Endogeneous Household Interaction[1]

Daniela Del Boca[2] Christopher Flinn[3]

August 2010
JEL Classification: C79,D19,J22
Keywords: Household Time Allocation; Grim Trigger Strategy; Household Production; Endogenous Mixing Distribution

[1]This research was partially supported by the National Science Foundation and by the C.V.Starr Center for Applied Economics at NYU. Luca Flabbi provided excellent research assistance at the early stages of this project. For helpful comments and discussions, we are grateful to two anonymous referees, Pierre-Andre Chiappori, Olivier Donni, Douglas Gale, Ahu Gemici, Arthur Lewbel, David Pearce, Yoram Weiss, workshop participants at Torino, Duke, Carnegie-Mellon, Penn, Tel Aviv, the 2007 SITE "Household Economics and the Macroeconomy" segment, the June 2008 Conference on Household Economics in Nice, the June 2008 IFS Conference "Modeling Household Behaviour," and the February 2009 Milton Friedman Institute Conference "New Economics of the Family." We are solely responsible for all errors, omissions, and interpretations.
[2]Department of Economics, University of Torino and Collegio Carlo Alberto
[3]Department of Economics, New York University, and Collegio Carlo Alberto

Abstract

Most econometric models of intrahousehold behavior assume that household decision-making is efficient, i.e., utility realizations lie on the Pareto frontier. In this paper we investigate this claim by adding a number of participation constraints to the household allocation problem. Short-run constraints ensure that each spouse obtains a utility level at least equal to what they would realize under (inefficient) Nash equilibrium. Long-run constraints ensure that each spouse obtains a utility level at least equal to what they would realize by cheating on the efficient allocation and receiving Nash equilibrium payoffs in all successive periods. Given household characteristics and the (common) discount factor of the spouses, not all households may be able to attain payoffs on the Pareto frontier. We estimate these models using a Method of Simulated Moments estimator and data from one wave of the Panel Study of Income Dynamics. We find that both short- and long-run constraints are binding for sizable proportions of households in the sample. We conclude that it is important to carefully model the constraint sets household members face when modeling household allocation decisions, and to allow for the possibility that efficient outcomes may not be implementable for some households.

1 Introduction

Despite there being a vast theoretical and empirical literature on household decision-making processes, there remains a lack of consensus regarding the best way to model household behavior. In Becker's early approach to analyzing household behavior, the social welfare function approach of Bergson-Samuelson was used to define a single objective of the household that represented an aggregated version of individual utilities or that was viewed as representing the preferences of a benevolent dictator directing the household. This "unitary" version of household preferences proved satisfactory in a large number of theoretical analyses, but yielded empirical implications broadly inconsistent with consumption and time allocation decisions observed in household-level data. For this reason, this specification of household objectives has largely been abandoned as a framework for conducting empirical analyses of household behavior.

Over the past several decades, there has been a movement to view the family as a collection of agents with their own clearly-demarcated preferences who are united through the sharing of public goods, emotional ties, and production technologies. Household members are often viewed as behaving strategically with respect to one another given their rather complicated and interconnected resource constraints. Analysis of these situations has focused on describing and analyzing cooperative equilibrium outcomes. Though models using the cooperative approach (e.g., Manser and Brown (1980), McElroy and Horney (1981), Chiappori (1988)) differ in many respects, they share the common characteristic of focusing on outcomes that are Pareto-efficient (the primary distinction between them being the method for selecting a point on the Pareto frontier). The main alternative approach, which uses noncooperative Nash equilibrium as the solution concept (e.g. Leuthold (1968), Bourgignon (1984), Del Boca and Flinn (1995), Chen and Woolley (2001)), leads to outcomes that are generally Pareto-dominated. The analytic attractiveness of noncooperative equilibrium models lies in the fact that equilibria are often unique, an especially distinct advantage when formulating an econometric model and conducting empirical analysis.

In this paper we develop and estimate a model of household time allocation which allows for both efficient and inefficient intrahousehold behavior. Gains from marriage are taken to arise from the presence of a publicly-consumed good that is produced in the household with time inputs from the spouses and goods purchased in the market. In the presence of a public good, efficient outcomes, which by definition lie on the Pareto frontier, must weakly dominate the value of the noncooperative equilibrium for each spouse. So why would some households fail to determine allocations in an efficient manner? We focus on implementation issues in this paper. We consider the case in which household interactions are repeated over an indefinitely long horizon, and determine whether cooperative behavior can be supported using Folk Theorem-inspired arguments. In this view, a key determinant of whether or not Pareto-efficient outcomes can be implemented is the forward-lookingness of the spouses.

There have been few empirical studies to date that have attempted to actually estimate a collective model of household labor supply (some notable exceptions include Kapteyn and Kooreman (1992), Browning et al. (1994), Fortin and Lacroix (1997), and Blundell et al. (2005)). Two of the more important reasons for the paucity of empirical studies are the stringent data requirements for estimation of such a model and lack of agreement regarding the "refinement" to utilize when selecting a unique equilibrium when a multiplicity exist (as is the case in virtually all cooperative models). Some researchers have advocated using the assumption of Pareto efficiency in nonunitary models as an identification device (see, e.g., Bourguignon and Chiappori (1992) and Flinn (2000)).[1] We view household time allocation decisions as either being associated with a particular utility outcome on the Pareto frontier, or to be associated with the noncooperative (static Nash) equilibrium point. In reality there are a continuum of points that dominate the noncooperative equilibrium point and that do not lie on the Pareto frontier, however developing an estimable model that allows such outcomes to enter the choice set of the household seems beyond our means. Our paper expands the equilibrium choice set to two focal points, but it still represents a very restrictive view of the world.

Even under an assumption of efficiency, there is wide latitude in modeling the mechanism by which a specific efficient outcome is implemented, as is evidenced by the lively debate between advocates of the use of Nash bargaining or other axiomatic systems (e.g., McElroy and Horney (1981,1990), McElroy (1990)) and those advocating a more data driven approach (e.g., Chiappori (1988)). The use of an axiomatic system such as Nash bargaining requires that one first specify a "disagreement outcome" with respect to which each party's surplus can be explicitly defined. It has long been appreciated that the bargaining outcome can depend critically on the specification of this threat point.[2] Most often (in the household economics literature) the threat point has been assumed to represent the value to each agent of living independently from the other. Lundberg and Pollak (1993) provide an illuminating discussion of the consequences of alternative specifications of the threat point on the analysis of household decision-making. In particular, instead of assuming the value of the divorce state as the disagreement point for each partner, they consider this point to be determined by the value of the marriage to each given some default mode of behavior, which they call "separate spheres." In this state, each party takes decisions and generally acts in a manner in accordance with "customary" gender roles. Lundberg and Pollak state that households will choose to behave in this customary way when the

[1] An argument sometimes given for this assumption relies on the Folk Theorem. As household members interact frequently and can observe many of each other's constraint sets and actions, for reasonable values of a discount factor efficient behavior should be attainable. The most general behavioral specification estimated here allows us to examine this claim empirically, and we find that efficient allocations cannot be sustained for a small percentage of households.

[2] Lugo-Gil (2003) demonstrates this point nicely. In her case, spouses decide on consumption allocations in a cooperative manner after the outside option is optimally chosen. All "intact" households chose a threat point either of divorce or noncooperative behavior. The choice of threat point has an impact on intrahousehold allocations. In her case, all household allocations are determined efficiently (using a Nash bargaining framework), whereas in ours, some allocations are determined in an inefficient manner.

“transactions costs” they face are too high.

The approach taken in this paper is something of a synthesis of the standard bargaining and “sharing rule” approaches to modeling household behavior. We introduce outside options that household members must recognize and meet, if possible, when choosing efficient allocations. These side conditions on the household’s optimization problem are interpretable more as participation constraints than “threat points,” and these options do not serve as a basis for conducting bargaining in the axiomatic Nash sense. Practically speaking, we view the household allocation decisions as emanating from maximization of the sum of the utilities of the two spouses where the Pareto weight associated with the utility of spouse 1 is α. Adding side constraints to the household optimization problem restricts the set of α-generated time allocations that are implementable, and, depending on the nature of the constraints, may make it impossible to implement any efficient solution. Mazzocco (2007) follows a somewhat similar strategy in his dynamic contracting formulation of the household’s intertemporal allocation problem. In his set up, though, participation constraints may or may not result in reallocations of the cooperative surplus over time. All households are assumed to behave efficiently at all points in time, however.

We build four distinct models of household behavior that are taken to the household time allocation data. They are not strictly nested for the most part, but do have a reasonably natural ordering in terms of complexity. The models are described and developed in a general manner in the following section. They are, in general order of complexity:

1. *(Static) Nash equilibrium*, in which case equilibrium is defined as the intersection of the spouses’ reaction functions.

2. *Pareto-weight* formulation, in which case household allocations are the solutions to the maximization of the Pareto-weighted sum of spousal utilities.

3. *Constrained Pareto-weight* specification, in which case household allocations are solutions to the maximization of the Pareto-weighted sum of spousal utilities subject to the constraint that each spouse receive at least the utility they would receive in Nash equilibrium.

4. *Endogenous Interaction model*, in which case household allocations are solutions to the maximization of the Pareto-weighted sum of spousal utilities subject to the constraint that each spouse receive at least the utility they would receive by deviating from the efficient allocation, if any such outcome is possible. If no such outcome is possible, the household allocations are determined under the (static) Nash equilibrium.

We view the contribution of the paper as bringing short- and long-run implementation issues into the estimation of models of household behavior. While basic versions of the first and second models described above, those of inefficient Nash equilibrium and the "collective" model, have been estimated on numerous occasions, estimation of the collective model subject to the constraint that no spouse's welfare level be less than what they can obtain under the Nash equilibrium has not been performed.[3] We demonstrate empirically that adding such a constraint has significant effects on the point estimates of the model and, consequently, on the welfare inferences that can be drawn.

The addition of the incentive compatibility constraint in our final model specification is noteworthy in that it allows spouses the choice of their mode of interaction. Some households, given their state variables, will choose to behave in an inefficient manner, while others will behave in a "constrained" efficient manner. This carries the important implication that small changes in state variables, such as in wages or nonlabor income, may potentially have large changes on actions if these small changes prompt a change in behavioral regime.

While we have added two sets of constraints to the household optimization problem, it is obviously the case that other types of constraints could be added instead of, or in addition to, the two we have included. For example, it is common to specify the value of each spouse in the divorce state as the disagreement point when using a Nash bargaining framework to analyze household decisions. Clearly, this constraint could be added to those considered here when determining the set of implementable values of the Pareto weight α. These types of generalizations are left to future research, and are probably best considered in a truly dynamic model of the household that allows for divorce outcomes.

The plan of the paper is as follows. In Section 2 we lay out the theoretical structure of the model. Section 3 presents the specific functional form assumptions that underlie the empirical analysis. Section 4 contains a discussion of estimation issues and develops the nonparametric and parametric estimators used in the empirical analysis. We provide a description of the data, present model estimates, and conduct a small comparative statics exercise in Section 5. Section 6 concludes.

[3] Mazzocco (2007) is a valuable contribution to the literature that also considers implementation issues explicitly. The focus of his analysis is on determining whether intertemporal household allocations are consistent with *ex ante* efficient allocations, or whether welfare weights have to be continually adjusted to meet short-run participation issues that arise when household members cannot credibly make life-long commitments to a given *ex ante* efficient allocation. He finds evidence supporting the lack of commitment hypothesis.

There are a number of differences between his approach and the one taken here, the most salient of which are the following. First, the dynamic setting he considers is not nearly as stylized as the one employed here. Second, participation constraints change over the life cycle, though they are modeled as exogenous random processes, whereas in our case the outside option is explicitly modeled. Third, in his model all households behave efficiently in every period, that is, the outside option is never chosen. In our case, the outside option of inefficient behavior is chosen in some states of the world.

Another valuable paper that examines commitment issues in a household setting using a dynamic contracting approach is Ligon (2002). As in Mazzocco (2007), all final household allocation decisions are constrained efficient, which is not the case in our EI specification.

2 The Household's Decision Problem

In the first part of this section we describe the objectives and constraints facing the spouses. We then turn to a consideration of the manner in which a household equilibrium allocation is determined under our four different modeling assumptions.

2.1 The General Environment

While the empirical application will require us to make a number of functional form assumptions, the model of household behavior we develop is applicable quite generally. We assume egoistic preferences, with the current period payoff function of spouse i given by the function

$$u_i(a_1, a_2),$$

where a_i is a vector of actions available to spouse i. These actions may consist of purchases of inputs for a household production technology, the supply of time to that technology, or purchases of market goods, the consumption of which directly impacts the satisfaction of one or both spouses. All of the actions considered in the paper are continuous, and choices will be a continuous function of the state variables facing the household. To be able to invert decision rules to recover primitive parameters, we will require the existence of a 1-1 mapping between a subset of (unobserved) state variables and the (observed) decisions. This will require us to rule out corner solutions.

We assume the presence of intrahousehold externalities, which for us means that the well-being of a spouse or their (household) productivity is directly affected by the actions of the other spouse. Given the presence of a household externality, assumed throughout, there exists a welfare-improving set of feasible actions that differ from the (static) Nash equilibrium actions.

2.2 The Behavioral Regimes

2.2.1 Nash Equilibrium

We make the following assumptions.

Assumption 1 There exists a unique Nash equilibrium in actions

$$(\hat{a}_1^N \ \hat{a}_2^N)(u_1, u_2, CS)$$

where U is a family of period-payoff functions and CS denotes the choice set.

Assumption 1 rules out, most importantly, various types of nonconvexities in the choice sets of the household members, such as the existence of fixed time or money costs associated with supplying time to the labor market or in household production. The Nash equilibrium is derived as follows. Let $a_i^*(a_{i'}, u_i, CS_i)$ denote the best response functions of spouse

i given that spouse i' chooses actions $a_{i'}$ when they have choice set CS_i and objective function u_i. Then the Nash equilibrium actions are given by $(\hat{a}_1^N\ \hat{a}_2^N)(u_1, u_2, CS)$, where $\hat{a}_1^N = a_1^*(\hat{a}_2^N, u_1, CS_1)$ and $\hat{a}_2^N = a_2^*(\hat{a}_1^N, u_2, CS_2)$, and $CS = CS_1 \cup CS_2$.

2.2.2 Pareto Optimal Decisions

We write the Benthamite social welfare function for the household as

$$W_\alpha(a_1, a_2) = \alpha u_1(a_1, a_2) + (1-\alpha) u_2(a_1, a_2).$$

Maximizing the value of $W_\alpha(a_1, a_2)$ with respect to the actions of the two spouses given CS traces out the Pareto frontier. In particular,

$$(\hat{a}_1^P\ \hat{a}_2^P)(\alpha, u_1, u_2, CS) = \arg \max_{(a_1, a_2) \in CS} W_\alpha(a_1, a_2).$$

Assumption 2 The Pareto optimal actions $(\hat{a}_1^P\ \hat{a}_2^P)(\alpha, u_1, u_2, CS)$ are unique.

The Pareto frontier is defined by the set of points

$$\begin{gathered} PF(u_1, u_2, CS) \equiv \\ \{u_1(\hat{a}_1^P(\alpha, u_1, u_2, CS), \hat{a}_2^P(\alpha, u_1, u_2, CS)), u_2(\hat{a}_1^P(\alpha, u_1, u_2, CS), \hat{a}_2^P(\alpha, u_1, u_2, CS))\}, \\ \alpha \in [0,1] \end{gathered}$$

It follows that u_1 is nondecreasing in α and u_2 is nonincreasing in α along PF. Assuming differentiability along the Pareto frontier with respect to α,

$$\left.\frac{du_2}{du_1}\right|_{\{u_2, u_1\} \in PF} < 0.$$

Our analysis is based on the existence of externalities within the household. It follows that the pair of Nash equilibrium utility levels, $\{V_1^N, V_2^N\}$, defined by

$$V_i^N = u_i(\hat{a}_1^N, \hat{a}_2^N),\ i = 1, 2,$$

is not a point on the Pareto frontier.

2.2.3 Constrained Pareto Outcomes

Pareto efficient outcomes have the desirable feature that one spouse's utility cannot be improved without decreasing the other's, but may or may not meet certain reasonable "fairness" criteria. We singled out Nash equilibrium play as a natural focal point for the behavior of spouses since it involves no coordination or policing of actions due to the fact that strategies are best responses. This gives the Nash equilibrium a type of stability not possessed by efficient outcomes.

At a minimum, then, it seems reasonable to restrict attention to efficient outcomes that yield each spouse at least the level of welfare they can attain under the Nash equilibrium actions. We think of this as a "short-run" participation constraint, where short-run refers to the fact that it ensures that each party is at least as well-off under the efficient allocation as they would be under Nash equilibrium in the current period. The following proposition follows directly from the relatively weak assumptions made to this point. To reduce notational clutter, we drop the explicit conditioning on the utility functions u_i and the choice set CS.

Proposition 1 *There exists a nonempty interval* $I^C(V_1^N, V_2^N) \equiv [\underline{\alpha}(V_1^N), \overline{\alpha}(V_2^N)] \subset (0,1)$, *with* $\underline{\alpha}(V_1^N) < \overline{\alpha}(V_2^N)$, *such that*

$$\begin{aligned} u_1(\hat{a}_1^P(\alpha), \hat{a}_2^P(\alpha)) &\geq V_1^N \\ u_2(\hat{a}_1^P(\alpha), \hat{a}_2^P(\alpha)) &\geq V_2^N \end{aligned}$$

if and only if $\alpha \in I^C(V_1^N, V_2^N)$.

