider an experiment where short-term unconditional cash transfers of sizable amount are allocated randomly to the wife or husband. They find no statistical difference in response between households in which the wife received the transfer and households in which the transfer was allocated to the husband.

Finally, the FIC model accounts for the effects of the cash-transfer program in a different way. Like the LIC model, it can easily accommodate the effect of the program on the budget constraint. But the assumption that household members can commit to future plans implies

²⁹Equivalently, “*Progresa* money” should have no effect on z -conditional demands, where the z -conditional demand functions are the collective version of the conditional demand functions defined in Browning and Meghir (1991) for the unitary model. See Bourguignon, Browning, and Chiappori (2009) for a precise definition.

that this model can only allow for changes in decision power for couples that marry *after* the implementation of the program. Households formed before the program is implemented have committed to a plan that is based on the spouses' initial decision power and cannot be changed by subsequent policies. This argument clearly suggests that the LIC and the FIC models can be empirically distinguished. Under the FIC model, the cash transfer should have no significant effect on decision powers in the short run, since all couples are already committed. But its impact should grow with time, since a larger fraction of couples are formed after the introduction of the program. Empirical evidence does not seem to support this prediction, although to the best of our knowledge, no formal test has been performed yet.

Joint versus Individual Taxation

The second policy, a change from a joint to an individual taxation system, has more complicated effects. The direct impact is straightforward. Typically, it increases the hourly wage of at least one spouse, and possibly of both, since after the change individuals are taxed only on their own income. But the exact extent of the effect depends on institutional details, such as the form and size of family deductions, and on other characteristics, such as other incomes and number of children. The policy may also have indirect effects, both in the short and long run. In the short run, the reform may affect both the distribution of powers at the time of household formation and the spouses' outside options in subsequent periods, with effects that are likely to be different for husband and wife, particularly if their labor incomes are dissimilar. From a long-term perspective, the prospect of higher wages and different outside options may affect educational choices of some individuals, the type of jobs and careers they select, their participation in the labor market, and ultimately their stock of human capital. In turn, these changes will influence both the marital surplus and the spouses' outside options, possibly in asymmetric ways.

Similarly to the cash-transfer policy, the static collective model does a better job than the static unitary model in capturing the short-term effects generated by the tax reform. It can account for the changes in decision power generated by the reform and the corresponding reallocation of resources and time across household members. Da Costa and Diniz (2014) provide some insight on the ability of the static collective model to study a tax reform. They analyze from a collective viewpoint the income tax system in the United States, where individuals may choose to file independently or jointly. The main finding of their theoretical analysis is that, although joint filing is generally welfare improving, the possibility of individual filing operates as an outside option by establishing a lower bound on each person's utility. They conclude that the redistribution across spouses due to this effect may be significant, even when the option of individual filing is never chosen.³⁰

However, the ability of static models, whether unitary or collective, to evaluate the tax reform is limited. A significant component of the policy is related to changes over time in wages, human capital, decision power, and intra-household specialization, which are outside the scope of a static approach. Intertemporal models have the potential of capturing those

³⁰The first paper to use a collective model to study the effect of income taxation on individual welfare is Apps and Rees (1988).

effects. The intertemporal unitary model can account for some of the effects generated by the tax reform. By construction, it can account for static and intertemporal changes in wages, the accumulation of human capital, static and dynamic changes in labor supply and household production, and changes in educational choices. Its main limitation is that, since it is unrelated to the intra-household allocation stage, it cannot capture the effects on all those variables of the cross-sectional and longitudinal changes in intra-household decision power generated by the tax reform. If the changes in decision power in favor of one of the spouses are substantial, the intertemporal unitary model may underestimate the impact on years of education, accumulation of human capital, and degree of intra-household specialization. It is up to the researcher to determine whether these additional effects are sufficiently important to require an alternative specification that can also account for the intra-household allocation stage of household decisions.

The LIC model can, in principle, capture all the effects generated by the tax reform. It can easily account for the changes in the wage rates of the two spouses through a modification of the budget constraints. The changes in decision power at the time of marriage can be accounted for by variation in $\mu^1(\Theta)$ and $\mu^2(\Theta)$. The modifications in decision power during marriage can be captured by changes in the outside options and the corresponding changes in $M_{t,\omega}^1(W)$ and $M_{t,\omega}^2(W)$. The long-run effects on human capital, wages, and labor-force participation follow directly the changes in outside options and decision power. The version of the model that we have developed in the survey cannot generate different educational and occupational choices before marriage, but it can be easily generalized by adding a first stage where men and women make those decisions before starting a family. This is done, for instance, in Chiappori, Iyigun, and Weiss (2009), who argue that intrahousehold allocation issues may explain the remarkably different evolution of demand for higher education by men and women since the 1980s, and also in Chiappori, Salanie, and Weiss (2017), Low (2014), and Bronson (2015). The FIC framework can also capture those effects. However, if the reform also changes the spouses' outside options during marriage, as can be expected, it may trigger renegotiations and influence future behavior, an effect that the FIC model cannot account for.

