

Household Time Allocation and
Models of Behavior: A Theory of Sorts

Daniela Del Boca
Christopher J. Flinn

No.8, May 2006

www.collegiocarloalberto.it

Household Time Allocation and Modes of Behavior: A Theory of Sorts^{1,2}

Daniela Del Boca
Department of Economics and Fondazione Collegio Carlo Alberto
Università di Torino

Christopher J. Flinn
Department of Economics
New York University

First Draft: August 2005

This Draft: April 2006

¹This research was partially supported by the C.V. Starr Center for Applied Economics at New York University. We thank the Fondazione Collegio Carlo Alberto for its hospitality while portions of this research were completed. We are grateful to Olivier Bargain, Xiaohong Chen, James Mabli, Tom MaCurdy, and Yoram Weiss for helpful comments, as well as to participants in the August 2005 SITE workshop on “The Nexus Between Household Economics and the Macroeconomy” and the conference “Interactions within the Family” held at the Università di Torino. We remain responsible for all errors and omissions.

²© 2006 by Daniela Del Boca and Christopher Flinn. Any opinions expressed here are those of the authors and not those of the Fondazione Collegio Carlo Alberto.

Abstract

We make the point that a flexible specification of spousal preferences and household production technology precludes the possibility of using revealed preference data on household time allocations to determine the manner in which spouses interact. Under strong, but standard, assumptions regarding marriage market equilibria, marital sorting patterns can be used essentially as “out of sample” information that allows us to assess whether household behavior is cooperative. We use a sample of households drawn from a recent wave of the Panel Study of Income Dynamics, and find some evidence supporting the view that households behave in a cooperative manner.

JEL: D13, J12, J22

Keywords: Bilateral Matching, Household Time Allocation, Nash Bargaining

1 Introduction

Most analyses of household behavior conducted at the microeconomic level posit cooperative behavior by spouses (for an exception, see Chen and Woolley (2001)). In fact, Chiappori and his coauthors (e.g., Chiappori (1992), Browning and Chiappori (1998)) have argued that all such models should posit efficiency as an identifying assumption when attempting to estimate individualistic preferences using data on household allocations. Such an assumption, however, leads to other difficult identification issues since the dependent variables, which are household allocations, are not uniquely determined.

In order to “close” the cooperative model, analysts have resorted to one of two devices. Since some of the original bargaining approaches to household behavior relied on the Nash bargaining axiomatic solution (e.g., Manser and Brown (1980) and McElroy and Horney (1981)), the use of some sort of refinement to select one of the continuum of possible outcomes associated with points on the Pareto frontier was a jumping off point (McElroy (1990)). While the use of a refinement approach solves the multiple equilibria problem, it does so at the cost of the necessity of specifying outside options and bargaining power weights (in the case of nonsymmetric Nash bargaining).

Alternatively, Chiappori and his collaborators (e.g., Chiappori (1988,1992), Browning et al. (1994), Browning and Chiappori (1998)) have proposed a data-based strategy to estimate the household utility function $\mu U_1(x) + (1 - \mu)U_2(x)$, where μ is the Pareto weight attached to the individualistic utility of agent 1, and x is a vector of consumption choices. The solution to this problem is guaranteed to lie on the Pareto frontier for $\mu \in [0, 1]$. Model identification is achieved through restrictions regarding the arguments of the weighting and individualistic utility functions as well as functional forms. Identification is achieved without resort to a specific axiomatic solution, with the data given the power to solve the multiple equilibria problem within the particular model structure.

While each of these competing approaches to the estimation of cooperative equilibria have their own advantages, both clearly have some unappealing aspects as well. From an econometric perspective, noncooperative equilibria are attractive since it is often straightforward to demonstrate existence and uniqueness given common specifications of spousal objectives, household production technologies, and constraint sets. Though cooperative equilibria lead each spouse to a superior welfare outcome in the absence of transactions costs associated with attaining the Pareto frontier, some prior empirical evidence suggests that the welfare gain to cooperative behavior may be small. For example, under strong functional form assumptions on individualistic utilities, Del Boca and Flinn (2006) found generally small differences between welfare levels associated with cooperative and noncooperative behavior in a sample of Italian married couples within a framework that allowed for the choice of mode of behavior. In analyzing the behavior of divorced parents, Del Boca and Flinn (1994) whose welfare was interdependent due to the presence of their child, the authors found little difference in the welfare of the parents under cooperative and noncooperative behavior. These results are quite idiosyncratic of course, and do not imply that

expenditure patterns cannot be very different under alternative behavioral assumptions. However, they do raise the question of whether, in the presence of implementation costs associated with cooperative outcomes, spouses should choose cooperative solutions to the allocation problem.