Proof. The points on the Pareto frontier are monotone and continuous functions of α. By Assumption 2, the Nash equilibrium payoffs are dominated by a set of points on the Pareto frontier. Define the unique value $\underline{\alpha}(V_1^N)$ by

$$u_1(\hat{a}_1^P(\underline{\alpha}(V_1^N), \hat{a}_2^P(\underline{\alpha}(V_1^N))) = V_1^N,$$

and, similarly, define $\overline{\alpha}(V_2^N)$ by

$$u_2(\hat{a}_1^P(\overline{\alpha}(V_2^N)), \hat{a}_2^P(\overline{\alpha}(V_2^N))) = V_2^N.$$

Define the set of all α values that yield an efficient payoff at least as large as V_1^N for spouse 1 as

$$I_1^C(V_1^N),$$

the minimum element of which is $\underline{\alpha}(V_1^N)$. Define the set of all values of α that yield an efficient payoff at least as large as V_2^N to spouse 2 as

$$I_2^C(V_2^N),$$

the maximum element of which is $\overline{\alpha}(V_2^N)$. Then

$$\begin{aligned} I^C(V_1^N, V_2^N) &= I_1^C(V_1^N) \cap I_2^C(V_2^N) \\ &= [\underline{\alpha}(V_1^N), \overline{\alpha}(V_2^N)] \\ &\neq \varnothing \end{aligned}$$

■

Given problems associated with the identification of the welfare weight α, which are discussed in detail below, we will typically assume that there exists one value of α, common to all marriages, which could be culturally determined. The constrained Pareto optimal allocation is determined by first determining whether $\alpha \in I^C(V_1^N, V_2^N)$. If so, each spouse's utility level using the Pareto weight of α exceeds their static Nash equilibrium utility level, and the constraint is not binding. Instead, if $\alpha < \underline{\alpha}(V_1^N)$, the Pareto efficient solution yields less utility to spouse 1 than the Nash equilibrium solution. To get this spouse to participate in the Pareto efficient solution, it is necessary to provide them with at least as much utility as they would obtain in the static Nash equilibrium, which means adjusting the Pareto weight up to the value $\underline{\alpha}(V_1^N)$. Conversely, if $\alpha > \overline{\alpha}(V_2^N)$, then the Pareto weight has to be adjusted downward to $\overline{\alpha}(V_2^N)$ to provide the incentive for the second spouse to participate in the household efficient allocation. The formal statement of the actions in this environment is as follows. If $\alpha \in I^C$, then

$$(\hat{a}_1^C \ \hat{a}_2^C)(\alpha) = (\hat{a}_1^P \ \hat{a}_2^P)(\alpha), \ \alpha \in I^C, \tag{1}$$

since the "participation constraint" is not binding. If $\alpha < \underline{\alpha}(V_1^N)$, so that spouse 1 would have a higher payoff in Nash equilibrium, the α must be "adjusted" up so that he has the same welfare in either regime. In this case,

$$(\hat{a}_1^C \ \hat{a}_2^C)(\alpha) = (\hat{a}_1^P \ \hat{a}_2^P)(\underline{\alpha}(V_1^N)), \ \alpha < \underline{\alpha}(V_1^N). \tag{2}$$

Conversely, if spouse 2 suffers utility-wise in the efficient allocation associated with α, the α must be adjusted downward, and we have

$$(\hat{a}_1^C \ \hat{a}_2^C)(\alpha) = (\hat{a}_1^P \ \hat{a}_2^P)(\overline{\alpha}(V_2^N)), \ \alpha > \overline{\alpha}(V_2^N). \tag{3}$$

Note that under this behavioral rule, there is still, in general, a continuum of possible solutions, associated with all values of α belonging to I^C.

2.2.4 Endogenous Interaction

The time allocations in the Pareto optimal and constrained Pareto optimal cases may or may not satisfy another set of constraints, one that involves implementation. The essential issue is that utility levels that lie along the Pareto frontier are not associated with time allocation choices by either spouse that are "best responses" (in the static sense) to the choices of their partner. As we know, only the static Nash equilibrium has that property, and is associated with utility outcomes that are dominated by those associated with the constrained Pareto optimal choices we have just discussed.

Why might spouses cheat on an efficient agreement that improves the welfare of both with respect to the Nash equilibrium outcome? The temptation to cheat in this case may arise from purely self-interested behavior, as it does when we study the incentives of firms engaged in collusive behavior to deviate from their assigned production quotas (e.g., Green and Porter, 1984). In that case, the welfare of firms is linked through a common market for outputs or inputs, though firms' objectives are typically taken to be solely the maximization of their own monetary profits. In the case of households, the objectives of spouses may be considerably more complex than those of firms, and may include altruism, for example. However, the existence of caring preferences, in and of itself, does not make implementation of an efficient allocation a foregone conclusion. Indeed, spouses may care so much about each other that an efficient solution may involve each behaving in what would appear to be a more self-interested manner to an observer. In this case, "cheating" on the efficient outcome may imply a spouse spends more of their resources on goods of direct value only to the other spouse. In the household context cheating may be prevalent, but due to the presence of household production technologies and interconnected preferences, it is difficult to detect without strong assumptions on preferences or technologies.

As is well-known from Folk theorem results, in order to implement equilibrium outcomes that are not best responses in a static sense, it is necessary to provide an intertemporal context to household choices. Accordingly we define the welfare of each spouse to be

$$J_i = \sum_{t=1}^{\infty} \beta^{t-1} u_i(a_1(t), a_2(t)),$$

where β is a discount factor taking values in the unit interval, and $a_j(t)$ are the actions chosen by spouse j in period t. For reasons related to data availability and computational feasibility, we restrict our attention to the case in which the stage game played by the spouses has the same structure in every period. That is, preferences and technology parameters are fixed over time, as well as wage offers and nonlabor income levels.

We assume that the couple utilize a *grim trigger strategy,* with the punishment phase being perpetual Nash equilibrium play.[4] We assume that the allocation is determined by

$$\hat{a}_i^E(t) = \begin{cases} \hat{a}_i^C(\alpha) & \text{if} \quad a_{i'}(t') = \hat{a}_{i'}^C(\alpha),\ t' = 1, ..., t-1 \\ \hat{a}_i^N & \text{if} \quad a_{i'}(t') \neq \hat{a}_{i'}^C(\alpha) \text{ for any } t' = 1, ..., t-1 \end{cases} \tag{4}$$

in the equilibrium. In this case, a divergence by either spouse from their prescribed action $\hat{a}_i^C(\alpha)$ leads to the play of Nash equilibrium in all subsequent periods. $\hat{a}_i^C(\alpha)$ denotes the prescribed efficient allocation, which is determined using the Pareto weight α.

To determine whether there exists an implementable cooperative equilibrium in the household, we must check whether there is sufficient patience among the spouses to sustain the efficient outcome given the size of the penalty they face for deviation. Each spouse's objective is to maximize the present discounted value of their sequence of payoffs given the

[4] We are aware that there are more 'efficient' punishment strategies available to the household members, but the incorporation of these punishments into the econometric model is a difficult task. The important point for the analysis is that given our modeling set up, a measurable subset of the state vector space will result in a lack of implementability of efficient outcomes, the main point of our analysis.

state variables characterizing the household and the history of past actions of the spouses. To determine whether or not cooperation is an equilibrium outcome, define the value of spouse 1 cheating on the cooperative agreement given that spouse 2 does not by

$$V_1^R(\alpha) + \beta \frac{V_1^N}{1-\beta}, \tag{5}$$

where

$$V_1^R(\alpha) = \max_{a_1} u_1(a_1, \hat{a}_2^C(\alpha)),$$

and where the second term on the right hand side of (5) is the discount rate multiplied by the present value of the noncooperative equilibrium, which is the outcome of a deviation from $\hat{a}_1^C(\alpha)$ under (4). If the spouse chooses to implement the cooperative outcome (and it is assumed that spouse 2 chooses $\hat{a}_2^C(\alpha)$), then the payoff from this action is

$$\frac{V_1^C(\alpha)}{1-\beta}.$$

Spouse 1 is indifferent between reneging and implementing the cooperative equilibrium when

$$\frac{V_1^C(\alpha)}{1-\beta} = V_1^R(\alpha) + \beta \frac{V_1^N}{1-\beta}. \tag{6}$$

The discount factor β is not a determinant of stage game payoffs, so we can look for a critical value of the discount factor at which the equality (6) holds. This critical value is given by

$$\beta_1^*(\alpha) = \frac{V_1^R(\alpha) - V_1^C(\alpha)}{V_1^R(\alpha) - V_1^N}.$$

Note that since $V_1^C(\alpha) > V_1^N(\alpha)$, then $V_1^R(\alpha) > V_1^C(\alpha)$, and $V_1^R(\alpha) - V_1^C(\alpha) < V_1^R(\alpha) - V_1^N$, so that $\beta_1^*(\alpha) \in (0,1)$. Clearly, an exactly symmetric analysis can be used to determine a critical discount factor for the spouse 2, $\beta_2^*(\alpha)$. This leads us to the following result.

Proposition 2 *Under a grim trigger strategy and given constrained Pareto optimal actions* $(\hat{a}_1^C, \hat{a}_2^C)(\alpha)$, *the household implements the efficient outcome if and only if* $\beta \geq \max\{\beta_1^*(\alpha), \beta_2^*(\alpha)\}$, *where* β *is the common household discount factor and* α *is the given Pareto weight.*

Proof. Under complete information, the values $(\beta_1^*(\alpha), \beta_2^*(\alpha))$ are known to both spouses. If $\beta \geq \beta_i^*(\alpha)$ for $i = 1, 2$, each agent knows that the value of implementing the cooperative solution forever dominates the value of reneging for each, so playing cooperative in each period is a best response for each spouse and constitutes a Nash equilibrium. Say that $\beta_1^*(\alpha) \leq \beta$ but $\beta_{2'}^*(\alpha) > \beta$. The value of spouse 1 choosing $\hat{a}_1^C(\alpha)$ will be

$$u_1(\hat{a}_1^C(\alpha), \hat{a}_2^R(\alpha)) + \beta \frac{V_1^N}{1-\beta}$$

under the grim trigger strategy, since $\hat{a}_2^R(\alpha) \neq \hat{a}_2^C(\alpha)$ triggers the punishment phase. Given the reneging action $\hat{a}_2^R(\alpha)$, $\hat{a}_1^C(\alpha)$ will not maximize this payoff, and spouse 1 will best respond $a_1^*(\hat{a}_2^R(\alpha))$, say, to which spouse 2, will best respond, with the actions converging to those of the (unique) Nash equilibrium $(\hat{a}_1^N, \hat{a}_2^N)$. Thus both spouses will play the Nash equilibrium at every point in time. For the same reason, when $\beta < \beta_1^*(\alpha)$ and $\beta < \beta_2^*(\alpha)$, the sequence of best responses to the reneging behavior of the other spouse leads to the Nash equilibrium being played in each period. ■

We now turn to the consideration of the determination of the actions $(\hat{a}_1^E(\alpha), \hat{a}_2^E(\alpha))$ in the Endogenous Interactions case. After determining the efficient allocation of the household under CPO given an initial notional value of α, we can determine if this solution is *implementable.* For simplicity, let α_{CPO} denote the *ex post* value of α that satisfies the participation constraint for a household (given the choice set CS) under the CPO specification. If $\beta \geq \beta_1^*(\alpha_{CPO})$ and $\beta \geq \beta_2^*(\alpha_{CPO})$, then the CPO outcome is implementable, and the actions in the Endogenous Interactions case are the same as are specified under CPO.

In general, the *ex post* Pareto weight associated with the CPO regime is not implementable under the EI regime. This is clearly the case when α_{CPO} is determined in such a way that the participation constraint is binding for one of the spouses, which is always the case whenever $\alpha_{CPO} \neq \alpha_0$. In this case, there will be no long run welfare gains for the spouse with the binding participation constraint, and his or her best response will be to cheat on the efficient outcome. In such a case, to induce that spouse not to deviate from the efficient outcome, the Pareto weight associated with that spouse must be increased. If there is an implementable efficient outcome, for a given value of β, it will be the one for which the "long run" participation (i.e., no cheating) constraint is exactly satisfied. For a given value of β, there may be no value of the Pareto weight that simultaneously satisfies the "no cheating" constraint for both spouses, and in this case, no efficient allocation is attainable. The formal definition of implementability is the following:

Definition 3 *A household has an implementable outcome on the Pareto frontier if there exists an $\alpha \in (0,1)$ such that $\beta \geq \max\{\beta_1^*(\alpha), \beta_2^*(\alpha)\}$.*

Figure 1 contains the graph of the β_i^*, $i = 1, 2$, for a given set of state variables that fully characterize spousal preferences, household technology, and choice sets. We note that β_1^* is a decreasing function of α, since increasing (static) gains associated with the efficient allocation (as α increases) require lower levels of patience to sustain implementation on the part of spouse 1. Obviously, β_2^* is increasing in α for the opposite reason. We see that for this set of state variables, the household has an implementable efficient outcome if the common discount factor of the spouses exceeds β^{**}. If the discount factor is less than that, no outcome on the Pareto frontier can be implemented, even though there are a continuum of allocations with static payoffs that strictly exceed the static Nash equilibrium payoffs.

In Figure 1 we have also indicated the manner in which the *ex post* Pareto weight is determined when an implementable allocation exists. Since the value of the discount factor, β, exceeds β^{**}, an implementable solution exists. Starting from the *ex post* α associated with the static Nash equilibrium participation constraint, α_{CPO}, we see that at this value $\beta_1^*(\alpha_{CPO}) > \beta$ and $\beta_2^*(\alpha_{CPO}) < \beta$, so that at this low level of α spouse 1 would cheat on the efficient allocation, while spouse 2 would not, meaning that in equilibrium the $\alpha_{CPO}-$generated allocation could not be implemented. In this case, α is increased until it reaches the value α_{EI}, which is that value at which the no-cheating participation constraint is met for spouse 1.

In summary, we think of the determination of the EI allocations as consisting of the following steps.

1. Determine the functions $\beta_j^*(\alpha)$.
2. If $\beta < \beta^{**}$, then the household is not able to implement efficient allocations. Then $\hat{a}_j^E = \hat{a}_j^N$, $j = 1, 2$.
3. If $\beta \geq \beta^{**}$, the household is able to implement an efficient time allocation. Let α_0 denote the notational Pareto weight. Then

$$(\hat{a}_1^E, \hat{a}_2^E) = \begin{cases} (a_1^E(\alpha_0), a_2^E(\alpha_0)) & \text{if} \quad \beta_1^*(\alpha_0) \leq \beta \text{ and } \beta_2^*(\alpha_0) \leq \beta \\ (a_1^E((\beta_1^*)^{-1}(\beta)), a_2^E((\beta_1^*)^{-1}(\beta))) & \text{if} \quad \beta_1^*(\alpha_0) > \beta \\ (a_1^E((\beta_2^*)^{-1}(\beta)), a_2^E((\beta_2^*)^{-1}(\beta))) & \text{if} \quad \beta_2^*(\alpha_0) > \beta \end{cases},$$

where $(\beta_j^*)^{-1}$ is the inverse of β_j^*.

2.3 Summary

We summarize the results of this section with the aid of Figure 2. For any given household, there exists a unique (static) household Nash equilibrium of actions $\{\hat{a}_1^N, \hat{a}_2^N\}$ and payoffs $\{V_1^N, V_2^N\}$, with the pair of payoffs given by the intersection of the two lines in Figure 2. Varying the Pareto weight α over $(0, 1)$ in the weighted household utility function specification (PO) traces out the Pareto frontier. When we impose the side constraint that the Pareto weight must be chosen so that each spouse obtains at least their payoff V_i^N, this defines a lower bound $\underline{\alpha}(V^N)$ at which the participation constraint is just binding for spouse 1 and an upper bound $\overline{\alpha}(V^N)$ at which the participation constraint is just binding for spouse 2. If the notional value of the Pareto weight, α_0, falls in this interval, then that value is used to define the efficient outcome, which will be the same as in the unconstrained case. If the value of α_0 is less than $\underline{\alpha}(V^N)$, then the efficient outcome is determined using the Pareto weight $\underline{\alpha}(V^N)$. If, instead, $\alpha_0 > \overline{\alpha}(V^N)$, then the efficient outcome is determined using the Pareto weight $\overline{\alpha}(V^N)$.

The "dynamic" participation constraint imposes a tighter set of restrictions on the α choice problem than does the "static" participation constraint, except in the extreme case of $\beta = 1$. For any $\beta < 1$, there either exists a nonempty interval $[\underline{\alpha}(V^N, \beta), \overline{\alpha}(V^N, \beta)] \subset [\underline{\alpha}(V^N), \overline{\alpha}(V^N)] = [\underline{\alpha}(V^N, \beta = 1), \overline{\alpha}(V^N, \beta = 1)]$, or the set of implementable α is empty, and inefficient behavior results. Put another way, for any given household, there exists a critical value β^{**}, with any $\beta < \beta^{**}$ inducing the household to behave inefficiently. When there does exist a nonempty set of α that satisfy the dynamic participation constraint, the ultimate household allocation is determined in the same manner as it was when we imposed the static participation constraint.

3 Functional Form Assumptions

In order to estimate the model with the data available to us, a number of functional form assumptions are required. Our strategy in this first section is to present the assumptions made regarding the form of spousal preferences and the household production technology, and then to demonstrate that the resulting household behavior is consistent with Assumptions 1 and 2.

We assume that each spouse possesses a utility function defined over the consumption of a good produced in the household with time inputs of the household members and a good (or goods) purchased in the market, which we denote by τ_1, τ_2, and M, respectively. The household production technology is given by

$$K = \tau_1^{\delta_1} \tau_2^{\delta_2} M^{1-\delta_1-\delta_2},$$

where M is total household income, τ_i is the time supplied by spouse i in household production, and $\delta_1 \geq 0$, $\delta_2 \geq 0$, and $\delta_1 + \delta_2 \leq 1$. Thus the household production technology exhibits constant returns to scale.

Individuals supply time in a competitive market to generate earnings. The wage rate of spouse i is w_i, and the time they spend in market activities is h_i. The nonlabor income of the household is denoted by Y, and is the sum of the nonlabor incomes of the spouses, $Y_1 + Y_2$, plus any nonattributable nonlabor income that accrues to the household, $\tilde{Y}$. The sources of nonlabor income will play no role in our analysis. Total income of the household is then $M = w_1 h_1 + w_2 h_2 + Y$.

Each spouse has Cobb-Douglas preferences over a private good, leisure, and the household public good, so that

$$u_i = \lambda_i \ln(l_i) + (1 - \lambda_i) \ln K, \ i = 1, 2,$$

where l_i is the leisure consumed by spouse i, and $\lambda_i \in [0, 1]$, $i = 1, 2$, is the Cobb-Douglas preference parameter. Each spouse has a time endowment of T, so that

$$T = h_i + \tau_i + l_i, \ i = 1, 2,$$

defines the time constraint. Spouse i controls his or her time allocations regarding h_i and τ_i (which determine l_i, of course).

By substituting the household production technology into the utility functions of each spouse, we get a payoff function

$$\begin{aligned} u_1(l_1, \tau_1, \tau_2, M) &= \lambda_1 \ln l_1 + \tilde{\delta}_{11} \ln \tau_1 + \tilde{\delta}_{12} \ln \tau_2 + \tilde{\delta}_{13} \ln M \\ u_2(l_2, \tau_1, \tau_2, M) &= \lambda_2 \ln l_2 + \tilde{\delta}_{21} \ln \tau_1 + \tilde{\delta}_{22} \ln \tau_2 + \tilde{\delta}_{23} \ln M, \end{aligned}$$

where $\delta_3 = 1 - \delta_2 - \delta_1$ and $\tilde{\delta}_{ij} \equiv (1 - \lambda_i)\delta_j$, and where $\lambda_i + \tilde{\delta}_{i1} + \tilde{\delta}_{i2} + \tilde{\delta}_{i3} = 1$, $i = 1, 2$. The reaction function of spouse 1, for example, is

$$(h_1^* \; \tau_1^*)(h_2 \; \tau_2) = \arg \max_{h_1 \geq 0, \tau_1} \lambda_1 \ln(T - h_1 - \tau_1) + \tilde{\delta}_{11} \ln \tau_1 + \tilde{\delta}_{12} \ln \tau_2 + \tilde{\delta}_{13} \ln(w_1 h_1 + w_2 h_2 + Y),$$

and it is straightforward to establish that this best response is uniquely determined for all $(h_2 \; \tau_2)$ and for any values of the preference and production parameters satisfying the constraints imposed above. There exists a unique best response for the second spouse, $(h_2^* \; \tau_2^*)(h_1 \; \tau_1)$ by symmetry. Since h_i^* is a nonincreasing function of $h_{i'}$ and is independent of $\tau_{i'}$, and since τ_i^* is independent of both $h_{i'}$ and $\tau_{i'}$, there exists a unique Nash equilibrium

$$\begin{aligned} (\hat{h}_1^N \; \hat{\tau}_1^N) &= (h_1^* \; \tau_1^*)(\hat{h}_2^N \; \hat{\tau}_2^N) \\ (\hat{h}_2^N \; \hat{\tau}_2^N) &= (h_2^* \; \tau_2^*)(\hat{h}_1^N \; \hat{\tau}_1^N). \end{aligned}$$

Thus these functional form assumptions satisfy Assumption 1.