Again, which intertemporal collective model is preferable depends on the particular question being investigated. If households are likely to experience intertemporal variation in relative decision power, the LIC model is a better choice, since it can account for the corresponding changes in human capital and in the degree of intra-household specialization. The FIC model would underestimate the effect of the policy. If households are unlikely to renegotiate their initial plan, the FIC model is a natural choice, since it is simpler to work with and estimate.

Laws Governing Household Formation and Dissolution

The case of divorce laws is highly specific. Issues related to household formation and dissolution can hardly be considered by a unitary framework, if only because ignoring individuals' well being within the household does not seem a very promising path to analyzing their decision regarding staying or quitting the relationship. The only unitary approach that can avoid this criticism is based on transferable utility. Models that use this

approach have an explicit collective structure, since individuals have their own preferences and decisions are Pareto efficient. They boil down to a unitary representation only because a specific assumption on preferences is made. However, the TU approach requires strong and probably unrealistic assumptions. For instance, it posits that individuals are in complete agreement, and therefore do not need any bargaining, on all decisions regarding public consumption. Moreover, it generates very specific predictions, such as the Becker–Coase theorem, whose empirical relevance has been criticized (see for instance Chiappori, Iyigun, and Weiss 2015). Static collective models are equally irrelevant. By nature, issues related to household formation and dissolution involve long-term consequences that cannot be adequately captured by static frameworks. LIC and FIC models, on the other hand, are precisely designed to address such issues.

Indeed, several works explicitly analyze divorce decisions using a collective framework. Gemici and Laufer (2012) study the effect of family laws on marriage and cohabitation choices. Most papers that employ collective models of the household abstract from cohabitation decisions and assume that individuals are either single or married. Using the PSID, Gemici and Laufer (2012) document the following two patterns. First, a large and growing fraction of couples cohabit. Second, in comparison to marriage, cohabitation is associated with a lower degree of household specialization, higher relationship instability, and higher degree of positive assortative mating. With the objective of understanding the reasons behind these differences, the paper estimates an LIC model of the household in which single individuals choose whether to cohabit with or marry the current partner, and married and cohabiting individuals select whether to dissolve the household. In addition, single, cohabiting, and married individuals also choose labor supply and fertility. The only exogenous difference between marriage and cohabitation is that the cost of household dissolution is higher for married couples. The estimated model is then used to conduct counterfactual experiments. The most interesting counter-factual experiment examines the effect of increasing the cost of divorce for married couples. The main finding is that, with higher divorce costs, marriage rates fall, cohabitation rates increase, and married couples experience a higher degree of specialization.³¹

Bayot and Voena (2014) study the effect of changes in labor-force participation, wages, and consequent intra-household specialization patterns on the adoption of prenuptial contracts by married couples. The paper uses administrative data on marriage and divorce from Italy to document a couple of interesting facts about prenuptial contracts. In Italy, community property is the default regime at the time of marriage. But couples may choose to opt for a separate property regime, which is equivalent to adopting a prenuptial contract, at no cost. Bayot and Voena (2014) find that from 1995 to 2011, the majority of newlyweds chose to forgo the community-property regime and maintain separate property. The fraction was 67 percent in 2011. The data also suggest that couples choosing the community-property regime do so to provide insurance against divorce to wives who select to drop out of the labor market and specialize in household production. To rationalize these findings, the paper estimates an LIC model in which individuals make marriage, labor supply, savings, and

³¹Another paper that considers the possible choice between cohabitation and marriage is Matouschek and Rasul (2008).

divorce decisions. The estimated model, which can match the documented patterns, indicates that the gains from adopting the separate-property regime increase with the reduction in the gender–wage gap and with the increase in female labor-force participation. The authors therefore conclude that the fraction of married couples that will select the separate-property regime will increase over time if that choice is available.