In this paper we explore the issue of the “mode” of household behavior, and for simplicity focus on only two alternatives, Nash equilibrium (NE) and symmetric Nash bargaining (NB). We first show that after allowing for general forms of population heterogeneity in preferences, household productive ability, market productivity, and time endowments, it is not possible to distinguish between NE and NB on the basis of household time allocation decisions. To do so requires imposing homogeneity restrictions that may not be justifiable and are rarely tested.

Nevertheless, we show that the patterns of *marital sorting* observed in the data do contain information on the manner in which household members interact. We are by no means the first to point this out. Following the view of Becker (1991) that marriage is a partnership for joint production and consumption, several authors have analyzed aspects of the marriage market to explore marital behavior and the gains to marriage (e.g., Choo and Siow (2006), Dagsvik et al. (2001), Pollack (1990)). Other research has explored the effects of the marriage market on household behavior. While Aiyagari, Greenwood and Guner (2000) and Greenwood, Guner and Knowles (2003) have focused on the link between the marriage market and parental investments in children and patterns of intergenerational mobility, Fernandez et al. (2005) have studied the implication of marital sorting for household income inequality.

Micro analyses such as Browning et al. (2003), Seitz (1999), and Igiyun and Walsh (2004) have explored aspects of household formation that precede marriage to merge household models with marital sorting in order to explore the implications of spousal matching for intrahousehold allocations. While the objective of these papers is mainly to identify sharing rules and to consider with household allocations are efficient, we use marital sorting to investigate *what* type of interaction is most consistent with observed outcomes.

The basic idea of our approach can be summarized in the following way. We begin by assuming that spouses interact using some rule R , and then use the observed household time allocations, along with exogenously determined wages and nonlabor incomes, to “back out” the parameters characterizing both spouses within each household in the sample. Using these individual-specific parameters, we can then construct preference orderings for each male over all possible females in this marriage “sub-market” assuming the household allocations are chosen according to R , and we can construct the preference orderings for the females in a similar manner. Armed with these R -specific preference orderings, we then apply the Gale and Shapley (1962) - henceforth GS - bilateral matching algorithm to determine the predicted equilibrium matches under R . We then compare the correspondence between the predicted matches and the observed ones for R using a variety of metrics. This analysis is conducted for the two modes of behavior (R) we consider, symmetric Nash bargaining and Nash equilibrium, and we conclude by comparing the relative performance

of the two rules under the various metrics examined.

The plan of the paper is as follows. Section 2 contains the description of the model and the bilateral matching algorithm. In Section 3 we explore econometric issues, which are reasonably straightforward for the most part. Empirical results are presented in Section 4, and Section 5 contains a brief conclusion.

2 Model

A focus of our attention will be household formation. Without loss of *empirical* generality (as we shall see below), we will assume the following simple determination of household utility in a static context. We assume a Stone-Geary utility function for spouse i of the form

$$u_i(l_i, K) = \alpha_i \ln(l_i - \lambda_i) + (1 - \alpha_i) \ln(K - \underline{K}_i), \quad i = 1, 2,$$

where l_i is the leisure of spouse i , λ_i is their leisure “subsistence level,” K is a public good that is produced within the household, \underline{K}_i is the subsistence level of the public good for spouse i , and α_i is the preference weight attached to “discretionary” leisure. For purposes of model identification, we will normalize the subsistence level $\underline{K}_i = 0$, $i = 1, 2$. The household good K is produced according to a Cobb-Douglas technology

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

where τ_i is the time input of spouse i in household production, δ_i is the elasticity of K with respect to time input τ_i , and M is total income of the household, or

$$M = w_1 h_1 + w_2 h_2 + y_1 + y_2,$$

where w_i is the wage rate of spouse i , h_i is their hours of work, and y_i is the nonlabor income of spouse i . We assume that each of the production elasticities δ_1 and δ_2 is strictly positive, so that there are increasing returns to household production.¹ The “physical” time endowment of each spouse is T , and

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

It will be convenient to think of there being a “notional” time endowment specific to each individual in the population. This notational time endowment is equal to $\tilde{T}_i \equiv T - \lambda_i$, where λ_i can be positive, negative, or zero.

Each individual has their own value of market productivity, with the value of their time in the market given by w_i . Moreover, each individual has a nonlabor income level of y_i . Both of these quantities are determined outside of the model.

¹We have chosen not to impose constant returns to scale in this function for purposes of conducting the matching analysis conducted below.