When considering the Pareto weight allocations, we can write the household objective function as

$$\begin{aligned} W_\alpha(l_1, l_2, \tau_1, \tau_2, M) &= \alpha u_1(l_1, \tau_1, \tau_2, M) + (1 - \alpha) u_2(l_2, \tau_1, \tau_2, M) \\ &= \tilde{\lambda}_1(\alpha) \ln l_1 + \tilde{\lambda}_2(\alpha) \ln l_2 + \tilde{\phi}_1(\alpha) \ln \tau_1 + \tilde{\phi}_2(\alpha) \ln \tau_2 + \tilde{\phi}_3(\alpha) \ln M, \end{aligned}$$

where $\tilde{\lambda}_1(\alpha) \equiv \alpha\lambda_1$, $\tilde{\lambda}_2(\alpha) \equiv (1 - \alpha)\lambda_2$, $\tilde{\phi}_j(\alpha) \equiv \omega(\alpha)\delta_j$, $j = 1, 2, 3$, and $\omega(\alpha) \equiv (\alpha(1 - \lambda_1) + (1 - \alpha)(1 - \lambda_2)$. Since the objective function is concave in its arguments and the constraint set is convex, there exists a unique solution to the household's optimization problem, which, for a given value of α, is given by

$$\begin{aligned} (h_1^P \; \tau_1^P \; h_2^P \; \tau_2^P) &= \arg \max_{\substack{h_1 \geq 0, h_2 \geq 0. \\ \tau_1, \tau_2}} \Big\{ \tilde{\lambda}_1(\alpha) \ln(T - h_1 - \tau_1) + \tilde{\lambda}_2(\alpha) \ln(T - h_2 - \tau_2) \\ &\quad + \tilde{\phi}_1(\alpha) \ln \tau_1 + \tilde{\phi}_2(\alpha) \ln \tau_2 + \tilde{\phi}_3(\alpha) \ln(w_1 h_1 + w_2 h_2 + Y) \Big\}. \end{aligned} \tag{7}$$

Thus these functional forms are consistent with Assumption 2.

We conclude this brief section with a comment on the specific functional form assumptions adopted for the empirical analysis. While the assumptions on preferences and production technology do imply a number of strong restrictions on household behavior, we will allow for very general forms of between- and within-household heterogeneity that will enable us to capture most or all of the variation in the data. In fact, under the most general specification, we show that the distribution of parameters is just-identified given the PSID available to us. The point being made is that the restrictiveness of functional form assumptions can only be judged in conjunction with the type of restrictions imposed on the population distribution of parameters. That said, other functional form assumptions that satisfied Assumptions 1 and 2 could have been used as a basis for the empirical work, and the implications drawn regarding the behavioral choices of households and marital sorting patterns could well have been different using a different set of "basis functions."

4 Econometric Framework

A household "stage game" equilibrium is uniquely determined given a vector S of state variables that, given the functional form assumptions maintained, completely characterize the preferences of both spouses and the choice set of the household. The state variables are given by

$$S = \begin{pmatrix} \lambda_1 \\ \lambda_2 \\ \delta_1 \\ \delta_2 \\ w_1 \\ w_2 \\ Y \\ \alpha \end{pmatrix}.$$

The vector S uniquely determines the efficient and inefficient solutions to the household's time allocation problem *given* the mode of behavior: inefficient (Nash equilibrium), Pareto optimal, or constrained Pareto optimal. Adding the discount factor β allows us to determine which mode of behavior is observed, thus

$$D_b = D_b(S; \beta), \ b \in B,$$

uniquely determines the time allocation decisions of the household under behavioral specification b, which belongs to the set B of four household behavioral specifications we consider in this paper. The identification and estimation problems relate to our ability to recover the parameters that characterize a particular mapping D_b.

In terms of the econometrics of the problem, identification and estimator implementation will depend on assumptions regarding the observability of the elements of S. Given the data at hand, we consider the subvector $S_1 = (w_1 \, w_2 \, Y)'$ observable for all households. Since wages are only observed for working spouses, clearly this implies that our sample contains only dual-earner households. This restriction results in us losing about 12 percent of our sample, with the benefit of making the identification conditions for the model considerably more transparent. We will comment further on this assumption below.

The subvector $S_2 = (\lambda_1\ \lambda_2\ \delta_1\ \delta_2\ \alpha)$ contains the unobservable variables to the analyst. In our parametric estimation of the model, we specify a population distribution of S_2, the parameters of which, in addition to those describing the distribution of β in the population, constitute the primitive parameters of the model. We allow considerable flexibility in our parametric specification of the joint distribution of $(\lambda_1\ \lambda_2\ \delta_1\ \delta_2)$ through the following procedure. Let x be a four-variate normal vector, with

$$x \overset{i.i.d.}{\sim} N(\mu, \Sigma), \tag{8}$$

with μ a 4×1 vector of means and Σ a 4×4 symmetric, positive definite matrix.

A draw from this distribution, x, is mapped into the appropriate state space through the vector of known functions, M (which is 4×1). In our case, we have the following specification of the "link" function,

$$\begin{array}{ll} \lambda_1: & M_1(x) = \text{logit}(x_1) \\ \lambda_2: & M_2(x) = \text{logit}(x_2) \\ \delta_1: & M_3(x) = \frac{\exp(x_3)}{1+\exp(x_3)+\exp(x_4)} \\ \delta_2: & M_4(x) = \frac{\exp(x_4)}{1+\exp(x_3)+\exp(x_4)} \end{array}.$$

Thus, the joint distribution of these 4 household characteristics is described by a total of 14 parameters, 4 from μ and the 10 nonredundant parameters in Σ.[5]

The two other parameters upon which (some of) the model solutions depend are the Pareto weight parameter, α, and the discount factor β. As is well-known from the collective household model literature, estimation of the Pareto weight α is not possible without auxiliary functional form assumptions and/or exclusion restrictions. While our functional form assumptions in principle allow us to identify and estimate a constant value of α within the various model specifications in which it appears, in practice identification of this parameter is problematic, particularly in the EI model which also contains the free parameter β.[6] As a result, we restrict its value to $\alpha = 0.5$ in all of the estimation performed

[5] When estimating Σ, it is necessary to choose a parameterization that ensures that any estimate $\hat{\Sigma}$ is symmetric, positive definite. The most straightforward way of doing so is to use the Cholesky decomposition of Σ. There are 10 parameters to estimate, with

$$C = \begin{bmatrix} \exp(c_1) & c_2 & c_3 & c_4 \\ 0 & \exp(c_5) & c_6 & c_7 \\ 0 & 0 & \exp(c_8) & c_9 \\ 0 & 0 & 0 & \exp(c_{10}) \end{bmatrix},$$

and $\Sigma(c) = C'C$. The $\exp(\cdot)$ functions on the diagonal ensure that each of these elements are strictly positive, which is a requirement for the matrix to be positive definite (see Pinheiro and Bates, 1996).

[6] As we will show below, α is not nonparametrically identified in the PO and CPO specifications, whereas the "unrestricted" population of (other) state variables is. By restricting the distribution of state variables to be of a particular parametric form, α may be identified, but its identification hinges on the particular parametric assumptions made regarding the distribution of the other state variables. Even given these

below. This is not as a severe of a restriction as it appears, since in the Constrained Pareto Optimal and Endogenous Interaction models the side constraints that the efficient solution is required to satisfy produce, in general, a nondegenerate distribution of "*ex post*" values of α, . even if the "notational" value of α (α_0) is equal to 0.5 for all households. As we will see in the results reported below, most households end up using a value of α not equal to the notional value of 0.5 in the CPO and EI specifications.

It is possible to allow for variability of β in the population, though for a variety of reasons we have decided against this. Perhaps the key ingredient for forming a "cooperative" marriage is that both spouses be sufficiently forward-looking. Though we believe that allowing heterogeneous β within and across households should be an important goal in extending existing models of marriage market equilibrium, we sidestep sorting (on β) issues here by assuming a common value in the population. This restriction also allows us to draw some sharp conclusions when we compare the fits of the CPO and the EI models to the data.

4.1 Simulation-Based Estimation

Let the parameter vector of the model be given by $\Omega = (\mu' \ \text{vec}(\Sigma)' \ \omega)'$, where $\text{vec}(\Sigma)$ is a column vector containing all of the nonredundant parameters in Σ, ω is empty in the NE specification, contains α in the PO and CPO specifications, or contains (α, β) in the EI specification, so that Ω is a 16×1 vector in the most "heterogeneous" model we consider. We have access to a sample of married households taken from the Panel Study of Income Dynamics (PSID) from the 2005 wave, which we consider to a random sample from the population of married households in the U.S. within a given age range. In terms of the observable information available to us, we see the decision variables for household i,

$$A_i = (h_{1,i} \ \tau_{1,i} \ h_{2,i} \ \tau_{2,i}),$$

and we see the state variables

$$S_{1,i} = (w_{1,i} \ w_{2,i} \ Y_i).$$

Define the union of these two vectors, which is the vector of all of the observable variables of the analysis, by $Q_i = (A_i \ S_{1,i})$, so that the $N \times 7$ data matrix is

$$Q = \begin{bmatrix} Q_1 \\ Q_2 \\ \vdots \\ Q_N \end{bmatrix}.$$

restrictions, we often found that estimates of α tended to the boundary of the parameter space, which meant, in the case of CPO, that all outcomes where determined by the participation constraints. In such a case, the *ex ante* value of α was locally not identified. These experiences led us to the prudent choice of fixing the *ex ante* value at a point that might be consistent with social norms.

We choose m characteristics of the empirical distribution of Q upon which to base our estimator. Denote the values of these characteristics by the $m \times 1$ vector Z.

Simulation proceeds as follows. For each of the N households in the analysis, we draw NR values of x, and also β when it is included in a model specification and allowed to be heterogeneous. For simplicity, we will consider the set of simulation draws as generating $(\lambda_1 \ \lambda_2 \ \delta_1 \ \delta_2 \ \beta)$, even when β is treated as fixed in the population. Let a given simulation draw of these state variables be given by $\theta_{i,j},\ i = 1, ..., N;\ j = 1, ..., NR$. The draws $\theta_{i,j}$ are functions of the parameter vector Ω, which we emphasize by writing $\theta_{i,j}(\Omega)$. Given a value of Ω, for each household i we solve for household decisions under behavioral mode b, $(a^b_{1,i,j}, a^b_{2,i,j})(S_{1,i}, \theta_{i,j}) = D_b(S_{1,i}, \theta_{i,j}),\ j = 1, ..., NR$, where $a^b_{s,i,j}$ is the market labor supply and time in household production of spouse s in household i given draws $\theta_{i,j}$ under behavioral regime b. The time allocations associated with the simulation are stacked in a new matrix

$$\tilde{Q}_b(\Omega) = \begin{bmatrix} A^b_{1,1}(\Omega) & S_{1,1} \\ \vdots & \vdots \\ A^b_{1,NR}(\Omega) & S_{1,1} \\ A^b_{2,1}(\Omega) & S_{1,2} \\ \vdots & \vdots \\ A^b_{2,NR}(\Omega) & S_{1,2} \\ \vdots & \vdots \\ A_{N,1}(\Omega) & S_{1,N} \\ \vdots & \vdots \\ A_{N,NR}(\Omega) & S_{1,N} \end{bmatrix},$$

where $A^b_{i,j}(\Omega) = (a^b_{1,i,j} \ a^b_{2,i,j})$. The analogous value of $Z_b(\Omega)$ is computed from the $(N \times NR) \times 7$ matrix $\tilde{Q}_b(\Omega)$. Given a positive-definite, conformable weighting matrix W, the estimator of Ω is given by

$$\hat{\Omega}_b = \arg\min_{\Omega} (Z - Z_b(\Omega))' W (Z - Z_b(\Omega)).$$

The weighting matrix W is computed by resampling the original data Q a total of 5000 times, and for each resampling we compute the value of Z, which we denote $\hat{Z}_k$ for replication k. We then form

$$\hat{Z} = \begin{bmatrix} \hat{Z}_1 \\ \hat{Z}_2 \\ \vdots \\ \hat{Z}_{5000} \end{bmatrix}.$$

W is the covariance matrix of $\hat{Z}$. Since a major focus of our empirical investigation is a comparison of the behavioral models in B in terms of their ability to fit the sample moments Z, it is advantageous to utilize a weighting matrix W that is not model dependent.

4.2 Identification

As is true of many (or most) simulation-based estimators, especially those used to estimate relatively complex behavioral models, providing exact conditions for identification is not feasible. Nevertheless, it may be useful to understand what features of the data generating process are being used to obtain point estimates of the parameters in our 'flexible' parametric model of the household. With this goal in mind, we proceed through a fairly careful consideration of nonparametric identification of the first three behavioral specifications. We will then discuss the reasons that the Endogenous Interaction model is not nonparametrically identified, which is the primary reason for our interest in the estimation of a flexible parametric specification of the distribution of primitive parameters.

The identification arguments we present in this section condition on observed wages, and do not allow for measurement error in any of the variables included in Q_i, which contains the conditioning variables $S_{1,i} = (w_{1,i}\ w_{2,1}\ Y_i)$, as well as the four time allocation measures $A_i = (h_{1,i}\ \tau_{1,i}\ h_{2,i}\ \tau_{2,i})$. We begin by considering the nonparametric identification case. We assume that there exists a joint distribution of $F_S(s)$, with the vector $S = (w_1\ w_2\ Y\ \lambda_1\ \lambda_2\ \delta_1\ \delta_2\ \alpha\ \beta)'$ in the most general model. An individual household in the PSID subsample is considered to be an i.i.d. draw from the distribution F_S. No parametric assumptions on F_S are made, at this point.

4.2.1 All Households Behave Inefficiently (Static Nash equilibrium)

In the case of Nash equilibrium, it is straightforward to show that the model is nonparametrically identified in the sense that we can define a nonparametric, maximum likelihood estimator (NPMLE) of F_S in the following (constructive) manner.

Proposition 4 *The distribution F_S is nonparametrically identified from Q when the behavioral rule is static Nash equilibrium and there are no corner solutions.*

Proof. The Nash equilibrium is the unique fixed point of the reaction functions of spouse 1 and spouse 2 given the time allocations of the other spouse. Given that both spouses work, we observe the vector $(w_{1,j}\ w_{2,j}, Y_j)$ for all households, so that the marginal distribution F_{S_1} is nonparametrically identified by construction. Given the observation Q_j, we can invert the reaction functions for household j to yield the two equation linear system

$$\begin{aligned}\delta_{1,j} &= \frac{B_{1,j}(1-B_{2,j})}{1-B_1,jB_{2,j}}\\ \delta_{2,j} &= \frac{B_{2,j}(1-B_{1,j})}{1-B_1,jB_{2,j}},\end{aligned}$$

where $B_{i,j} = w_{i,j}\tau_{i,j}/(M_j + w_{i,j}\tau_{i,j})$, with $M_j = w_{1,j}h_{1,j} + w_{2,j}h_{2,j} + Y_j$. Given constant returns to scale in household production, $\delta_{3,j} = 1 - \delta_{1,j} - \delta_{2,j}$, and we can invert the remaining two equations in the system of reaction functions to obtain

$$\lambda_{i,j} = \frac{\delta_{3,j} w_{i,j}(T - h_{i,j} - \tau_{i,j})}{M_j + \delta_{3,j} w_{i,j}(T - h_{i,j} - \tau_{i,j})},\ \ i = 1, 2.$$

Since these values of $(\delta_{1,j}, \delta_{2,j}, \alpha_{1,j}, \alpha_{2,j})$ are uniquely determined by $(h_{1,j}, \tau_{1,j}, h_{2,j}, \tau_{2,j}, w_{1,j}, w_{2,j}, Y_j)$, we "observe" the complete vector of values $S_j = (w_{1,j}\ w_{2,j}\ Y_j\ \lambda_{1,j}\ \lambda_{2,j}\ \delta_{1,j}\ \delta_{2,j})$, and the nonparametric maximum likelihood estimator of F_S is the empirical distribution of $\{S_j\}_{j=1}^{N}$.
■

Note that the restriction of no corner solutions is essential in our ability to nonparametrically identify the model. Say that, for example, $h_{1,j} = 0$. Even if the offered wage $w_{1,j}$ were available, an unlikely event, there would exist a set of values of $\lambda_{1,j}$ consistent with the observed choices and the observed state variables, with no way to assess the likelihood of any value in the set relative to that of any other. In the presence of any kind of truncation or censoring, nonparametric point identification of the complete distribution is typically impossible, since some functional form assumptions are required to assign likelihoods over sets of values of parameters consistent with observed outcomes.