Lafortune et al. (2012) use legislative changes affecting cohabiting couples to test the LIC model against the FIC model. Specifically, they study the introduction of a legislation that granted the right to petition for alimony upon separation for cohabiting couples in Canada. These new laws, enacted between 1978 and 2000, were implemented at different times in different provinces with different eligibility rules; moreover, one can easily distinguish between couples who started their relationships before and after the legislative changes. The empirical analysis thus compares the causal estimates of granting alimony rights to partnerships that were already in existence when the new rules were implemented with those that potentially reflect how individuals responded to these changes before entering a union. According to an FIC model, the reform should not affect couples who predate its implementation. An LIC framework generates more complex predictions. For existing couples, the reform improves the outside options of the less wealthy partner, usually the wife, and should therefore translate into an increase in the share of resources she receives; if leisure is a normal good, this should in turn reduce her labor supply. The case of couples who start cohabiting after the reform is entirely different. The authors argue that, since the market for marriage is not globally impacted by the reform, the latter should have a negligible effect on lifetime utilities. However, it drastically changes the distribution of individual utilities over time. The reform improves the wife's prospect after divorce and, from an LIC perspective, in the longer term, whether divorce actually takes place or not; this will be compensated by a *decline* in female utility at the beginning of the relationship, hence an increase in her labor supply. Using an estimation strategy based on triple differences, the paper shows that this prediction is strongly supported by the data.

6. Conclusion

In this review, we have provided an overview of models that have been used to characterize household behavior. The main distinctions we drew are between unitary versus collective models and static versus intertemporal models. We have tried to provide some insight on which features should be considered when choosing one model instead of another. A general insight given in the survey is that a unitary model should be avoided when the question asked requires a good understanding of the changes in intra-household decision power. If that is not the case, the unitary model is still a good choice, since it is less complicated than a collective model. We have then discussed proposed tests, identification results, and estimation findings for the models we have surveyed.

Where to go from here? We believe that this field of economics would benefit from further research in two areas. The first area of research is intertemporal collective models. These models have experienced significant developments over the recent years. Still, much remains to be done. One aspect of household behavior that is particularly relevant in a dynamic setting is the possible existence of information asymmetries between household members.

Dubois and Ligon (2011) provide evidence that incomplete information plays a role in the allocation of food across household members. Ashraf's (2009) field experiment suggests that asymmetric information has an effect on the financial decisions of households. In principle, incomplete information about the spouse's preferences, income, and savings could be an important dimension of household behavior. As discussed in the survey, there is little work about this subject. Developing a general framework that can incorporate different types of asymmetric information, deriving implications to test it, obtaining identification results, and estimating the model to evaluate relevant policies are all important topics for future research.

A second aspect related to intertemporal models that we believe requires more research is the choice of the outside options in intertemporal collective models. In the survey, we have outlined two possible choices: the value of being divorced for at least one period and the value of noncooperation while married. Which outside option better characterizes household decisions? Should the choice be environment specific? Which factors determine in which environment one choice is better than the other? Should both outside options be used at the same time? We believe all these questions should be the subject of future research projects. Moreover, the modeling of divorce as the outside option gives rise to additional research questions. The welfare of an individual after divorce depends not only on her or his well-being but also on her or his remarriage probabilities and the welfare allocation in a potential new marriage. This suggests that, following Becker's early insight, the interaction between divorce, marriage, and intrahousehold allocations should be studied from a general-equilibrium perspective. Only in this way, economists will be able to capture the full effect of the outside options on household behavior.

This takes us to the second area in which we believe further research should be undertaken: the introduction of general-equilibrium considerations in household models. On this front, recent advances in matching models of the marriage market, whether based on frictionless matching (Choo and Siow 2006; Chiappori, Salanie, and Weiss 2011; Salanie and Galichon 2015; and Chiappori 2017) or on search (Jacquemet and Robin 2013; Gousse 2014), open promising perspectives, especially when extended beyond a strict, transferable-utility setting (Chiappori, Costa Dias, and Meghir forthcoming). But these projects are only the starting point of what should be a productive area of research.

References

- Adams, Abi, Cherchye, Laurens, De Rock, Bram, Verriest, Ewout. Consume Now or Later? Time Inconsistency, Collective Choice, and Revealed Preference. *American Economic Review*. 2014; 104(12):4147–83.
- Alderman, Harold, Chiappori, Pierre-André, Haddad, Lawrence, Hoddinott, John, Kanbur, Ravi. Unitary versus Collective Models of the Household: Is it Time to Shift the Burden of Proof? *World Bank Research Observer*. 1995; 10(1):1–19.
- Apps, Patricia F., Rees, Ray. Taxation and the Household. *Journal of Public Economics*. 1988; 35(3): 355–69.
- Apps, Patricia F., Ray, Rees. Collective Labor Supply and Household Production. *Journal of Political Economy*. 1997; 105(1):178–90.
- Ashraf, Nava. Spousal Control and Intrahousehold Decision Making: An Experimental Study in the Philippines. *American Economic Review*. 2009; 99(4):1245–77.