4.2.2 All Households Behave Efficiently

Proposition 5 *The distribution F_S is nonparametrically identified from Q when the behavioral rule is Pareto efficiency, there are no corner solutions, data points are consistent with the model, and α is known.*

Proof. Household time allocation is determined by solving the system of four first order conditions associated with (7). We find that

$$\begin{aligned}\delta_{1j} &= \frac{B_{1j}(1 - B_{2,j})}{1 - B_{1j}B_{2,j}} \\ \delta_{2j} &= \frac{B_{2j}(1 - B_{1,j})}{1 - B_{1j}B_{2,j}},\end{aligned}$$

the same as in the Nash equilibrium case. Conditional on these values of δ_{1j} and δ_{2j} ($\Rightarrow$ $\delta_{3j} = 1 - \delta_{1j} - \delta_{2j}$ under the CRS assumption) and a value of α, we find

$$\begin{aligned}\lambda_{1j} &= \frac{1}{\alpha} R_{1j} \\ \lambda_{2j} &= \frac{1}{1-\alpha} R_{2j}\end{aligned}$$

where

$$\begin{aligned}R_{1j} &\equiv \frac{C_{1j}(1 - C_{2j})}{1 - C_{1j}C_{2j}} \\ R_{2j} &= \frac{C_{2j}(1 - C_{1j})}{1 - C_{1j}C_{2j}}.\end{aligned}$$

and

$$\begin{aligned} C_{1j} &= \frac{\delta_{3j} w_{1j}(T - h_{1j} - \tau_{1j})}{\delta_{3j} w_{1j}(T - h_{1j} - \tau_{1j}) + M_j} \\ C_{2j} &= \frac{\delta_{3j} w_{1j}(T - h_{2j} - \tau_{2j})}{\delta_{3j} w_{2j}(T - h_{2j} - \tau_{2j}) + M_j}. \end{aligned}$$

The values of C_{1j} and C_{2j} lie in the unit interval for all j, which implies that λ_{1j} and λ_{2j} are always positive. However, for given values of α and all other state variables and choice variables, either or both λ_{1j} and λ_{2j} may not belong to the open unit interval. In this case, the data for household j are not consistent with the model and are not used in the estimation of F_S, which is the empirical distribution of $\{S_j\}_{j \in \varkappa}$, where $\varkappa$ is the set of household indices for which λ_{1j} and λ_{2j} both belong to the open unit interval. ■

We found that the implied values of λ_{1j} and λ_{2j} belonged to the unit interval for all of the sample cases. In the proposition, in constructing the NPMLE for F_S we specified that only implied values of the preference parameters that satisfied our theoretical restrictions would be utilized. One could argue that the satisfaction of theoretical restrictions should be treated as a necessary condition for defining the estimator, instead. In this particular application, we did not have to explicitly confront this problem.

4.2.3 Constrained Efficient Case

The constrained efficient case imposes a side constraint on the efficient solution, one which insures that each spouse attains a utility value at least as large as what could be obtained in the inefficient, Nash equilibrium case. This makes the mapping from the observed time allocations and observed state variables into the unobserved state variables more complex. However, we are able to establish the following result

Proposition 6 *The distribution F_S is nonparametrically identified from Q when the behavioral rule is Constrained Pareto efficiency, there are no corner solutions, data points are consistent with the model, and α is known.*

Proof. As we showed in the proof of the nonparametric identification of the Pareto weight case, inverting the system of first order conditions for household j yields solutions for δ_{1j} and δ_{2j} which are not functions of α, while

$$\begin{aligned} \lambda_{1j} &= \frac{R_{1j}}{\alpha} \\ \lambda_{2j} &= \frac{R_{2j}}{1-\alpha}, \end{aligned}$$

where R_{ij} is a positive-valued function of the data, $i = 1, 2$. A feasible value of α for household j must satisfy the conditions

$$U_{ij}^P(\alpha) \geq U_{ij}^N, \; i = 1, 2,$$

where $U_{ij}^P(\alpha)$ is individual $i's$ utility at the efficient solution when the Pareto weight is α and U_{ij}^N is that individual's utility in the inefficient equilibrium. This condition can be written

$$D_{ij}(\alpha) = \lambda_i(\ln l_{ij}^P - \ln l_{ij}^N) + (1-\lambda_i)(\ln K_j^{PW} - \ln K_j^N) \geq 0, \;\; i = 1, 2.$$

Now, conditional on δ_{1j} and δ_{2j}, which are not a function of α, and R_{1j} and R_{2j}, any given value of α implies values of λ_{1j} and λ_{2j}, with

$$\begin{aligned}
\frac{\partial \lambda_{1j}}{\partial \alpha} &= -\frac{R_{1j}}{\alpha^2} < 0 \\
\frac{\partial \lambda_{2j}}{\partial \alpha} &= \frac{R_{2j}}{(1-\alpha)^2} > 0.
\end{aligned}$$

Partially differentiating D_{ij} with respect to α, we have

$$\begin{aligned}
\frac{\partial(U_{ij}^P(\alpha) - U_{ij}^N)}{\partial \alpha} = & \frac{\partial \lambda_{ij}}{\partial \alpha}(\ln l_{ij}^P - \ln l_{ij}^N) - \frac{\lambda_{ij}}{l_{ij}^N}\{\frac{\partial l_{ij}^N}{\partial \lambda_{1j}}\frac{\partial \lambda_{1j}}{\partial \alpha} + \frac{\partial l_{ij}^N}{\partial \lambda_{2j}}\frac{\partial \lambda_{2j}}{\partial \alpha}\} \\
& - \frac{\partial \lambda_{ij}}{\partial \alpha}(\ln K_j^P - \ln K_j^N) - \frac{1-\lambda_{ij}}{K_j^N}\{\frac{\partial K_j^N}{\partial \lambda_{1j}}\frac{\partial \lambda_{1j}}{\partial \alpha} + \frac{\partial K_j^N}{\partial \lambda_{2j}}\frac{\partial \lambda_{2j}}{\partial \alpha}\}, \;\; i = 1, 2.
\end{aligned}$$

Note that

$$\left.\frac{\partial U_{ij}^P(\alpha)}{\partial \alpha}\right|_{R_{1j}, R_{2j}, B_{i1}, B_{2j}} = 0,$$

since the first order conditions for the Pareto weight problem are characterized in terms of these values, not the preference and production parameters and α individually. This is the reason that only partial derivatives involving the Nash equilibrium values of household choices are present. For any values of $(\lambda_1, \lambda_2, \delta_1, \delta_2, w_1, w_2, Y)$, $l_{ij}^P \leq l_{ij}^N$, $i = 1, 2$, and $K_j^P \geq K_j^N$. Consider individual 1, for example. In this case,

$$\begin{aligned}
\frac{\partial l_{1j}^N}{\partial \lambda_{1j}} &\geq 0 \\
\frac{\partial l_{ij}^N}{\partial \lambda_{2j}} &\leq 0,
\end{aligned}$$

while

$$\begin{aligned}
\frac{\partial K_j^N}{\partial \lambda_{1j}} &\leq 0 \\
\frac{\partial K_j^N}{\partial \lambda_{2j}} &\leq 0,
\end{aligned}$$

which implies that

$$\frac{\partial D_{1j}(\alpha)}{\partial \alpha} \geq 0.$$

A similar argument yields

$$\frac{\partial D_{2j}(\alpha)}{\partial \alpha} \leq 0.$$

Finally, it must be the case that if one individual is indifferent at the efficient solution, the other individual must be strictly better off. Thus $D_{1j} = 0 \Rightarrow D_{2j} > 0$ and $D_{2j} = 0 \Rightarrow D_{1j} > 0$. Then D_{1j} and D_{2j} intersect at a value of α, $\hat{\alpha}$ say, where $D_{1j}(\hat{\alpha}) > 0$ and $D_{2j}(\hat{\alpha}) > 0$. The set of all feasible values of α is connected, and is given by

$$\alpha_j^F = [\underline{\alpha}_j, \overline{\alpha}_j],$$

where $\underline{\alpha}_j$ is defined by $D_{1j}(\underline{\alpha}_j) = 0$ and $\overline{\alpha}_j$ is defined by $D_{2j}(\overline{\alpha}_j) = 0$. The values of all primitive parameters, given the known, notional Pareto weight value α_0, are determined from the *ex post* value of α, defined as

$$\tilde{\alpha} = \begin{cases} \underline{\alpha}_j & \text{if} \quad \alpha_0 < \underline{\alpha}_j \\ \alpha_0 & \text{if} \quad \alpha_0 \in \alpha^F \\ \overline{\alpha}_j & \text{if} \quad \alpha_0 > \overline{\alpha}_j \end{cases} .$$

The "data" for household j are given by the observed state variables w_{ij}, w_{2j}, and Y_j, and the primitive parameters $\lambda_{1j}(\tilde{\alpha}) = R_{1j}/\tilde{\alpha}$, $\lambda_{2j}(\tilde{\alpha}) = R_{2j}/(1-\tilde{\alpha})$, δ_{1j}, and δ_{2j}. Under our random sampling assumption, the empirical c.d.f. of these variables is the nonparametric maximum likelihood estimator of F_S. ■

4.2.4 Endogenous Interaction Case

We cannot show that F_S is nonparametrically identified in the EI case; in fact, it is straightforward to provide counterexamples to show that it is not. We can continue to assume that the technology parameters are uniquely determined without reference to the behavioral regime or value of α used in efficient cases. The fundamental identification problems then concern the discount factor β and the preference parameters, λ_1 and λ_2.

Since we could not identify the notional value of α_0 nonparametrically, it comes as no surprise that the same is true of β. We know that if $\beta = 0$, no efficient allocations can be supported, and the resulting time allocations are all generated in NE. Conversely, as $\beta \rightarrow 1$, all households make efficient time allocation decisions, and the preference parameters implied by the data are those generated under CPO. Both sets of values of the preference parameters are one-to-one mappings from the data, so we have two separate estimators of F_S.

Unfortunately, the lack of identification result continues to hold even after fixing β at some predetermined value β_0. Having a notional value of α, α_0, and a fixed value of β, β_0, we can find cases of observed state variable vectors, S_1, and decisions, A, that yield two valid implied values of (λ_1, λ_2), one under static Nash equilibrium and one other dynamic efficiency. By "valid" we mean that the implied values of preferences and technology assuming static Nash equilibrium are such that no efficient allocation is implementable given β_0. At the same time, using implied values of preferences and technology assuming an efficient allocation, we can determine values of the preference parameters that imply the existence of an implementable solution in the sense of satisfying the long-run participation constraint. Obviously, the two sets of preference parameters are not identical, and there is no way to differentiate between them when forming a NPMLE of F_S.

4.2.5 The Flexible Parametric Case

Under the parametric specification of the unobserved state variables described at the beginning of this Section, the estimation problem becomes one of estimating a set of parameters θ assumed to completely characterize the distribution $F_S(\theta)$, instead of the function F_S itself. In the case of the NE behavioral specification we showed that F_S was nonparametrically identified, so that a (nonredundantly) parameterized F_S will also be identified. In the case of the PO and CPO behavioral specifications, the parameter vector θ is clearly identified given a known value of α since we showed that F_S itself was nonparametrically identified when α was known. In this case, using the nonparametric MLE for F_S, $\hat{F}_S$, we can define estimators for $\hat{\theta}_M = \arg\min_\theta M(\hat{F}_S, F_S(\theta))$ for some distance function M, and the argument for the identification of θ will be dependent on properties of M.

Under parametric assumptions on F_S, it becomes theoretically possible to estimate α depending on the restrictions imposed on F_S. For the reasons given above, we chose not to attempt to estimate the (*ex ante*) value of α, instead fixing it at 0.5 in all specifcations. This restriction ensures identification of θ in the PO, CPO, and EI specifications.

In terms of the EI specification, the flexible parametric specification aids in overcoming the identification problems associated with this case by smoothing the density over regions in which unique solutions to the inversion problem associated with the nonparametric estimator do not exist. While we cannot establish formal identification of θ using the MSM estimator, the parameter estimates we obtain are reasonable and bootstrap standard error estimates imply an acceptable level of precision.

There are a some additional identification issues associated with the EI specification that are not present in the PO and CPO versions of the model. Letting β_0 denote the true value of β in the population, we saw that as $\beta \to 1$, all outcomes are identical to the CPO outcome associated with a notional Pareto weight of α_0. Furthermore, as $\beta_0 \to 0$, all outcomes are produced under static Nash equilibrium. Then when all population members are myopic, i.e., $\beta_0 = 0$, all outcomes are Nash equilibrium outcomes, and α_0 is not identified. For the household's decision problem to be well-defined, we must restrict β_0 to be strictly less than one. However, assume that the distribution of S is such that the probability of any draw of a state vector from this distribution has probability zero of resulting in a static Nash equilibrium outcome for any value of $\beta \in (1-\varepsilon, 1)$, where ε is a small positive constant. Then if $1-\varepsilon < \beta_0 < 1$, all observed outcomes will be associated with the CPO solution, α_0 will be identified, but β_0 will not be point-identified. Thus, we have two extreme cases of local nonidentification in the EI specification. When β_0 is 0, or abitrarily close to zero, α_0 is not identified. When β_0 is arbitrarily close to 1, β_0 is not point identified, but α_0 is identified. When β_0 assumes "intermediate" values, so that the measure of noncooperative households in the population is not equal to 0 or 1, then both are, in principle at least, identified under parametric restrictions on F_S. In the estimates of the EI specification reported below, we find that the estimate of β is one which produces a proper mixture distribution, i.e., the estimates imply that a measurable set of households are in the cooperative and noncooperative regimes.

5 Empirical Results

We begin this section by presenting the sample selection criteria used in creating the final sample from the PSID with which we work. This is followed by a discussion of the estimates of the distributions of primitive parameters in the "nonparametric" analysis for the three identified models: Nash equilibrium, Pareto efficiency, and Constrained Pareto efficiency. We then move on to our focus of interest, which are the estimates from the flexible parametric analysis.

5.1 Sample Selection Criteria and Descriptive Statistics

We use sample information from the 2005 wave of the PSID. All models are essentially static, and therefore we only utilize cross-sectional information from this wave of the survey. We only considered households in which the head was married, with the spouse present in the household. In this wave, the PSID obtained the standard information regarding usual hours of work over the previous year for both spouses, and this information corresponds to h_1 and h_2 in the model. Every few years, the PSID also includes a question regarding the usual hours devoted to housework by each spouse, and this information is included in the 2005 wave. The responses to these items are interpreted as τ_1 and τ_2 in the model. These time allocation questions are regarded as referring to the same time period. We use total hours worked from the previous year and labor earnings for each spouse to infer a wage rate, w_i for spouse i. In addition, information is available on the nonlabor income of the spouses over the previous year, and we divide this amount by 52 to obtain a weekly nonlabor income level, Y.

We only utilize information from households in which both spouses are between the ages of 30 and 49, inclusive. In addition to this age requirement, we excluded all households with any child less than 7 years of age, since the household production function is likely to be far different when small children are present. We also excluded couples with what we considered to be excessively high time allocations to housework and the labor market, namely, those with over 100 hours combined in these two activities. We selected this amount since we set $T = 112$, which we arrived at by assuming 16 hours to allocate to leisure, housework, and the labor market for each of the seven days in a week. We also excluded households reporting a nonlabor income level of more than \$1000 a week, on the grounds that such people were likely to be generating a significant amount of self-employment income, making the labor supply information they supplied difficult to interpret. Not many households were lost to this exclusion criterion.

By far the most significant sample selection criterion we imposed was the one requiring both spouses to work. If a spouse does not work, then clearly we have no wage information for that spouse, making the nonparametric analysis we discuss above and report on below impossible to implement. Within the flexible parametric estimator we implement, it would be possible to allow for corner solutions in labor supply if we are willing to impose a parametric assumption regarding the wage offer process.[7] While this allows for more model generality, in principle, it comes at the expense of having to take a position on the partially unobservable wage process. We chose to follow the route of ruling out corner solutions, allowing us to condition all of our analysis on observed wages. Because we restricted our attention to households without small children, imposing the condition that both spouses supply time to the market resulted in a reduction of 12 percent in our (otherwise) final sample. We were left with 823 valid cases, which were those satisfying the conditions stated above and with no missing data on any of the state and decision variables included in the analysis.

A description of the decisions and state variables is contained in Table 1. As has been often remarked upon in other analyses of household behavior that include time in housework, the average time spent in the both housework and labor supply to the market is very similar for husbands and wives. On average, husbands spend approximately 7 hours more per week in the labor market than their (working) wives, but devote 7 hours less to housework. Under our assumption that each spouse has 112 'disposable' hours of time to allocate each week, on average each spouse spends slightly more than one-half of their time consuming leisure. We also note that the wives' distributions of hours in the market and housework are much more disperse than the corresponding distributions of husbands. In terms of market work, this is undoubtedly due to the fact that married women are much more likely to be employed in part-time work than their husbands (see, e.g., Mabli (2007)). The limited amount of variation in the distribution of husbands' housework is mainly due to the floor effect - most observations are clustered in the neighborhood of zero.

In terms of the observed state variables of the analysis, the mean wage of husbands is approximately 39 percent greater than the mean wage of wives, and exhibits considerably more dispersion, some of it due to the presence of a few wage outliers among the husbands (the maximum wage of wives is \$80.50, while the maximum wage of husbands is \$144.93).

[7] This was precisely what was done in an earlier version of this paper, when we only considered labor supply in a model without a household production component.

Average weekly nonlabor income of the household is \$118.15, and this distribution is quite disperse, even though the sample is restricted to households receiving no more than \$1000 of nonlabor income per week. No nonlabor income is reported by 27 percent of sample households.

Table 2 contains the zero-order correlation matrix of the variables reported in Table 1. There is no correlation between the labor supply and housework of husbands, while there is a reasonably strong negative correlation (-0.189) between them for wives. There is a strong positive correlation (0.321) between the times spent in housework by husbands and wives. The wage of a husband and the labor supply of his wife have a negative correlation of -0.132, while wives with high wages tend to spend less time in housework. There are no particularly noteworthy correlations between household nonlabor income and other variables in the analysis, with the possible exception of the husband's wage (0.115). The correlation between the wages of the spouses (0.294) indicates positive assortative mating in the marriage market.

5.2 Nonparametric Estimation of the Distribution of State Variables

Under the Nash equilibrium, Pareto efficient, and constrained Pareto efficient modeling assumptions, we were able to obtain estimates of the distributions of S in our sample. In all cases other than static Nash equilibrium, we established that the Pareto weight parameter α was not identified. Accordingly, in all of these models, we simply assume that the 'notional' Pareto weight is 0.5. Of course, the Nash equilibrium solution is not a function of the parameter α.

In Section 4.2.2, we noted that the mapping from the time allocation decisions and the observed state variables, the wages of the spouses and household nonlabor income, did not necessarily produce values of the preference parameters λ_1 and λ_2 contained in $(0, 1)$. Nevertheless, all of our 823 sample cases generated values of λ_1 and λ_2 in the unit interval, so that all cases are used to generate 'data' on preferences and household production parameters that are used to form the nonparametric estimator of F_S.

Table 3 contains estimates of the means and standard deviations of the marginal distributions of preference and production parameters of the model under the three estimable behavioral specifications. As discussed above, under our functional form assumptions on preferences and household production, the implied value of the production parameters δ_{1j} and δ_{2j} for household j are the same functions of the decisions of household j, D_j, and the observed state variables for household j, S_{1j}, for each of the three behavioral models for which we obtain nonparametric estimates of F_S. This explains the fact that the estimated means and standard deviations of the production parameters are identical across the three behavioral specifications. We note that wives have a higher average productivity in household production than husbands, with the mean for wives being about 41 percent larger. There is also slightly more dispersion in the wives' productivity parameter.

Large differences are observed across the three specifications in terms of the distribution of the preference parameters. Given that the Nash equilibrium outcomes are inefficient, it is not surprising to find that the means of the preference parameters under Nash equilibrium are considerably less than they are under constrained or unconstrained Pareto efficiency. In all three behavioral cases, the average weight placed on the private good, leisure, is smaller for wives than their husbands. In the unconstrained Pareto weight case, the average weight placed on leisure by husbands is 0.580, in comparison with an average leisure weight of 0.430 for wives. There are similar levels of dispersion in the distributions of preference parameters for husbands and wives across the three behavioral specifications.