- Ashworth, JS., Ulph, DT. Household Models. In: Brown, CV., editor. *Taxation and Labour Supply*. London: Allen and Unwin; 1981. p. 117-33.
- Attanasio, Orazio P., Lechene, Valerie. Tests of Income Pooling in Household Decisions. *Review of Economic Dynamics*. 2002; 5(4):720-48.
- Attanasio, Orazio P., Lechene, Valerie. Efficient Responses to Targeted Cash Transfers. *Journal of Political Economy*. 2014; 122(1):178-222.
- Attanasio, Orazio P., Weber, Guglielmo. Is Consumption Growth Consistent with Intertemporal Optimization? Evidence from the Consumer Expenditure Survey. *Journal of Political Economy*. 1995; 103(6):1121-57.
- Aura, Saku. Does the Balance of Power within a Family Matter? The Case of the Retirement Equity Act. *Journal of Public Economics*. 2005; 89(9-10):1699-717.
- Bargain, Olivier, et al. Does the Representation of Household Behavior Matter for Welfare Analysis of Tax-Benefit Policies? An Introduction. *Review of Economics of the Household*. 2006; 4(2):99-111.
- Basu, Kaushik. Gender and Say: A Model of Household Behaviour with Endogenously Determined Balance of Power. *Economic Journal*. 2006; 116(511):558-80.
- Bayot, Denrick, Alessandra, Voena. Prenuptial Contracts, Labor Supply and Household Investments. 2014 Unpublished.
- Becker, Gary S. Investment in Human Capital: A Theoretical Analysis. *Journal of Political Economy*. 1962; 70(5 Part 2):9-49.
- Becker, Gary S. A Theory of the Allocation of Time. *Economic Journal*. 1965; 75(299):493-517.
- Becker, Gary S. A Theory of Social Interactions. *Journal of Political Economy*. 1974; 82(6):1063-93.
- Becker, Gary S. *A Treatise on the Family*. Cambridge, MA and London: Harvard University Press; 1981.
- Beninger, Denis, et al. Evaluating the Move to a Linear Tax System in Germany and Other European Countries. *Review of Economics of the Household*. 2006; 4(2):159-80.
- Bergstrom, Theodore C. A Fresh Look at the Rotten Kid Theorem—And Other Household Mysteries. *Journal of Political Economy*. 1989; 97(5):1138-59.
- Bergstrom, Theodore C. Economics in a Family Way. *Journal of Economic Literature*. 1996; 34(4):1903-34.
- Bergstrom, Theodore C., Blume, Lawrence, Varian, Hal. On the Provision of Public Goods. *Journal of Public Economics*. 1986; 29(1):25-49.
- Bergstrom, Theodore C., Cornes, Richard C. Independence of Allocative Efficiency from Distribution in the Theory of Public Goods. *Econometrica*. 1983; 51(6):1753-65.
- Bergstrom, Theodore C., Varian, Hal R. When Do Market Games Have Transferable Utility? *Journal of Economic Theory*. 1985; 35(2):222-33.
- Blundell, Richard, Chiappori, Pierre-André, Meghir, Costas. Collective Labor Supply with Children. *Journal of Political Economy*. 2005; 113(6):1277-306.
- Blundell, Richard, Pashardes, Panos, Weber, Guglielmo. What Do We Learn about Consumer Demand Patterns from Micro Data? *American Economic Review*. 1993; 83(3):570-97.
- Bobonis, Gustavo J. Is the Allocation of Resources within the Household Efficient? New Evidence from a Randomized Experiment. *Journal of Political Economy*. 2009; 117(3):453-503.
- Boone, Jan, van der Karen, Wiel, van Arthur, Soest, Frederic, Vermeulen. Kinky Choices, Dictators and Split Might: A Noncooperative Model for Household Consumption and Labor Supply. *IZA Journal of Labor Economics*. 2014; 3
- Bourguignon, François, Browning, Martin, Chiappori, Pierre-André. Efficient Intrahousehold Allocations and Distribution Factors: Implications and Identification. *Review of Economic Studies*. 2009; 76(2):503-28.
- Bourguignon, François, Browning, Martin, Chiappori, Pierre-André, Lechene, Valérie. Intra Household Allocation of Consumption: A Model and Some Evidence from French Data. *Annals d'Economie et de Statistique*. 1993; 29:137-56.
- Bronson, Mary Ann. Degrees Are Forever: Marriage, Educational Investment, and Lifecycle Labor Decisions of Men and Women. 2015 Unpublished.