Comparing estimates across columns two and three, it is interesting to note that the constraint that the payoffs under the efficient solution are at least as large as the payoffs under Nash equilibrium for both spouses is binding for a number of sample cases given the notional value of $\alpha = 0.5$. This is evidenced by the differences in the preference parameter distributions. Imposing this particular constraint narrows the difference in the mean of spousal preference parameters, while reducing dispersion as well.

Recall that all three estimates of F_S, are equally "valid," and no statistical criterion can be used to distinguish between the behavioral specifications given that they are all based on (different) one-to-one mappings from the data and observed state space to the unobserved state space. In the next section, when we make flexible parametric assumptions regarding the distributions of the parameters, we will be able to compare the performance of the various behavioral models, including the Endogenous Interaction specification.

5.3 Parametric Estimation of the Distribution of State Variables

Before looking at the estimates produced by the parametric estimator under the four behavioral specifications, it may be worthwhile to consider why we expect them to differ to some degree from the nonparametric estimators of F_S discussed in the preceding subsection. First, we have assumed that the distribution of the state variables (subvector) $S_2 = (\lambda_1\ \lambda_2\ \delta_1\ \delta_2)$ is independent of the state variables $S_1 = (w_1\ w_2\ Y)$. This is a strong assumption, but without specifying some form of parametric dependence between S_2 and S_1, it would not be possible to relax it. There are reasons to doubt the validity of the independence assumption. For example, a spouse i with a low value of leisure might have worked and invested more in the past, so that w_i and Y may be negatively related to λ_i. To fully account for such dependencies, we would require a life cycle household model with capital accumulation, which is beyond the scope of the current paper.

Second, while our parametric specification of the distribution of S_2 is reasonably flexible, it does impose restrictions on the data. These restrictions are what allow us to say something about the relative abilities of the four different behavioral specifications to fit the data. Nevertheless, different parametric specifications of the distribution of S_2 could lead to different inferences concerning which behavioral framework is most consistent with the data features chosen for the MSM estimator.

In implementing any moments-based estimator, the analyst faces the problem of choosing which moments to utilize in defining the estimator. We have used a set of 23 moments throughout, and never made any change to this set in the course of the estimation process. The list of moments appears in Table A.1, and is reasonably standard. It includes the means and standard deviations of all of the endogenous variables, as well as some cross-products terms involving the endogenous variables, and some involving the observed state variables of the household, wages and nonlabor income. The last few moments listed in the table pick up some distributional features of the labor market hours distributions not captured solely by the first two moments of the distribution.

Table 4 contains the MSM estimates of the four behavioral specifications. The estimates presented were computed as follows. We obtained point estimates of the 14 parameters used to characterize the distribution of S_2 for each of the four specifications. We then took a large number of draws (one million) from the estimated distribution of S_2, and computed the means and standard deviations of each of the components of the vector S_2. In the EI specification, we also estimated the discount factor β, which was constrained to be homogeneous in the population. For the CPO and EI specifications, we also computed the (*ex post*) Pareto weight at which the efficient outcome was implemented for each draw in the simulation process. The means and standard deviations of the *ex post* α distribution are presented for the CPO and EI specifications..The last row in the table reports the value of the distance metric for the model; obviously, a lower value indicates that the model is able to better fit the selected moments at the optimally-chosen parameter estimates.

The Nash equilibrium specification produces estimates of the mean values of the preference parameters, λ_1 and λ_2, roughly in accord with those produced by the nonparametric estimator, though the mean value of the preference weight on leisure of the wife is almost equal to that of the husband. While the mean value of λ_1 is similar for the two estimators, the estimated dispersion in λ_1 is about 50 percent greater when the parametric estimator is used. The estimated mean of λ_2 is significantly higher in the parametric case, as is the estimated dispersion in this parameter. The estimated distributions of production function parameters are quite a bit different under the flexible parametric estimator and the nonparametric one. We still find that, on average, wives are more productive in housework than their husbands, but the difference in the means is now a bit larger than was the case using the nonparametric estimator (0.068 versus 0.046). The parametric estimator produces a lower amount of dispersion in δ_1 and δ_2 than does the nonparametric estimator. Our sense is that a substantial amount of these differences are due to the restriction that the state variables in S_2 are independently distributed with respect to S_1 that is imposed using the parametric estimator and that is not imposed under the nonparametric estimator, rather than arising from the parametric restrictions on the distribution of S_2, *per se*.[8]

The first two moments of S_2 estimated under the assumption of Pareto efficiency are

[8] We performed some experimentation with making the mean vector of the 4 variate normal that generates the primitive parameters a linear function of the 3 observed state variables w_1, w_2, and Y.We concluded that it was too difficult to precisely estimate this conditional mean specification, and so stayed with the specification that assumes independence between S_1 and S_2.

presented in the third column. As regards the preference parameter distributions, we see that the mean estimated leisure weights for both spouses are considerably larger than we found in the Nash equilibrium case in Table 4. The estimated mean λ_2 for wives is virtually identical using either the nonparametric or parametric estimator, while the estimated value for the husbands' mean is smaller using the nonparametric estimator. The estimated dispersion in λ_2 is notably larger using the parametric estimator. The estimated mean value of the production parameters is roughly similar to what was obtained using the nonparametric estimator, and there are no large differences in the estimates of these parameters from the noncooperative Nash equilibrium case, and the same is true regarding the amount of dispersion in these parameters.

Column 5 of the table contains the estimated parameter moments for the constrained efficient model. There are no large changes in the estimates of the first and second moments of the preference and production parameters from the PO case. The big difference between the CPO and PO specification estimates is in the *ex post* α distribution. In the PO case, the *ex post* Pareto weight is identical to the *ex ante* Pareto weight, the value of which is fixed at 0.5. In the CPO case, there is considerable dispersion in the distribution of the *ex post* value of α. We see that the mean value of α is 0.515, while the standard deviation is 0.042. Only in 41.5 percent of cases is the *ex post* value equal to the *ex ante* value. The distribution of the *ex post* value of α, after deleting the cases for which it is equal to 0.5, is shown in Figure 3.a. Most of the mass of the distribution lies to the right of 0.5, indicating that it is the husbands whose participation constraint is binding in the majority of cases at the *ex ante* value of 0.5.

We now turn to our focus of interest, the Endogenous Interaction specification. The EI specification is also based on a fixed, notional value of the Pareto weight of 0.5, but includes the discount factor, β, a parameter not included in the PO an CPO specifications. Opening up the possibility of cheating on the efficient outcome introduces a more stringent form of a participation constraint than the one that exists in the CPO specification. Perhaps the most interesting result reported in column 4 is the proportion of sample cases that achieve utility realizations that lie on the Pareto frontier, which we estimate to be 0.753. Approximately three-fourths of households do manage to implement efficient time allocations, however, only 7.4 percent of these efficient households utilize the notional Pareto weight of 0.5. The distribution of the *ex post* value of α, excluding *ex post* values of α equal to 0.5, is exhibited in Figure 3.b. The shape of this distribution is similar to the one shown in the panel above it, with a slight negative skew. The average value of *ex post* α among efficient households in the EI specification is 0.527. The dispersion in *ex post* α is about the same in the CPO and EI specifications.

In terms of the estimates of the two first moments of the marginal distributions of preference and production parameters, the EI estimates of the mean values of these parameters are similar to those from the CPO specification. Except for the case of δ_2, the estimated means of the parameters from the EI specification lie between those associated with the NE and the CPO models. The variance of the distribution of the preference parameters is less than under any of the other specifications estimated. This may be attributable to the fact that a lower level of heterogeneity in the primitive parameters is capable of generating the same amount of variability in household allocations given the additional variability produced by the two sets of decision rules (efficient and inefficient).

To induce any households to behave inefficiently, a relatively low value of β is required, and our point estimate is 0.524, which we interpret as referring to a yearly period, since the data refer to a representative week in 2004, and we think of participation decisions being made on a yearly basis. While the estimate of β is 'low,' it is not very much out of line with respect to other estimates of the subjective rate of discount found in the experimental and microeconomics literature (see, for example, Hausman (1979) and Thaler (1981)). The compilation of estimates of time preference performed by Frederick et al. (Table 1, 2002) is striking for the huge range of values of the subjective discount rate that have been found using both experimental and empirical methods. To our knowledge, this is the first application to attempt to use a formal model with a grim trigger strategy to estimate a discount factor, so there are no other studies with which we can directly compare our estimate.[9,10]

We now compare the model fits across the four specifications. The first three specifications contain the same number of free parameters, which are 14. The EI specification also includes the parameter β. Thus the fits of the first three models are very much directly comparable, since the parameter space is identical. We see that, under the restriction that the notional value of $\alpha = 0.5$, the fit of the PO specification is the worst, followed by the NE specification, and then CPO. The EI specification produces a marked increase in the ability of the model to fit the data features we have selected. Recall that the EI specification nests the NE and CPO models as special cases. As $\beta \to 0$, no efficient solutions could be supported, so all households would behave in an innefficient manner, with allocations given by the Nash equilibrium values. As $\beta \to 1$, all households will behave efficiently, with the only constraint on the allocations being that they satisfy the participation constraint, which imposes the restriction on α associated with the CPO specification. Moving β from a value of 1 (implicit in the CPO specification) to 0.524 results in a reduction in the distance metric of 14 percent. Moving β from a value of 0 (implicit in the NE specification) to 0.524 results in a reduction of almost 20 percent. The estimate of β we obtained suggests that the "cheating" problem is an important one in determining observed household time allocations.

In order to conduct a more proper "horse race" between the models, we need to compare

[9]Porter (1983) and Lee and Porter (1983) estimate a switching regressions model motivated by the trigger price strategy model of Green and Porter (1984). In that model of collusive behavior with imperfect signals regarding other agents' actions, a noncooperative punishment period is enterred whenever public signals indicate a high probability of cheating. The punishment period is determined endogenously, and at its termination another collusive regime is begun. The econometric framework used in the two empirical papers cited does not allow one to back out an estimate of the discount factor of firms.

[10]In addition, as one referee has noted, the discount factor may incorporate the "death" rate of marriages, in the simple case where we consider this to be exogenous and constant.

the values of the distance metrics under each of the resamplings we performed in computing the bootstrap standard errors for the model parameters. In constructing the bootstrap standard errors, we generated 25 replication samples and reestimated all of the models on each replication sample. Thus, in all, we estimated the models 26 times when we include the original data sample. The fitted values are displayed in Figure 4, where models 1.through 4 are NE, PO, CPO, and EI, respectively. As we can see, no single model clearly dominates all of the others across all of the replications, although it is apparent that the EI specification is superior to the PO specification in every case. The EI specification bests NE in almost always, and dominates the CPO specification in the majority of cases. If one were forced to choose one of the specifications on the basis of this evidence, the choice would clearly seem to be EI.

We conclude this section by describing Figures 5-8, which use the flexible parametric specification to plot bivariate relationships between production and preference parameters within and across spouses by household. In each case, we used the point estimates of the parameters that characterized the flexible multivariate distribution of $(\lambda_1\ \lambda_2\ \delta_1\ \delta_2)$, in conjunction with a large number of pseudo-random number draws from the underlying standard normal distribution, to generate pseudo-random number draws from the (estimated) joint distribution of the preference and technology parameters.

Figure 5.a contains the scatter plot of draws of λ_1 and λ_2 obtained from the Nash equilibrium specification. There is a strong, almost linear relationship between the preferences of the spouses over most of the support of the distribution, indicating a substantial degree of (positive) assortative mating with respect to preferences. The scatter plot of δ_1 and δ_2 under the Nash equilibrium specification is presented in Figure 5.b. In this case as well, there is indication of positive assortative mating, though the relationship is far more disperse. This is particularly true at higher values of δ_1 and δ_2.

The last two plots in Figure 5 exhibit the relationship between the preference and technology parameters of each spouse. These are not produced by "assortative" mating, *per se*, but the estimated distributions are, to be sure, functions of the assumed form of behavior within the marriage. For husbands (Figure 5.c) and wives (Figure 5.d), there is little evidence of a systematic relationship between productivity and preference parameters, in part due to the limited amount of variability .

Figure 6 contains the analogous scatter plots for the unconstrained Pareto weight case, with the Pareto weight set at 0.5. While some of the same general shape patterns are exhibited here as we saw under the assumption of Nash equilibrium, there are some subtle differences. For example, while the preference parameters of the spouses (Figure 6.a) continue to exhibit a strong positive dependence, there is more dispersion in the distribution of λ_2 conditional on λ_1 than we observed in Figure 5.a. There is also somewhat less of a systematic association between the spousal production function parameters (Figure 6.b) under the Pareto weight model.

In Figure 7 we present the scatter plots for the Constrained Pareto Optimal case. Adding the side constraint that efficient solution payoffs must exceed inefficient Nash equilibrium payoffs has a negligible impact on the relationships between the parameters that we have already observed for the previous two cases. The strong positive relationship between preference parameters of the spouses is similar to what was observed in the NE case (Figure 7.a), while the relationship between productivity parameters (Figure 7.b) is somewhat intermediate to the two cases already discussed. Once again, there is some indication of a negative relationship between the husband's productivity and preference parameters (Figure 7.c).

Our preferred specification, that of Endogenous Interaction, yields implied associations between parameters somewhat intermediate to the others we have examined to this point. The association between preference parameters is positive and approximately linear, as was true in the other cases. The support of the distribution is smaller than was the case previously. There is fairly clear positive dependence between the productivity parameters of the spouses (Figure 8.b). Once again, we see that there is a negative relationship between λ_1 and δ_1 in the population (Figure 8.c) and no indication of a relationship between λ_2 and δ_2 (Figure 8.d).

5.4 Welfare Implications of the Analysis

Our estimates of the distributions of the unobserved state variables, used in conjunction with the observed state variables in the data, allow us to examine the implied joint distribution of spousal welfare within our sample. We follow the methodology used in computing the scatter plots described above to compute the intrahousehold welfare levels. For household j in the sample, defined in terms of $(w_{1j}\ w_{2j}\ Y_j)$, we draw 1000 values of the unobserved state variable from the estimated distribution under behavioral regime k. Given the entire state variable vector, we compute time allocations under behavioral rule k, and then the utility level of each spouse. We then plot the utility levels (u_1, u_2) for each state variable vector under the four behavioral regimes. The results are shown in Figure 9.

In all four behavioral regimes, there is a very strong relationship between the attained utility levels of the spouses. This is not totally unexpected given the specifications of the utility and household production functions, which posit that all consumption in the household, aside from leisure, is public. Nonetheless, the specification in and of itself does not specify the preference weights associated with the public good, which, in principle, could have been small.

Within each figure we see that a strong majority of the points lie above the 45-degree line, indicating that wives have a somewhat higher payoff on average under our cardinal utility representation. If husbands and wives were perfectly symmetric, in the sense that $\lambda_1 = \lambda_2$, $\delta_1 = \delta_2$, and $w_1 = w_2$, then all utility outcomes in each figure should lie on the 45-degree line. In Figures 8.b through 8.d, where the outcomes involve the Pareto weight parameter α, even under perfect symmetry of preferences, productivity, and wages, values of α different than 0.5 will produce outcomes off of the 45-degree line. Since the notional Pareto weight is always set to 0.5, utility realizations are produced by asymmetry in spousal characteristics, both observed and unobserved.

The plot of utility payoffs for the Endogenous Interaction case (Figure 9.d) is somewhat interesting, in that the utility values are somewhat more highly related than in the other "cooperative" cases (Figures 9.b and 9.c). We attribute this to the presence of a "tighter" participation constraint for both partners, which limits the ability of either spouse to extract an inordinate share of the surplus. The correlation between welfare outcomes is highest in the Nash equilibrium case and lowest in the unconstrained Pareto environment.

5.5 Comparative Statics Exercise

We conclude the empirical analysis with another look at the importance of the discount factor in determining household outcomes. As we have emphasized at several points, the NE and CPO specifications are special cases of EI. When β is approximately equal to 0, all outcomes are determined in noncooperative Nash equilibrium, and the solutions are identical to those associated with NE. When β is arbitrarily close to 1, all outcomes are associated with CPO, meaning all households behave in an efficient manner. Our estimate of β, taken together with our estimates of the distribution of primitive parameters, implied that about one-fourth of households behaved inefficiently.

In this exercise, we utilize the estimate of the distribution of the primitive parameters under the EI specification to examine the impact of changing β on household decisions and (current period) utility. To do this, we generated data under the model estimates for 99 values of β, beginning with 0.01 and ending with 0.99. For each simulation, we computed the mean values of household time allocations and current period utility levels for each of the spouses. The results are presented in Figure 10.

Figure 10.a plots the average hours in the market by spouse over the range of β. We notice that the curves are flat until β increases to approximately 0.42; this is due to the fact that all households are in the noncooperative regime and time allocations in this regime are independent of β. As β increases beyond 0.42, the proportion of households able to implement efficient allocations is steadily increasing; beyond a β value of 0.52, all (simulated) households are able to behave efficiently. As a result, we see the largest change in market labor supply over the range 0.42 to 0.52, where average hours of work in the market goes up by over one-third for husbands and wives. For $\beta > 0.52$, when all households are behaving efficiently, changes in β merely change the constraint set facing the household members. The average effect of these changes is more beneficial for wives than for husbands, as we can see from Figure 10.c, which plots average (period) utility. As β increases over the "cooperative" range, the labor supply of wives is falling while that of husbands is increasing. In terms of time spent in housework (Figure 10.b), the time spent in housework by wives is increasing over this range, until it flattens out when β reaches 0.8. The housework of males is constant when β reaches the value 0.55 (approximately).

These results indicate the importance of the discount factor in household labor supply decisions. Even though we found that all households were able to reach the Pareto frontier for reasonably low values of β, β continues to have an impact on time allocation decisions and welfare outcomes even when all households behave cooperatively through its effect on the participation constraints of the spouses.

6 Conclusion

In this paper we have examined a general household time allocation model in a variety of behavioral frameworks, including one, that of Endogenous Interaction, that nests efficient and inefficient behavioral choices within it. We have worked within a very specific specification of preferences and household production technology when carrying out the empirical application, but we have considered very general forms of household heterogeneity. The point of this portion of the analysis, which is always probably worth making, is that strong functional form assumptions and restrictions on the distributions of state variables in the population are always required to identify the parameters characterizing any particular behavioral model. Testing between behavioral models cannot be done without resort to a number of nontestable identifying restrictions. We showed that a model that assumed Cobb-Douglas preferences with a CRS Cobb-Douglas household good production function which, along with spousal wages and household nonlabor income, were distributed according to an unspecified distribution F in the population could be consistently estimated with the data available in the PSID for any of three behavioral specifications, NE, PO, or CPO. Most importantly, there was no case that could be made for preferring one behavioral specification to another, since all were simply different representations of a saturated statistical model, in which there was a one-to-one mapping between data points and parameters.