- Browning, Martin. The Saving Behavior of a Two-Person Household. *Scandinavian Journal of Economics*. 2000; 102(2):235–51.
- Browning, Martin, Bourguignon, François, Chiappori, Pierre-Andre, Lechene, Valérie. Income and Outcomes: A Structural Model of Intrahousehold Allocation. *Journal of Political Economy*. 1994; 102(6):1067–96.
- Browning, Martin, Chiappori, Pierre-André. Efficient Intra-household Allocations: A General Characterization and Empirical Tests. *Econometrica*. 1998; 66(6):1241–78.
- Browning, Martin, Chiappori, Pierre-André, Lechene, Valérie. Distributional Effects in Household Models: Separate Spheres and Income Pooling. *Economic Journal*. 2010; 120(545):786–99.
- Browning, Martin, Chiappori, Pierre-André, Lewbel, Arthur. Estimating Consumption Economies of Scale, Adult Equivalence Scales, and Household Bargaining Power. *Review of Economic Studies*. 2013; 80(4):1267–303.
- Browning, Martin, Pierre-André, Chiappori, Yoram, Weiss. *Economics of the Family*. Cambridge and New York: Cambridge University Press; 2014.
- Browning, Martin, Valérie, Lechene. Caring and Sharing: Tests between Alternative Models of Intra-household Allocation. University of Copenhagen Department of Economics Discussion Paper 01–07; 2001.
- Browning, Martin, Meghir, Costas. The Effects of Male and Female Labor Supply on Commodity Demands. *Econometrica*. 1991; 59(4):925–51.
- Casanova, Maria, Maurizio, Mazzocco. Individual Rather than Household Euler Equations: Identification and Estimation of Individual Preferences Using Household Data. 2014 Unpublished.
- Chen, Zhiqi, Woolley, Frances. A Cournot–Nash Model of Family Decision Making. *Economic Journal*. 2001; 111(474):722–48.
- Cherchye, Laurens, Demuynck, Thomas, De Rock, Bram. Revealed Preference Analysis of Non-cooperative Household Consumption. *Economic Journal*. 2011; 121(555):1073–96.
- Cherchye, Laurens, Thomas, Demuynck, De Bram, Rock, Frederic, Vermeulen. Household Consumption When the Marriage Is Stable. Catholic University Center for Economic Studies Discussion Paper 14.08; 2014.
- Cherchye, Laurens, De Rock, Bram, Vermeulen, Frederic. The Collective Model of Household Consumption: A Nonparametric Characterization. *Econometrica*. 2007; 75(2):553–74.
- Cherchye, Laurens, De Rock, Bram, Vermeulen, Frederic. Opening the Black Box of Intra-household Decision Making: Theory and Non-parametric Empirical Tests of General Collective Consumption Models. *Journal of Political Economy*. 2009; 117(6):1074–104.
- Cherchye, Laurens, De Rock, Bram, Vermeulen, Frederic. The Revealed Preference Approach to Collective Consumption Behaviour: Testing and Sharing Rule Recovery. *Review of Economic Studies*. 2011; 78(1):176–98.
- Chiappori, Pierre-André. Collective Labor Supply and Welfare. *Journal of Political Economy*. 1992; 100(3):437–67.
- Chiappori, Pierre-Andre. Introducing Household Production in Collective Models of Labor Supply. *Journal of Political Economy*. 1997; 105(1):191–209.
- Chiappori, Pierre-Andre. *Matching with Transfers: The Economics of Love and Marriage*. Princeton: Princeton University Press; 2017.
- Chiappori, Pierre-André, Dias, Monica Costa, Meghir, Costas. The Marriage Market, Labor Supply and Choice. *Journal of Political Economy*. Forthcoming.
- Chiappori, Pierre-André, Donni, Olivier, Komunjer, Ivana. Learning from a Piece of Pie. *Review of Economic Studies*. 2012; 79(1):162–95.
- Chiappori, Pierre-André, Ekeland, Ivar. The Micro Economics of Group Behavior: General Characterization. *Journal of Economic Theory*. 2006; 130(1):1–26.
- Chiappori, Pierre-André, Ekeland, Ivar. The Microeconomics of Efficient Group Behavior: Identification. *Econometrica*. 2009; 77(3):763–99.
- Chiappori, Pierre-André, Fortin, Bernard, Lacroix, Guy. Marriage Market, Divorce Legislation, and Household Labor Supply. *Journal of Political Economy*. 2002; 110(1):37–72.