The main theoretical contribution of the paper was the development of the model of Endogenous Interaction, which had households endogenously sorting into inefficient and efficient time allocation regimes. This model cannot be used to formulate a nonparametric estimator of F, and so can only be estimated using classes of estimators that rely on some parametric assumptions. We worked with a reasonably flexible parametric model of the distribution of the unobserved state variables, always conditioning on the set of state variables that were observable for each household (i.e., wages and nonlabor income). We found that the homogeneous discount factor was low ($\hat{\beta} = 0.524$), which in combination with our estimates of the distribution of the state variables in the population, implied that one-fourth of households behaved inefficiently.

We believe that there is a methodological contribution in our approach to creating "endogenous" mixing distributions. An obvious way to proceed when considering behavioral heterogeneity in the population is simply to allow some proportion of households to behave according to rule A and the rest according to rule B, where this probability is considered to be a primitive of the model. In our case, the mixing probability is not a primitive parameter, but is determined by the distribution of the state variables in the population, the *ex ante* value of α, and the value of the discount factor. Changes in the distribution of state variables, α, or β, in general.

When constraints are imposed on the Pareto weight formulation of the household time allocation problem, a constant population value of the "notional" Pareto weight must be adjusted to satisfy the time constraints. This produces what we might term model-induced "structural" heterogeneity in the *ex post* Pareto weights associated with the efficient outcomes in the population. We find that the Pareto models with side constraints produce significant amounts of heterogeneity in the *ex post* Pareto weight distributions. Under the Constrained Pareto specification, virtually all households implemented efficient outcomes at Pareto weights not equal to the *ex ante* value; in the EI specification, only 7 percent of cooperative households implemented a solution with a Pareto weight of 0.5, the *ex ante* value. On the basis of these results, we conclude that it is extremely import to consider the *ex post* heterogeneity induced by behavioral constraints which produce substantial heterogeneity in implied Pareto weights even when the notional Pareto weight is constant in the population. This finding is consistent with that of Mazzocco's (2007) analysis, which supports the dynamic adjustment of Pareto weights to satisfy evolving participation constraints.

In this paper we have only considered the impact of adding two particular side constraints to the efficient allocation problem. In terms of the Constrained Pareto specification, we added the constraint that each spouse receive at least what they would in Nash equilibrium. A number of bargaining-based models of household behavior assume that the outside option for each spouse is the value of being single. On a conceptual level, adding further constraints to the efficient allocation problem is straightforward, and, as we have seen, adding such constraints allows for a better correspondence between household time allocations observed in the data and those generated by the model. Extending such frameworks to a realistic dynamic setting which allowed for the possibility of inefficient household allocations would also considerably increase the appeal of the Pareto-weight approach to the analysis of household behavior

References

[1] Aaberge, R., U. Colombino, and S. Strom (1999). "An Empirical Analysis of Joint Household Decisions, with Taxes and Quantity Constraints." *Journal of Applied Econometrics* 14: 403-422.

[2] Becker, G. (1981). *A Treatise on the Family.* Cambridge: Harvard University Press.

[3] Blundell R. W. P. A Chiappori, T. Magnac, and C. Meghir (2007) "Collective Labour Supply: Heterogeneity and Non-Participation." *Review of Economic Studies* 74: 417-445.

[4] Bourguignon, F. (1984) "Individual or Strategic Rationality: The Case of Family Labor Supply." (in French) *Revue Economique* 35: 147-162.

[5] Bourguignon, F. and P.A. Chiappori (1992) "Collective Models of Household Behavior: An Introduction." *European Economic Review* 36: 355-364.

[6] Brown, M. and C. Flinn. (2007) "Investment in Child Quality over Marital States." Mimeo, UW-Madison and NYU.

[7] Brien, M., L. Lillard, and S. Stern (2006) "Cohabitation, Marriage, and Divorce in a Model of Match Quality." *International Economic Review* 47: 451-494.

[8] Chen, Z. and F. Woolley (2001) "A Cournot-Nash Model of Family Decision-Making." *Economic Journal* 111: 722-748.

[9] Chiappori P.A. (1988). "Rational Household Labor Supply" *Econometrica* 56: 63-89.

[10] Del Boca, D. and C. Flinn (1995). "Rationalizing Child Support Orders." *American Economic Review* 85: 1241-1262.

[11] Flinn, C. (2000) "Modes of Interaction Between Divorced Parents." *International Economic Review* 41: 545-578.

[12] Fortin, B. and G. Lacroix (1997) "A Test of Neoclassical Model and Collective Models of Household Labour Supply." *Economic Journal* 107: 933-955.

[13] Frederick,S.,G.Loewenstein,and T. O'Donoghue.(2002) "Time Discounting and Time Preference:A Critical Review" 41:545-578.

[14] Green, E. and R. Porter. (1984) "Noncooperative Collusion under Imperfect Price Information." *Econometrica* 52: 87-100.

[15] Hausman,J. (1979) "Individual Discount Rates and the Purchase and Utilization of Energy-Using Durables"Bell Journal of Economics 10:33-54.

[16] Kapteyn, A. and P. Kooreman (1990). "On the Empirical Implementation of Some Game-Theoretic Models of Household Labor Supply." *Journal of Human Resources* 25: 584-598.

[17] Lee, L.-F., and R. Porter (1984) "Switching-Regression Models with Imperfect Sample Separation Information - With an Application to Cartel Stability." *Econometrica* 52: 391-418.

[18] Leuthold, J. H. (1968) "An Empirical Study of Female Income Transfers and the Work Decision of the Poor." *Journal of Human Resources* 3: 312-323.

[19] Ligon, E. (2002) "Dynamic Bargaining in Households (with an Application to Bangladesh)". CUDARE Working Paper 972, UC-Berkeley Department of Agricultural and Resource Economics.

[20] Lugo-Gil, J. (2003). "Intrahousehold Bargaining and Welfare in Intact Families." Mimeo, NYU.

[21] Lundberg, S. and R. Pollak (1993). "Separate Spheres Bargaining and the Marriage Market." *Journal of Political Economy* 101: 988-1010.

[22] Mabli, J. (2007) "Life Cycle Inequality and Labor Supply: A Household Search Approach over Wages and Hours." Mimeo, NYU.

[23] Manser, M. and M. Brown (1980). "Marriage and Household Decision Making: A Bargaining Analysis." *International Economic Review* 21: 31-44.

[24] Mazzocco, M. (2007) "Household Intertemporal Behaviour: A Collective Characterization and a Test of Commitment." *Review of Economic Studies* 74: 857-895.

[25] McElroy M. and M. Horney (1981). "Nash-Bargained Household Decisions: Towards a Generalization of the Theory of Demand." *International Economic Review* 22: 333-349.

[26] McElroy, M. and M. Horney (1990). "Nash-Bargained Household Decisions: Reply." *International Economic Review* 31: 237-242.

[27] Pinheiro, J. and D. Bates (1996). "Unconstrained Parameterizations for Variance-Covariance Matrices." *Statistics and Computing* 6: 289-296.

[28] Porter, R. (1983) "A Study of Cartel Stability: The Joint Executive Committee, 1880-1886." *Bell Journal of Economics* 14: 301-314.

[29] Thaler,R.(1981) "Some Empirical Evidence on Dynamic Inconsistency"*Economics Letters* 8:201-207.

Table 1
PSID 2005 Sample
Means and (Standard Deviations)

$N = 823$

Variable	*Husband*	*Wife*
h	45.706	38.588
	(8.546)	(10.512)
τ	7.787	14.920
	(6.418)	(9.428)
w	22.009	15.823
	(13.626)	(9.327)
Y	118.151	
	(182.526)	

Table 2
Correlation Matrix of Observables

	h_1	τ_1	h_2	τ_2	w_1	w_2	Y
h_1	1.000	−0.017	0.093	0.060	0.029	−0.004	0.084
τ_1		1.000	0.081	0.321	−0.031	−0.026	0.024
h_2			1.000	−0.189	−0.132	0.084	0.011
τ_2				1.000	−0.018	−0.137	0.066
w_1					1.000	0.294	0.115
w_2						1.000	0.026

Table 3
Moments of Primitive Parameters
Fixed Effects Distributions

Parameter	*Behavioral Specification*		
	NE	*PO*	*CPO*
$\mu(\lambda_1)$	0.369	0.580	0.531
$\sigma(\lambda_1)$	0.095	0.166	0.106
$\mu(\lambda_2)$	0.302	0.430	0.456
$\sigma(\lambda_2)$	0.102	0.158	0.120
$\mu(\delta_1)$	0.075	0.075	0.075
$\sigma(\delta_1)$	0.057	0.057	0.057
$\mu(\delta_2)$	0.106	0.106	0.106
$\sigma(\delta_2)$	0.066	0.066	0.066
α		0.500	0.500

Table 4
Estimates of Primitive Parameter Moments
Flexible Parametric Specification
with Bootstrap Standard Errors

	NE		*PO*		*CPO*		*EHI*	
	Estimate	S.E.	Estimate	S. E.	Estimate	S. E.	Estimate	S.E.
$\mu(\lambda_1)$	0.349	0.016	0.469	0.012	0.483	0.005	0.449	0.006
$\sigma(\lambda_1)$	0.142	0.010	0.152	0.033	0.140	0.006	0.102	0.017
$\mu(\lambda_2)$	0.342	0.022	0.419	0.023	0.450	0.005	0.419	0.006
$\sigma(\lambda_2)$	0.245	0.022	0.233	0.022	0.227	0.009	0.150	0.024
$\mu(\delta_1)$	0.053	0.015	0.064	0.004	0.065	0.004	0.060	0.003
$\sigma(\delta_1)$	0.038	0.010	0.033	0.011	0.034	0.006	0.037	0.008
$\mu(\delta_2)$	0.121	0.003	0.113	0.004	0.109	0.002	0.116	0.005
$\sigma(\delta_2)$	0.039	0.003	0.029	0.008	0.026	0.016	0.047	0.015
$\mu(\alpha)$	-	-	-	-	0.515		0.527	
$\sigma(\alpha)$	-	-	-	-	0.042		0.039	
$P(\alpha = 0.5 \vert C)$	-		-		0.415		0.074	
$P(C)$	0		1		1		0.753	
β	-	-	-	-	-	-	0.524	0.004
Distance Measure	4682.370		4926.074		4389.950		3768.480	

Table A.1

Moments Used in the MSM Estimator

Sample Characteristic	Sample Value
Average h_1	45.706
Average h_2	38.588
Average τ_1	7.787
Average τ_2	14.920
St. Dev. h_1	8.546
St. Dev. h_2	10.512
Average $(h_1 \times h_2)$	1772.074
Average $(h_1 \times Y)$	5530.755
Average $(h_2 \times Y)$	4580.739
Average $(h_1 \times w_1)$	1009.293
Average $(h_2 \times w_2)$	618.773
Average $(\tau_1 \times Y)$	947.620
Average $(\tau_2 \times Y)$	1876.535
St. Dev. τ_1	6.418
St. Dev. τ_2	9.423
Average $(h_1 \times w_2)$	722.880
Average $(h_2 \times w_1)$	830.367
Average $(h_1 \geq 40)$	0.955
Average $(h_2 \geq 40)$	0.694
Average $(25 \leq h_1 < 40)$	0.0346
Average $(25 \leq h_2 < 40)$	0.210

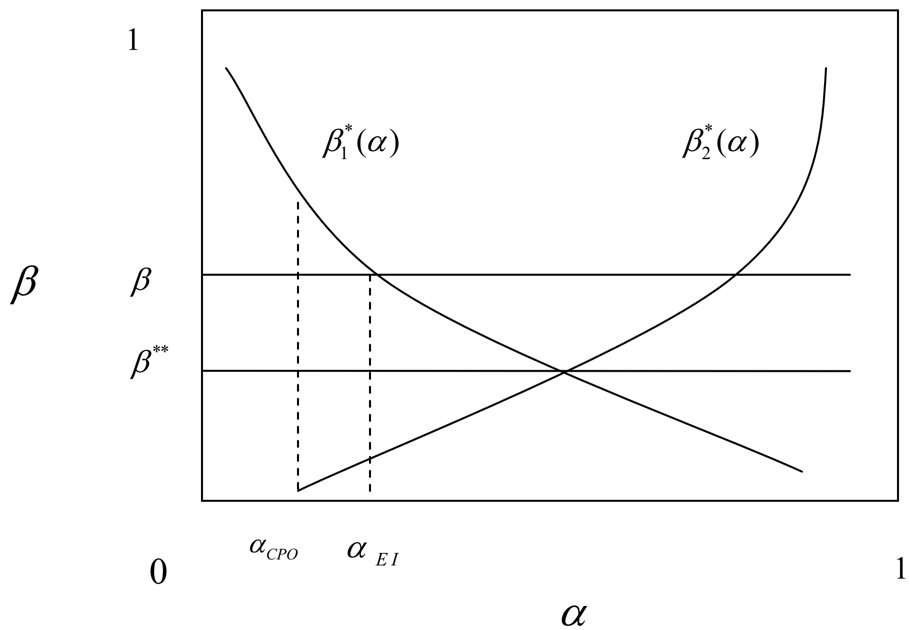

Figure 1
Critical β Values

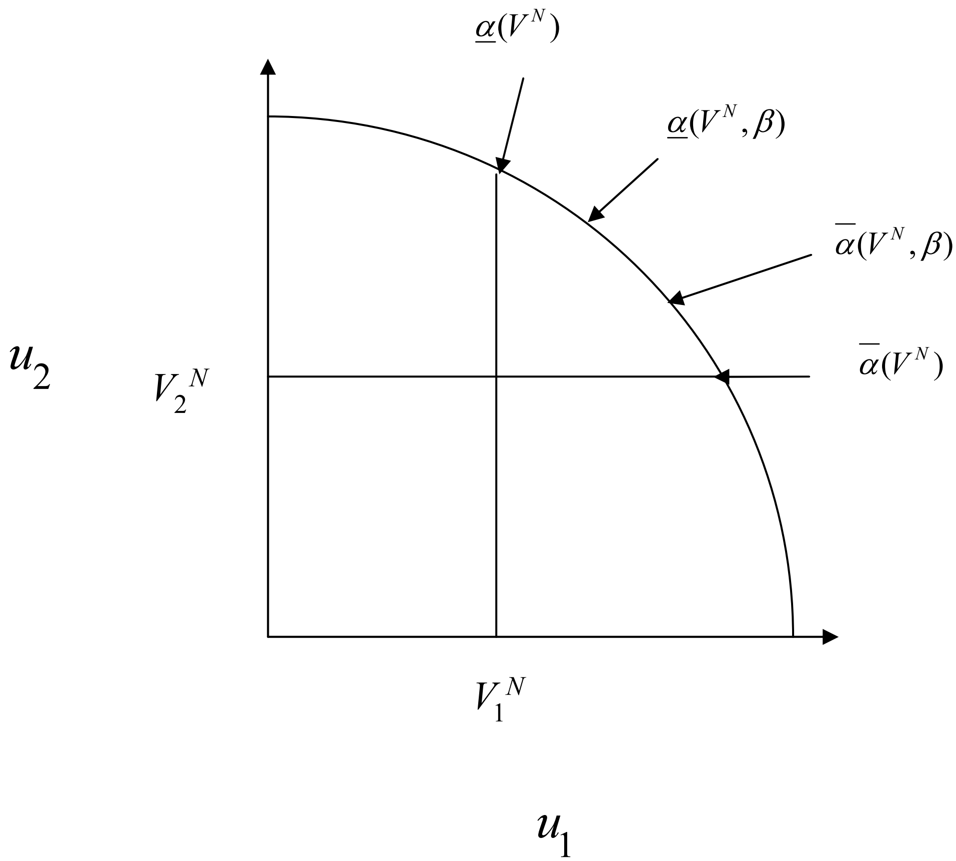

Figure 2
Pareto Frontier and Admissible Solutions

Figure 3.a
Distribution of α
Constrained Pareto Optimal Specification
(Excludes Cases with α = 0.5)

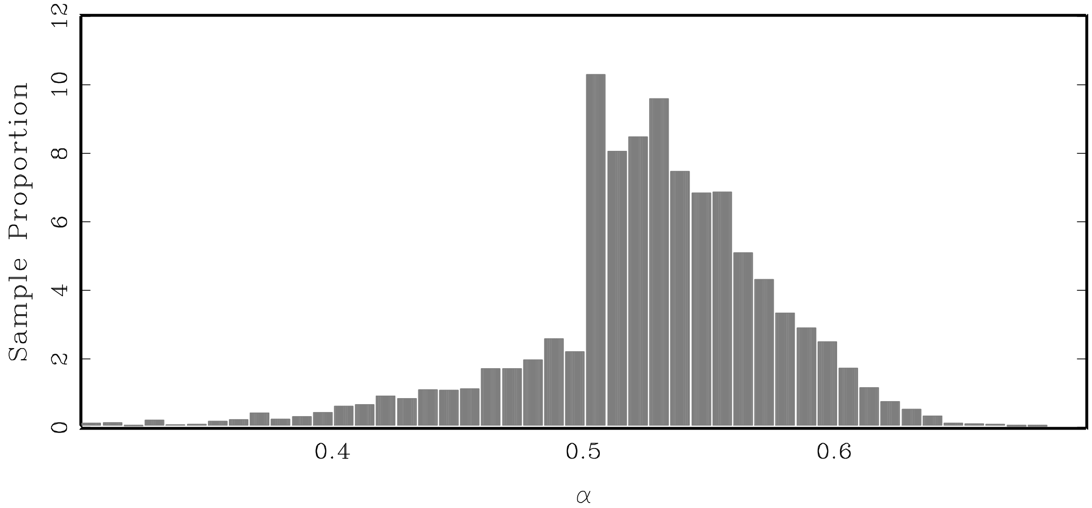

Figure 3.b
Distribution of α
Endogenous Interaction Specification
(Excludes Efficient Cases with α = 0.5)