- Chiappori, Pierre-André, Gugl, E. Necessary and Sufficient Conditions for Transferable Utility. 2015 Unpublished.
- Chiappori, Pierre-André, Iyigun, Murat, Weiss, Yoram. Investment in Schooling and the Marriage Market. *American Economic Review*. 2009; 99(5):1689–713.
- Chiappori, Pierre-André, Iyigun, Murat, Weiss, Yoram. The Becker–Coase Theorem Reconsidered. *Journal of Demographic Economics*. 2015; 81(2):157–77.
- Chiappori, Pierre-André, Lewbel, Arthur. Gary Becker's A Theory of the Allocation of Time. *Economic Journal*. 2015; 125(583):410–42.
- Chiappori, Pierre-André, Jesse, Naidoo. The Engel Curves of Noncooperative Households. 2015 Unpublished.
- Chiappori, Pierre-André, Salanie, Bernard, Weiss, Yoram. Partner Choice, Investment in Children, and the Marital College Premium. *American Economic Review*. 2017; 107(8):2109–67.
- Choo, Eugene, Siow, Aloysius. Who Marries Whom and Why. *Journal of Political Economy*. 2006; 114(1):175–201.
- da Costa, Carlos E., Érica, Diniz. Tax Filing Choices for the Household under Separable Spheres Bargaining. 2014 Unpublished.
- d'Aspremont, Claude, Dos Santos Ferreira, Rodolphe. Household Behavior and Individual Autonomy: An Extended Lindahl Mechanism. *Economic Theory*. 2014; 55(3):643–64.
- Dauphin, Anyck, El Lahga, Abdel-Rahmen, Fortin, Bernard, Lacroix, Guy. Are Children Decision-Makers within the Household? *Economic Journal*. 2011; 121(553):871–903.
- Dauphin, Anyck, Fortin, Bernard. A Test of Collective Rationality for Multi-person Households. *Economics Letters*. 2001; 71(2):211–16.
- Deaton, Angus, John, Muellbauer. *Economics and Consumer Behavior*. Cambridge and New York: Cambridge University Press; 1980.
- Del Boca, Daniela, Flinn, Christopher J. Endogenous Household Interaction. *Journal of Econometrics*. 2012; 166(1):49–65.
- Del Boca, Daniela, Flinn, Christopher J. Household Behavior and the Marriage Market. *Journal of Economic Theory*. 2014; 150:515–50.
- Doepke, Matthias, Michele, Tertilt. Does Female Empowerment Promote Economic Development?. National Bureau of Economic Research Working Paper 19888; 2014.
- Donni, Olivier. A Simple Approach to Investigate Intrahousehold Allocation of Private and Public Goods. *Review of Economics and Statistics*. 2009; 91(3):617–28.
- Dubois, Pierre, Ethan, Ligon. Incentives and Nutrition for Rotten Kids: Intrahousehold Food Allocation in the Philippines. California University Department of Agricultural and Resource Economics Working Paper 1114; 2011.
- Dunbar, Geoffrey R., Lewbel, Arthur, Pendakur, Krishna. Children's Resources in Collective Households: Identification, Estimation, and an Application to Child Poverty in Malawi. *American Economic Review*. 2013; 103(1):438–71.
- Dunbar, Geoffrey R., Arthur, Lewbel, Krishna, Pendakur. Identification of Random Resource Shares in Collective Households with an Application to Microcredit in Malawi. 2014 Unpublished.
- Edlund, Lena, Korn, Evelyn. A Theory of Prostitution. *Journal of Political Economy*. 2002; 110(1):181–214.
- Flavin, Marjorie A. The Adjustment of Consumption to Changing Expectations about Future Income. *Journal of Political Economy*. 1981; 89(5):974–1009.
- Fong, Yuk-fai, Zhang, Junsen. The Identification of Unobservable Independent and Spousal Leisure. *Journal of Political Economy*. 2001; 109(1):191–202.
- Fortin, Bernard, Lacroix, Guy. A Test of the Unitary and Collective Models of Household Labour Supply. *Economic Journal*. 1997; 107(443):933–55.
- Gemici, Ahu. Family Migration and Labor Market Outcomes. 2011 Unpublished.
- Gemici, Ahu, Steve, Laufer. Marriage and Cohabitation. 2012 Unpublished.
- Gersbach, Hans, Haller, Hans. Collective Decisions and Competitive Markets. *Review of Economic Studies*. 2001; 68(2):347–68.
- Gorman WM. Community Preference Fields. *Econometrica*. 1953; 21(1):63–80.