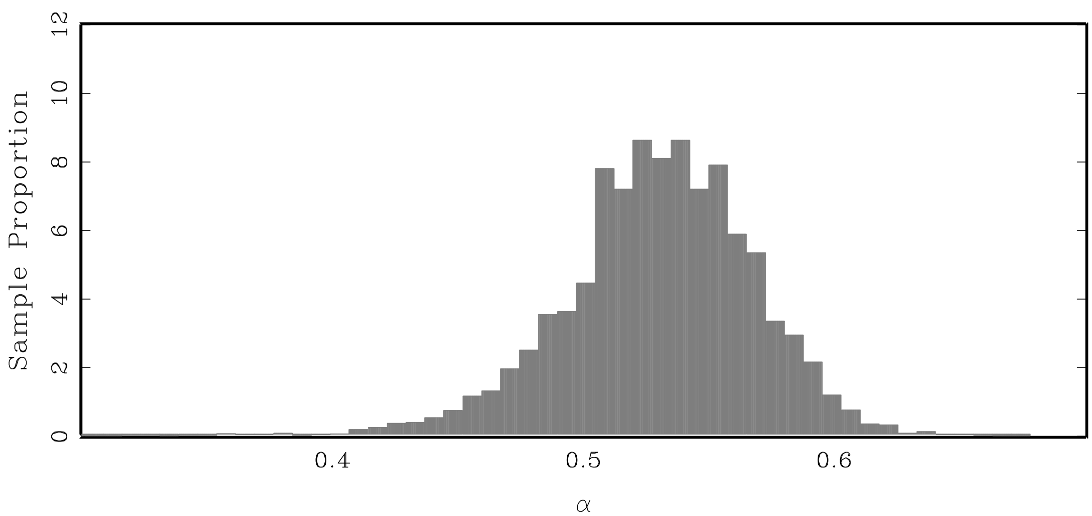

Figure 4
Distribution of Distance Measures by Model

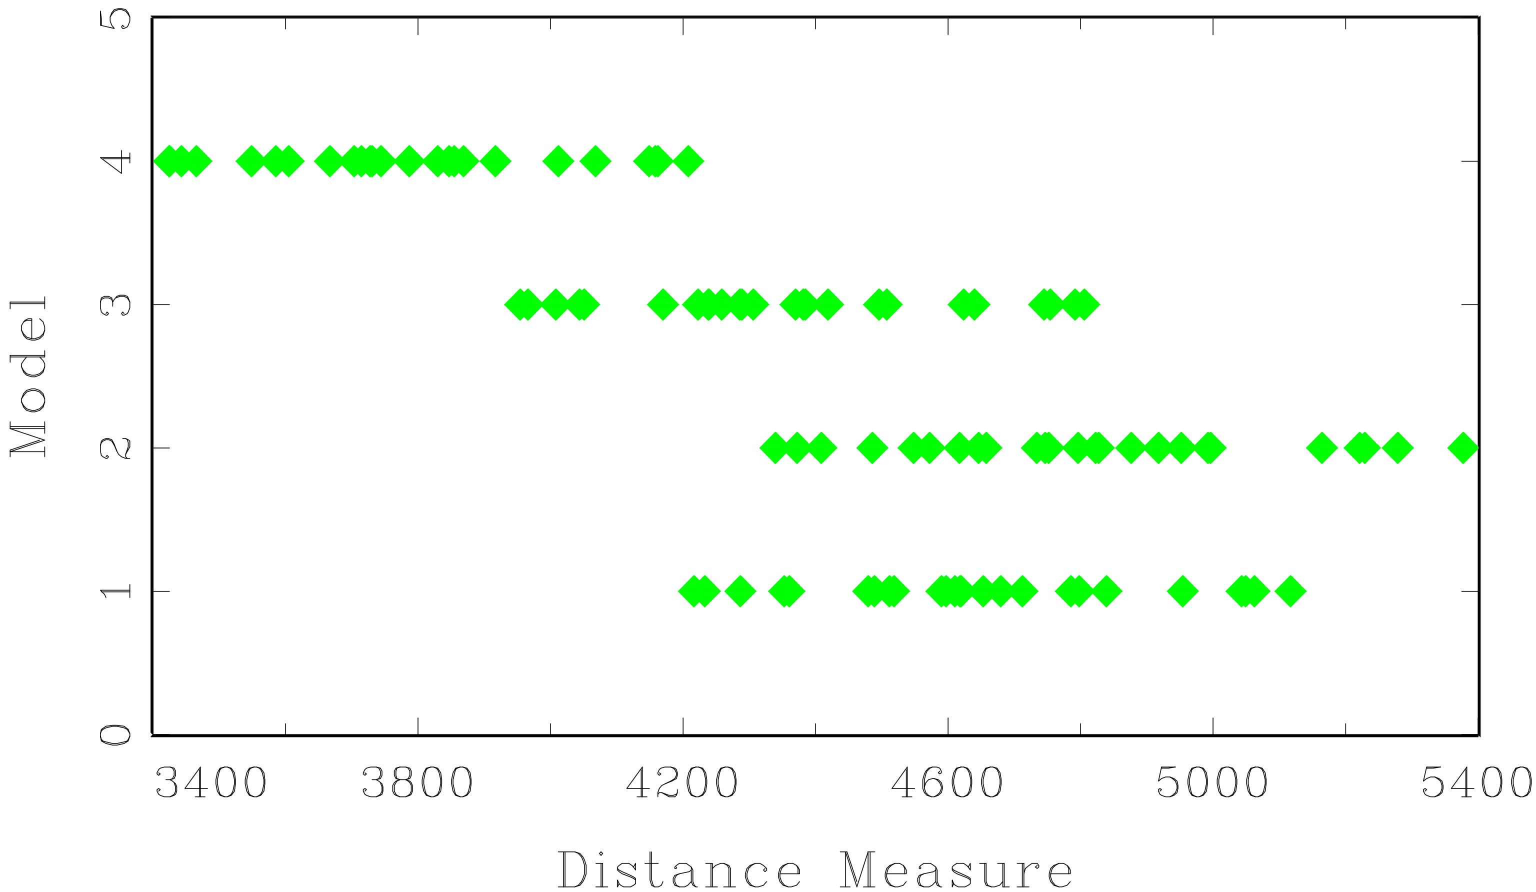

Figure 5.a
Spousal Preference Parameters
Nash Equilibrium

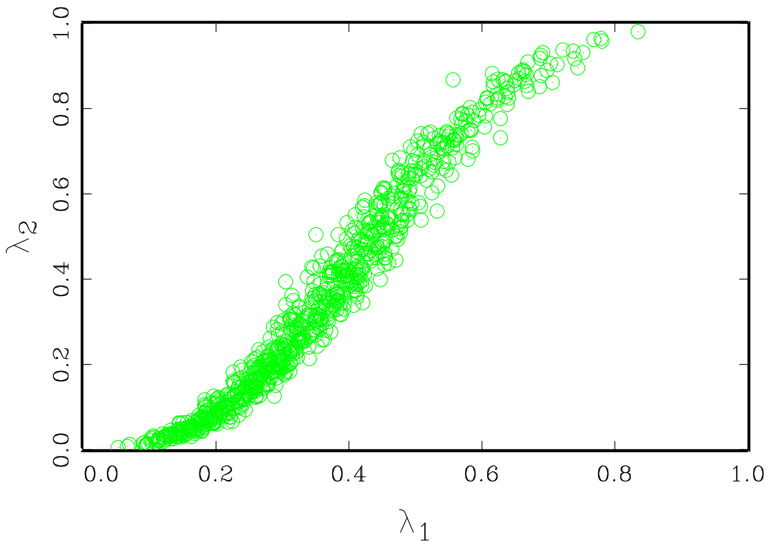

Figure 5.b
Spousal Productivity Parameters
Nash Equilibrium

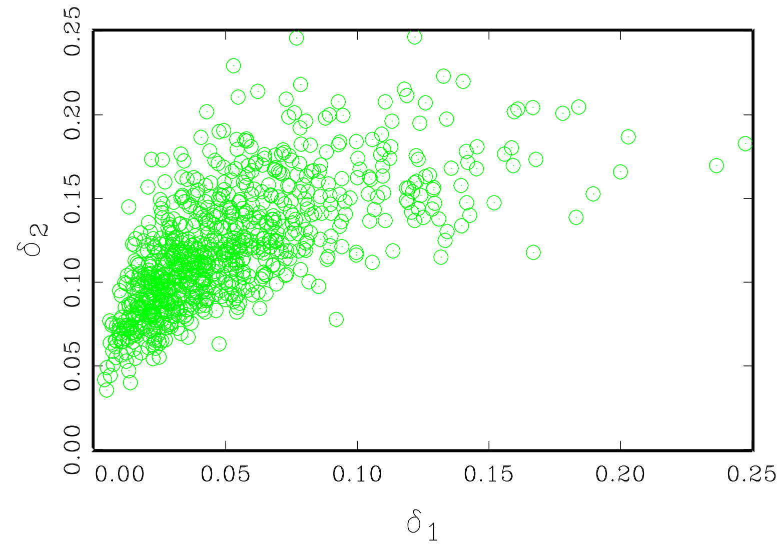

Figure 5.c
Husbands' Primitive Parameters
Nash Equilibrium

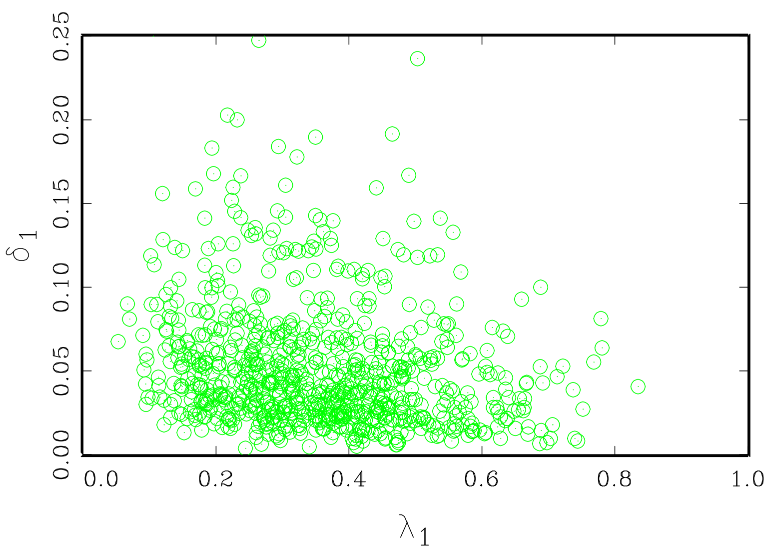

Figure 5.d
Wives' Primitive Parameters
Nash Equilibrium

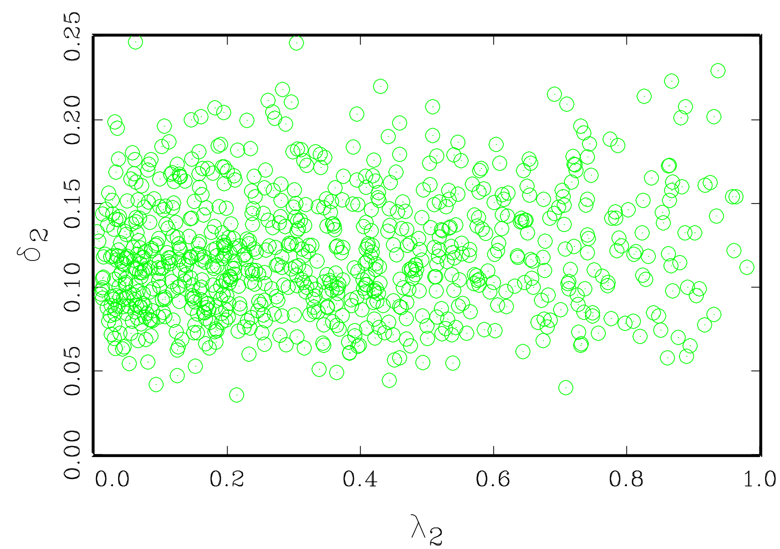

Figure 6.a
Spousal Preference Parameters
Pareto Weight

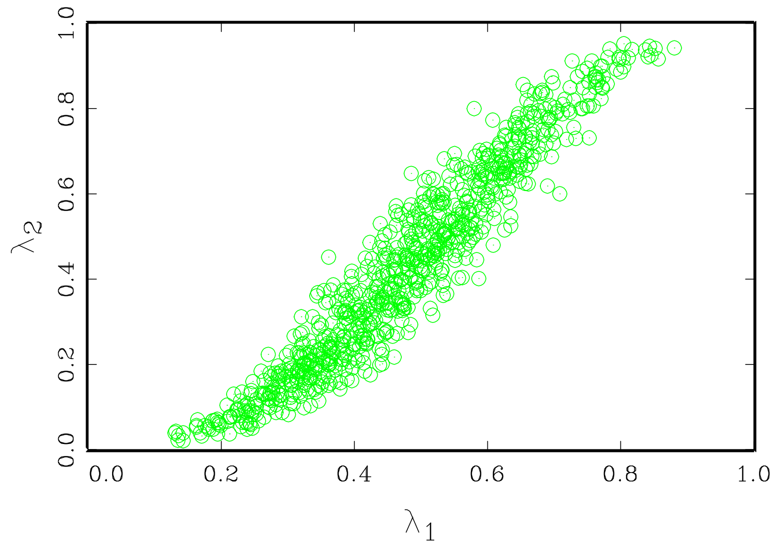

Figure 6.b
Spousal Productivity Parameters
Pareto Weight

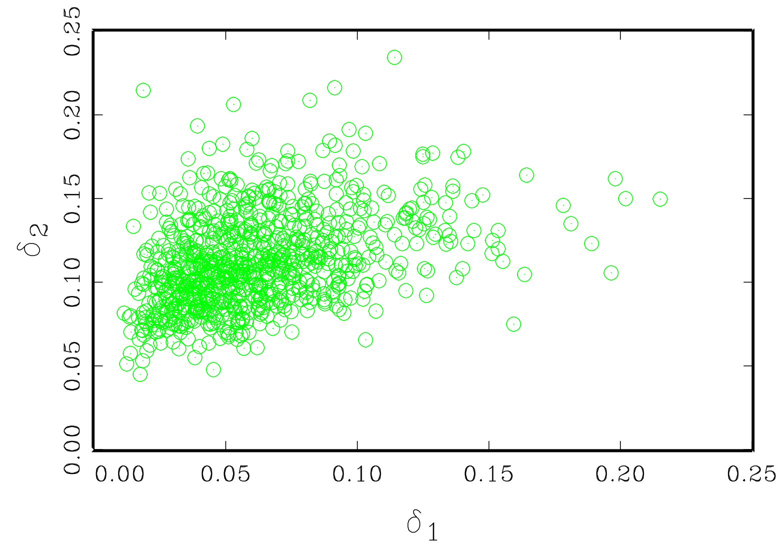

Figure 6.c
Husbands' Primitive Parameters
Pareto Weight

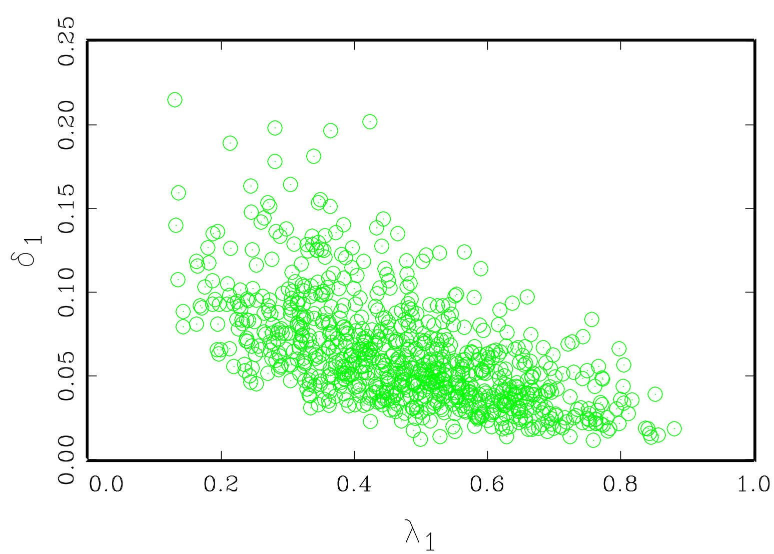

Figure 6.d
Wives' Primitive Parameters
Pareto Weight

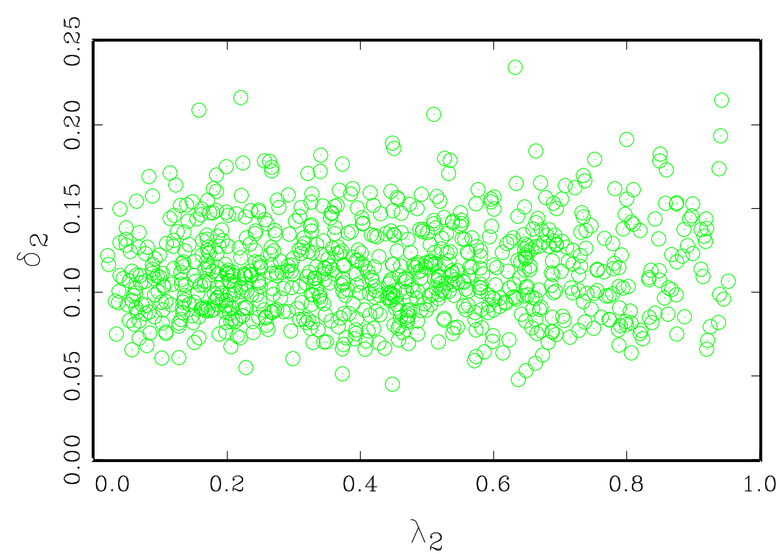

Figure 7.a
Spousal Preference Parameters
Constrained Pareto Optimal

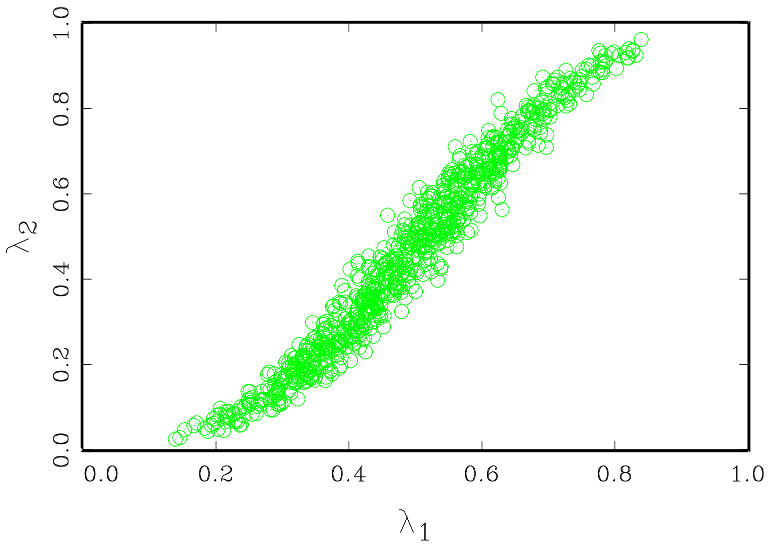

Figure 7.b
Spousal Productivity Parameters
Constrained Pareto Optimal

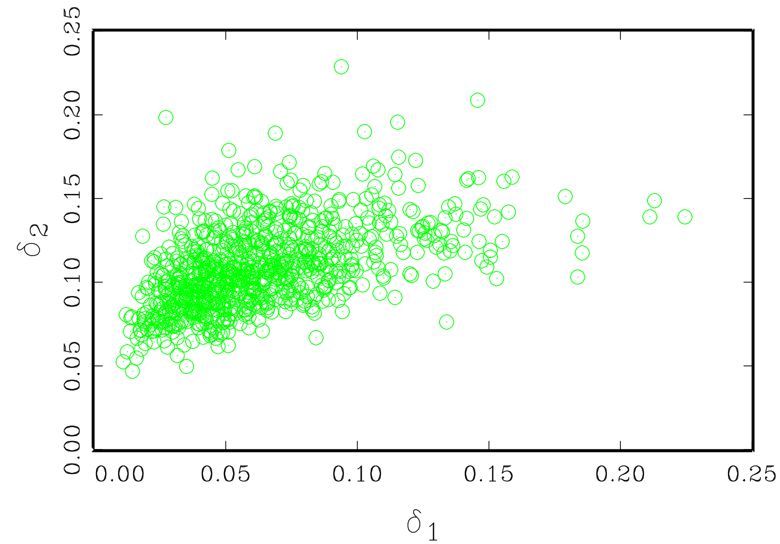

Figure 7.c
Husbands' Primitive Parameters
Constrained Pareto Optimal

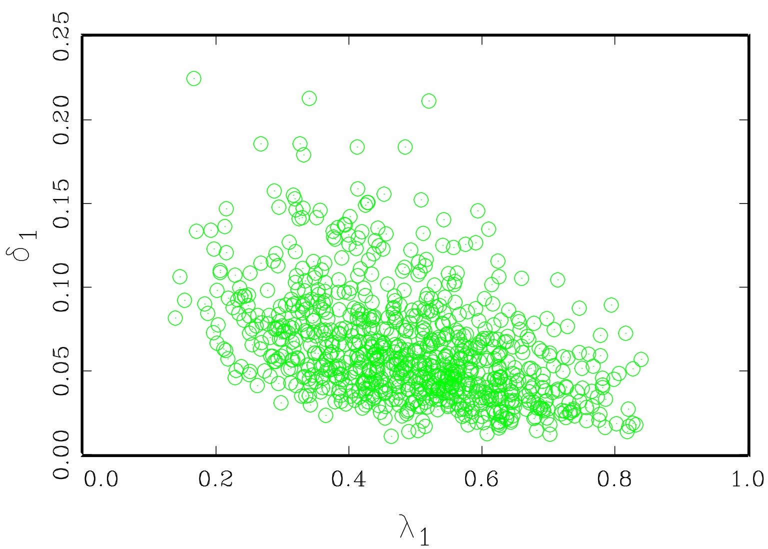

Figure 7.d
Wives' Primitive Parameters
Constrained Pareto Optimal

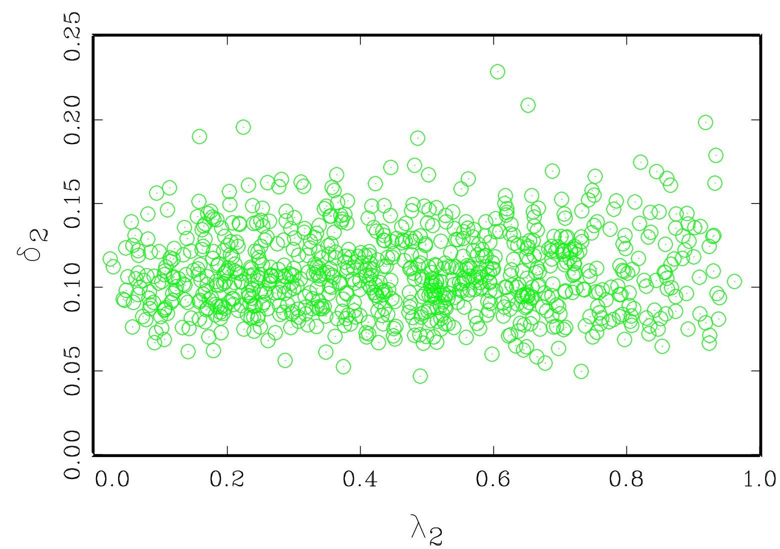

Figure 8.a
Spousal Preference Parameters
Endogenous Interaction

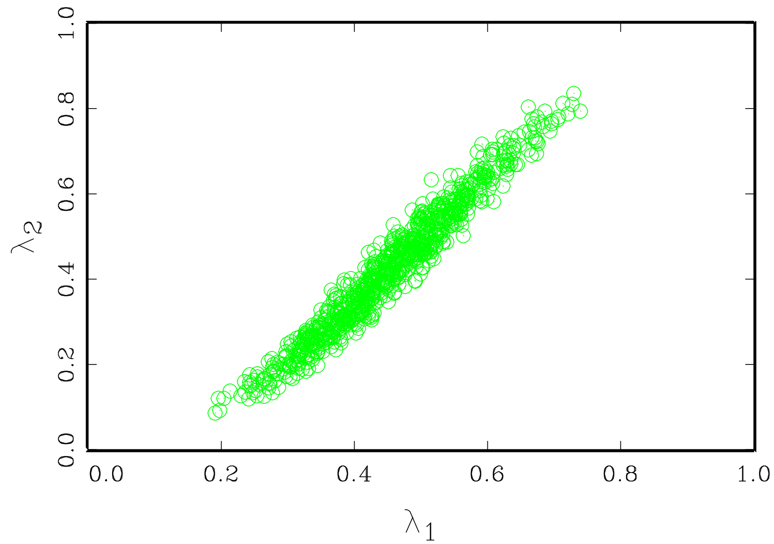

Figure 8.b
Spousal Productivity Parameters
Endogenous Interaction

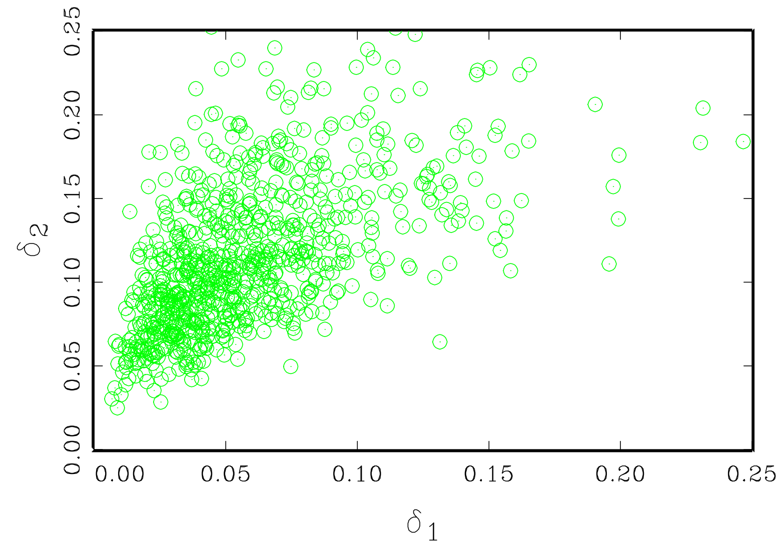

Figure 8.c
Husbands' Primitive Parameters
Endogenous Interaction

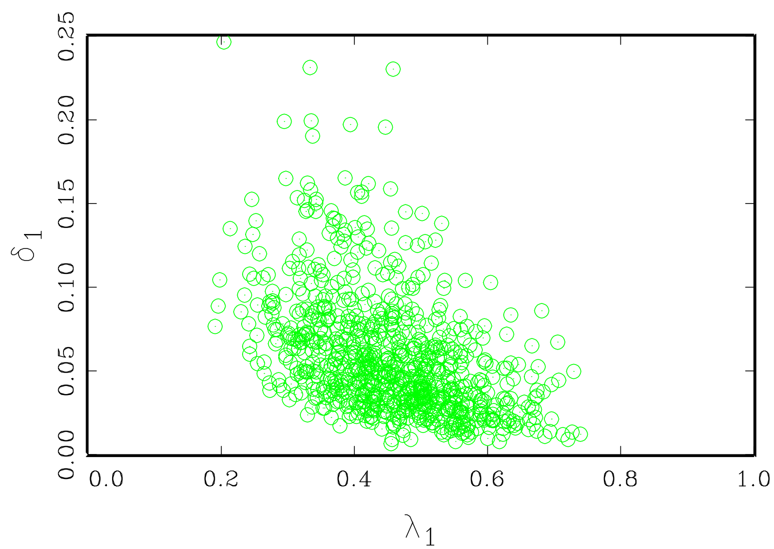

Figure 8.d
Wives' Primitive Parameters
Endogenous Interaction

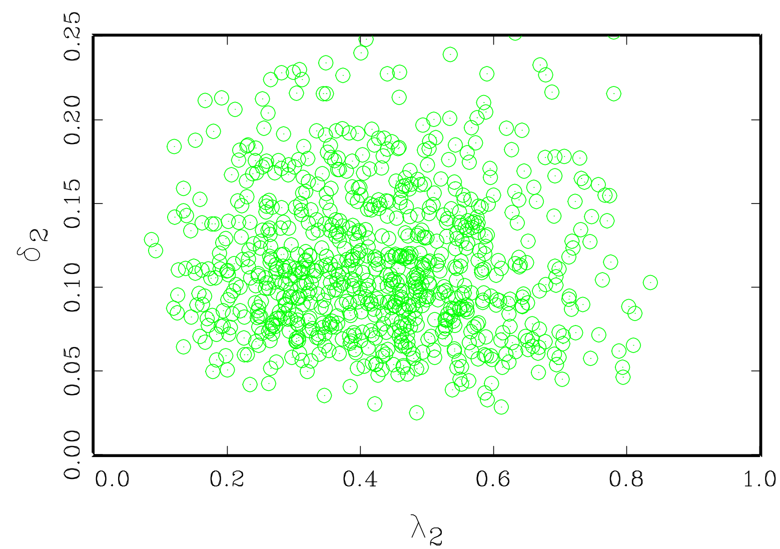

Figure 9.a
Welfare Outcomes
Nash Equilibrium

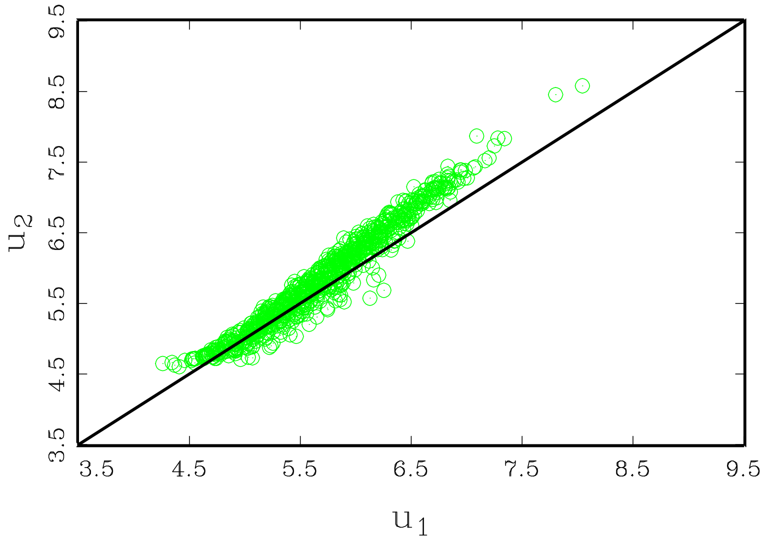

Figure 9.b
Welfare Outcomes
Pareto Optimal

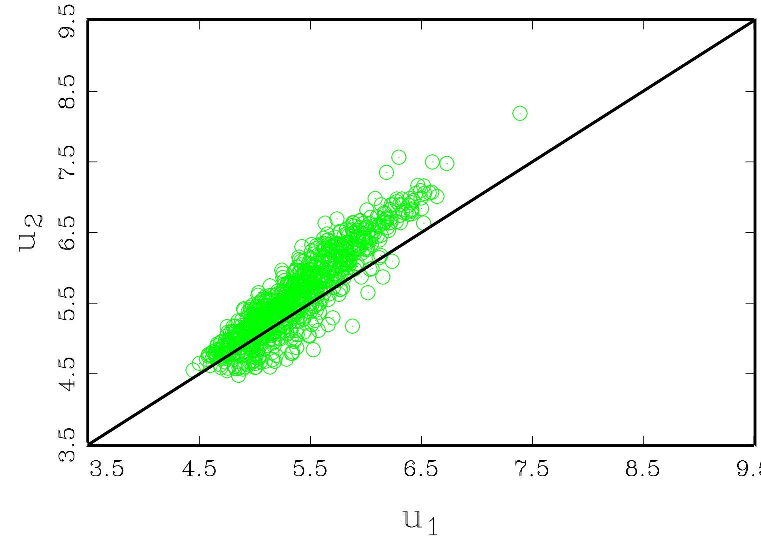

Figure 9.c
Welfare Outcomes
Constrained Pareto Optimal

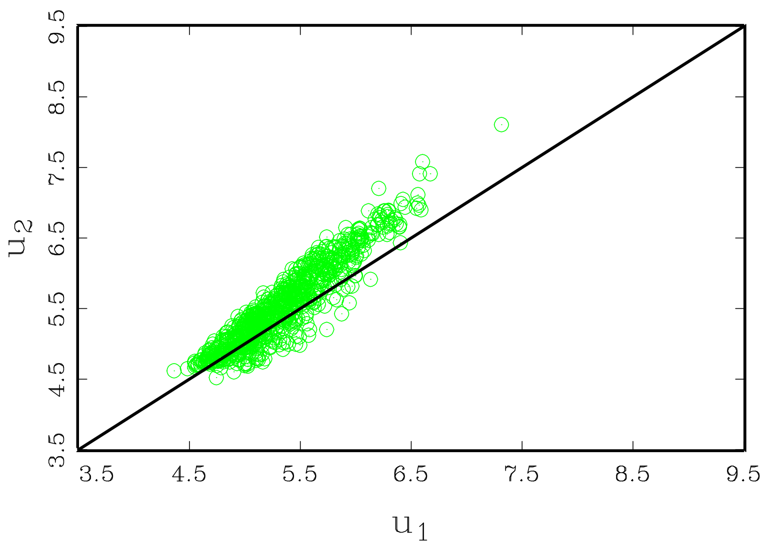

Figure 9.d
Welfare Outcomes
Endogenous Interaction

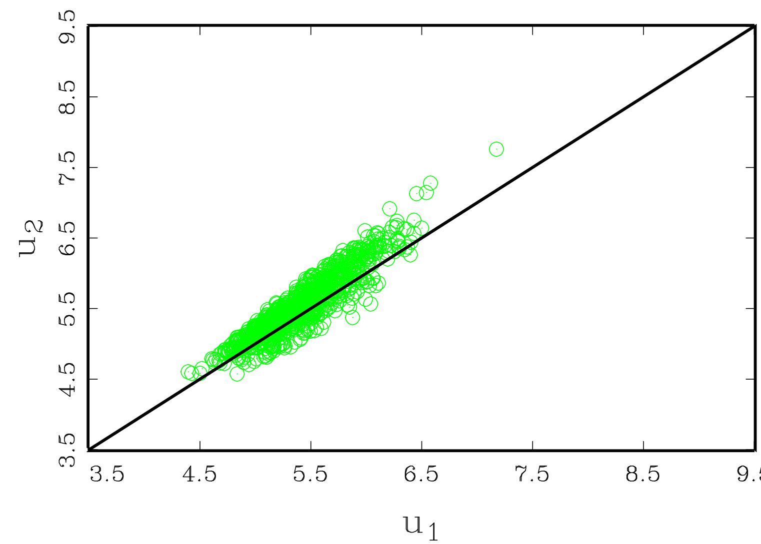

Figure 10.a
Mean Spousal Labor Supplies

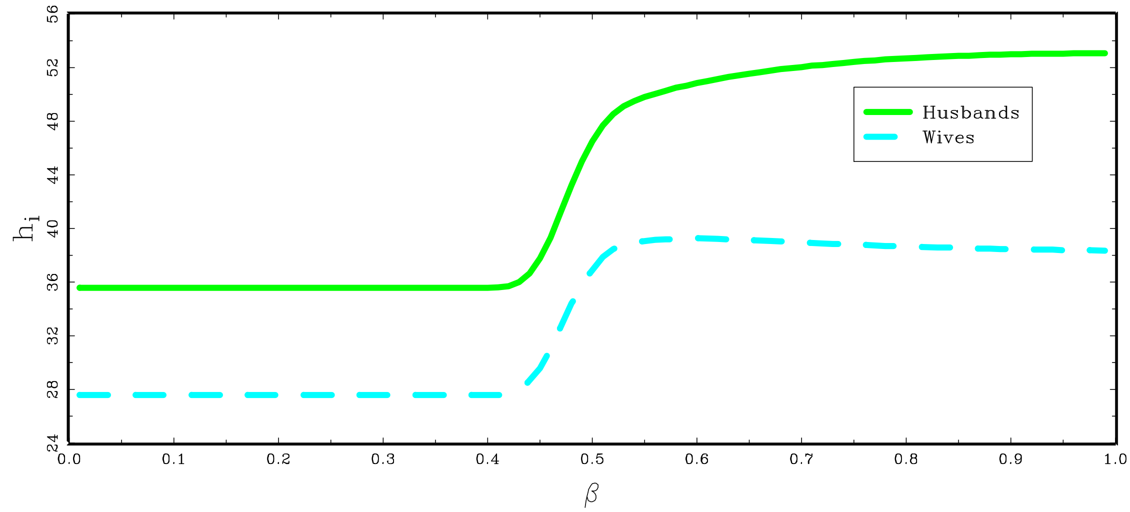

Figure 10.b
Mean Spousal Home Production

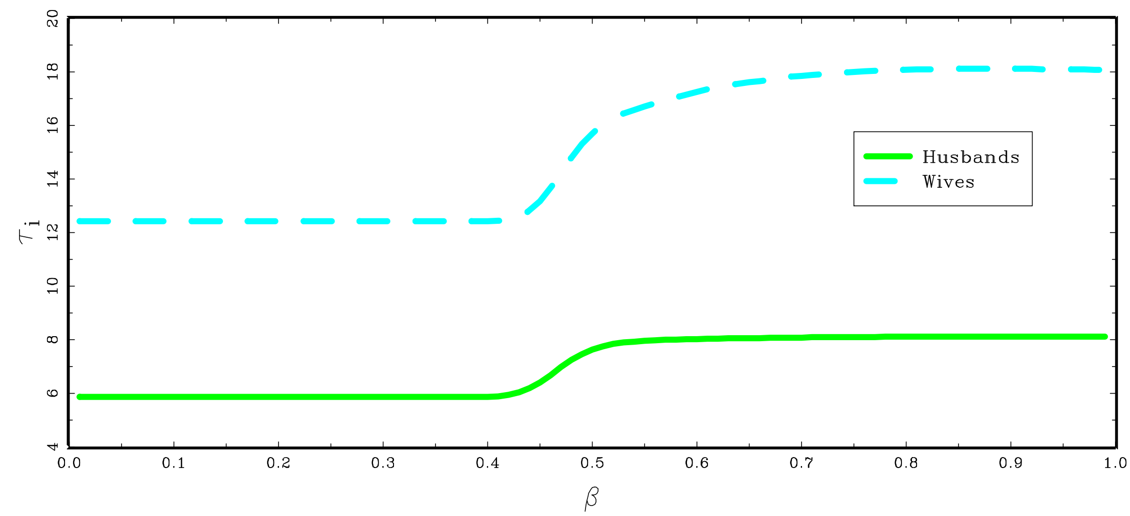

Figure 10.c
Mean Spousal Period Utility