- Gousse, Marion. Marriage Gains: Who Should You Marry?. 2014 Unpublished.
- Greenwood, Jeremy, Guner, Nezih, Knowles, John A. More on Marriage, Fertility, and the Distribution of Income. *International Economic Review*. 2003; 44(3):827–62.
- Grossbard-Shechtman, Amyra. A Theory of Allocation of Time in Markets for Labour and Marriage. *Economic Journal*. 1984; 94(376):863–82.
- Haag, Berthold R., Hoderlein, Stefan, Pendakur, Krishna. Testing and Imposing Slutsky Symmetry in Nonparametric Demand Systems. *Journal of Econometrics*. 2009; 153(1):33–50.
- Hall, Robert E. Stochastic Implications of the Life Cycle-Permanent Income Hypothesis: Theory and Evidence. *Journal of Political Economy*. 1978; 86(6):971–87.
- Hall, Robert E., Mishkin, Frederic S. The Sensitivity of Consumption to Transitory Income: Estimates from Panel Data on Households. *Econometrica*. 1982; 50(2):461–81.
- Haushofer, Johannes, Jeremy, Shapiro. Household Response to Income Changes: Evidence from an Unconditional Cash Transfer Program in Kenya. 2014 Unpublished.
- Heathcote, Jonathan, Storesletten, Kjetil, Violante, Giovanni L. The Macroeconomic Implications of Rising Wage Inequality in the United States. *Journal of Political Economy*. 2010; 118(4):681–722.
- Jacquemet, Nicolas, Jean-Marc, Robin. Assortative Matching and Search with Labour Supply and Home Production. Centre for Microdata Methods and Practice Working Paper 07/13; 2013.
- Kapan, Tumer. Property Division Laws: The Effects on Labor Supply and Household Bargaining. Society for Economic Dynamics Meeting Paper 1127; 2010.
- Kocherlakota, Narayana R. Implications of Efficient Risk Sharing without Commitment. *Review of Economic Studies*. 1996; 63(4):595–609.
- Konrad, Kai A., Lommerud, Kjell Erik. The Bargaining Family Revisited. *Canadian Journal of Economics*. 2000; 33(2):471–87.
- Krueger, Dirk, Perri, Fabrizio. Does Income Inequality Lead to Consumption Inequality? Evidence and Theory. *Review of Economic Studies*. 2006; 73(1):163–93.
- Lafortune, Jeanne, Pierre-André, Chiappori, Murat, Iyigun, Yoram, Weiss. Changing the Rules Midway: The Impact of Granting Alimony Rights on Existing and Newly-Formed Partnerships. Pontificia Universidad Catolica de Chile Instituto de Economia Documento de Trabajo 424; 2012.
- Lechene, Valérie, Preston, Ian. Noncooperative Household Demand. *Journal of Economic Theory*. 2011; 146(2):504–27.
- Leuthold, Jane H. An Empirical Study of Formula Income Transfers and the Work Decision of the Poor. *Journal of Human Resources*. 1968; 3(3):312–23.
- Lewbel, Arthur. Consistent Nonparametric Hypothesis Tests with an Application to Slutsky Symmetry. *Journal of Econometrics*. 1995; 67(2):379–401.
- Lewbel, Arthur. Demand Systems with and without Errors. *American Economic Review*. 2001; 91(3): 611–18.
- Ligon, Ethan. Dynamic Bargaining in Households (with an Application to Bangladesh). California University Department of Agricultural and Resource Economics Working Paper 972; 2002.
- Ligon, Ethan, Thomas, Jonathan P., Worrall, Tim. Informal Insurance Arrangements with Limited Commitment: Theory and Evidence from Village Economies. *Review of Economic Studies*. 2002; 69(1):209–44.
- Lise, Jeremy, Yamada, Ken. Household Sharing and Commitment: Evidence from Panel Data on Individual Expenditures and Time Use. Institute for Fiscal Studies Working Paper W14/05; 2014.
- Low, Corinne. Pricing the Biological Clock: Reproductive Capital on the US Marriage Market. 2014 Unpublished.
- Lundberg, Shelly, Pollak, Robert A. Separate Spheres Bargaining and the Marriage Market. *Journal of Political Economy*. 1993; 101(6):988–1010.
- Lundberg, Shelly, Pollak, Robert A. Efficiency in Marriage. *Review of Economics of the Household*. 2003; 1(3):153–67.
- Lundberg, Shelly, Pollak, Robert A. Marriage Market Equilibrium and Bargaining. Speech delivered at the 2009 Society of Labor Economics Meeting; Boston. 2009.

- Lundberg, Shelly, Pollak, Robert A., Wales, Terence J. Do Husbands and Wives Pool Their Resources? Evidence from the United Kingdom Child Benefit. *Journal of Human Resources*. 1997; 32(3): 463–80.
- Manser, Marilyn, Brown, Murray. Marriage and Household Decision-Making: A Bargaining Analysis. *International Economic Review*. 1980; 21(1):31–44.
- Matouschek, Niko, Rasul, Imran. The Economics of the Marriage Contract: Theories and Evidence. *Journal of Law and Economics*. 2008; 51(1):59–110.
- Mazzocco, Maurizio. PhD Dissertation. University of Chicago; 2001. Essays on Household Intertemporal Behavior.
- Mazzocco, Maurizio. Household Intertemporal Behaviour: A Collective Characterization and a Test of Commitment. *Review of Economic Studies*. 2007; 74(3):857–95.
- Mazzocco, Maurizio, Ruiz, Claudia, Yamaguchi, Shintaro. Labor Supply and Household Dynamics. *American Economic Review*. 2014a; 104(5):354–59.
- Mazzocco, Maurizio, Claudia, Ruiz, Shintaro, Yamaguchi. Labor Supply, Wealth Dynamics, and Marriage Decisions. 2014b Unpublished.
- McElroy, Marjorie B., Horney, Mary Jean. Nash-Bargained Household Decisions: Toward a Generalization of the Theory of Demand. *International Economic Review*. 1981; 22(2):333–49.
- Myck, Michal, et al. The Working Families' Tax Credit and Some European Tax Reforms in a Collective Setting. *Review of Economics of the Household*. 2006; 4(2):129–58.
- Phipps, Shelley A., Burton, Peter S. What's Mine Is Yours? The Influence of Male and Female Incomes on Patterns of Household Expenditure. *Economica*. 1998; 65(260):599–613.
- Pollak, Robert A. Additive Utility Functions and Linear Engel Curves. *Review of Economic Studies*. 1971; 38(4):401–14.
- Pollak, Robert A., Wachter, Michael L. The Relevance of the Household Production Function and Its Implications for the Allocation of Time. *Journal of Political Economy*. 1975; 83(2):255–78.
- Rey, Patrick, Salanié, Bernard. On the Value of Commitment with Asymmetric Information. *Econometrica*. 1996; 64(6):1395–414.
- Salanie, Bernard, Alfred, Galichon. Cupid's Invisible Hand: Social Surplus and Identification in Matching Models. 2015 Unpublished.
- Samuelson, Paul A. Social Indifference Curves. *Quarterly Journal of Economics*. 1956; 70(1):1–22.
- Sargent, Thomas J. Rational Expectations, Econometric Exogeneity, and Consumption. *Journal of Political Economy*. 1978; 86(4):673–700.
- Scholz, John Karl, Seshadri, Ananth, Khitatrakun, Surachai. Are Americans Saving 'Optimally' for Retirement? *Journal of Political Economy*. 2006; 114(4):607–43.
- Schultz, T Paul. Testing the Neoclassical Model of Family Labor Supply and Fertility. *Journal of Human Resources*. 1990; 25(4):599–634.
- Thomas, Duncan. Intra-household Resource Allocation: An Inferential Approach. *Journal of Human Resources*. 1990; 25(4):635–64.
- Thomas, Duncan, Dante, Contreras, Elizabeth, Frankenberg. Distribution of Power within the Household and Child Health. 1999 Unpublished.
- Udry, Christopher. Gender, Agricultural Production, and the Theory of the Household. *Journal of Political Economy*. 1996; 104(5):1010–46.
- Vermeulen, Frederic, et al. Collective Models of Labor Supply with Nonconvex Budget Sets and Nonparticipation: A Calibration Approach. *Review of Economics of the Household*. 2006; 4(2): 113–27.
- Voena, Alessandra. Yours, Mine, and Ours: Do Divorce Laws Affect the Intertemporal Behavior of Married Couples? *American Economic Review*. 2015; 105(8):2295–332.
- Ward-Batts, Jennifer. Out of the Wallet and into the Purse: Using Micro Data to Test Income Pooling. *Journal of Human Resources*. 2008; 43(2):325–51.
- Zeldes, Stephen P. Consumption and Liquidity Constraints: An Empirical Investigation. *Journal of Political Economy*. 1989; 97(2):305–46.

Figure 1.
Shocks and Household Potential Responses

TABLE 1

Relationship between the Three-stage Formulation and Different Models

	Static models		
	Unitary	Nash-bargaining	Collective
Stage 3:		✓	✓
Intra-household allocation stage		A subset of efficient solutions	Any efficient solution
Stage 2:	✓	✓	✓
Resource allocation stage			
Stage 1:			
Intertemporal stage			
	Intertemporal models		
	Unitary	Full efficiency	Limited commitment
Stage 3:		✓	✓
Intra-household allocation stage		Decision power is time-invariant	Decision power can vary over time
Stage 2:	✓	✓	✓
Resource allocation stage			
Stage 1:	✓	✓	✓
Intertemporal stage			

TABLE 2

Relationship between the Three-stage Formulation and Different Testable Implications

	Static models		
	Unitary	Nash bargaining	Collective
Stage 3:		Proportionality	Proportionality
Intra-household allocation stage		Slutsky symmetry + rank 1 Efficiency with production Collective GARP	Slutsky symmetry + rank 1 Efficiency with production Collective GARP
Stage 2:	Income		
Resource allocation stage	Pooling		
Stage 1:			
Intertemporal stage			
	Intertemporal models		
	Unitary	Full efficiency	No commitment
Stage 3:		Proportionality	
Intra-household allocation stage		Slutsky symmetry + rank 1 Efficiency with production Collective GARP	
Stage 2:	Income		
Resource allocation stage	Pooling		
Stage 1:	Euler test	Euler test	Euler test
Intertemporal stage	(Distribution factors have no effect)	(Only cross-sectional variation in distribution factors has effect)	(Both cross-sectional and longitudinal variation have effect)