Within our framework, all households in the population share the same preference and household production structure. The population is, however, characterized by heterogeneity in all of the parameters that appear in the functions defined above. The population consists of two types of agents, males (husbands) and females (wives). Each subpopulation is characterized by a distribution of characteristics particular to that type. The cumulative distribution function of characteristics of individuals of gender i is

$$G_i(\alpha_i, \delta_i, \tilde{T}_i, w_i, y_i).$$

Then a household is defined by the vector of state variables

$$S = (\alpha_1, \delta_1, \tilde{T}_1, w_1, y_1) \cup (\alpha_2, \delta_2, \tilde{T}_2, w_2, y_2).$$

Given a value of S , the household determines equilibrium time allocations and the resultant welfare distribution in the household according to some rule R . Thus R is a mapping from S into a vector of observable household choices, in our case given by the vector

$$C = (h_1, h_2, \tau_1, \tau_2).$$

Thus

$$C = R(S). \tag{1}$$

We will discuss specific properties of the mapping R below, but for now we assume that R assigns a unique value E to any vector $S \in \Omega_S$, where we will think of Ω_S as the parameter space of household characteristics.

2.1 Noncooperative Behavior

We begin our investigation of the time allocation decision of the household with the case of Nash equilibrium. Later we will turn our attention to cooperative models of household behavior.

The reaction function for spouse 1 in a household characterized by S is given by

$$\begin{aligned} (h_1, \tau_1)^*(h_2, \tau_2; S) &= \arg \max_{h_1, \tau_1} \alpha_1 \ln(\tilde{T}_1 - h_1 - \tau_1) \\ &\quad + (1 - \alpha_1)[\delta_1 \ln \tau_1 + \delta_2 \ln \tau_2 + \ln(y + w_1 h_1 + w_2 h_2)]. \end{aligned}$$

Assuming an interior solution for h ,² the solutions are given by continuously differentiable functions

$$\begin{aligned} h_1^* &= h_1^*(h_2, \tau_2; S) \\ \tau_1^* &= \tau_1^*(h_2, \tau_2; S). \end{aligned}$$

²Whenever $\alpha_1 > 0$ and $\delta_1 > 0$, an interior solution for τ_1 is assured by the Inada condition.

An analogous pair of reaction functions exists for the second individual. Under our specification of preferences and the production technology, there exists a unique Nash equilibrium

$$\begin{aligned} h_1^{**} &= h_1^*(h_2^{**}, \tau_2^{**}; S) \\ \tau_1^{**} &= \tau_1^*(h_2^{**}, \tau_2^{**}; S) \\ h_2^{**} &= h_2^*(h_1^{**}, \tau_1^{**}; S) \\ \tau_2^{**} &= \tau_2^*(h_1^{**}, \tau_1^{**}; S). \end{aligned}$$

Insuring that h_1^{**} and h_2^{**} are both greater than zero requires restricting the parameter space Ω_S . We will provide further discussion on this point in the econometrics section below.

Associated with the Nash equilibrium is a welfare pair $(V_1^{NE}(S), V_2^{NE}(S))$. These values will be used as outside options in the Nash Bargaining part of the analysis. After considering the marital sorting process, we will justify the use of these values as threat points.³

2.2 Symmetric Nash Bargaining

We consider the case of symmetric Nash bargaining, once again, without any loss of (empirical) generality. Denote the outside options of the husband and wife by $Q_1(S, Z_1)$ and $Q_2(S, Z_2)$, where Z_i represents environmental characteristics for individual i that influence the value of the alternative to behaving cooperatively within marriage S . Then the Nash bargained household time allocation is

$$\begin{aligned} &(h_1^{NB}, \tau_1^{NB}, h_2^{NB}, \tau_2^{NB})(S, Z_1, Z_2) \\ &= \arg \max_{h_1, \tau_1, h_2, \tau_2} (U_1(h_1, \tau_1, h_2, \tau_2; S) - Q_1(S, Z_1)) \times (U_2(h_1, \tau_1, h_2, \tau_2; S) - Q_2(S, Z_2)), \end{aligned}$$

where $U_i(h_1, \tau_1, h_2, \tau_2; S) = \alpha_i \ln(\tilde{T}_i - h_i - \tau_i) + (1 - \alpha_i)[\delta_1 \ln \tau_1 + \delta_2 \ln \tau_2 + \ln(y_1 + y_2 + w_1 h_1 + w_2 h_2)]$, $i = 1, 2$. Given our soon to be justified assumption that $Q_i(S, Z_i) = V_i^{NE}(S)$, we will dispense with the variables (Z_1, Z_2) , and write

$$\begin{aligned} &(h_1^{NB}, \tau_1^{NB}, h_2^{NB}, \tau_2^{NB})(S) \\ &= \arg \max_{h_1, \tau_1, h_2, \tau_2} (U_1(h_1, \tau_1; S) - V_1^{NE}(S)) \times (U_2(h_2, \tau_2; S) - V_2^{NE}(S)) \quad (2) \end{aligned}$$

We note that since we restrict the parameter space Ω_S so as to produce noncooperative time allocations that are strictly positive, the choices made under Nash bargaining, with the noncooperative equilibrium values serving as outside options, will be strictly positive as well.

³We will consider the case in which there are an equal number of males and females in the population. In the marriage equilibrium we define all agents will have the possibility of being married to an individual of the opposite sex. We find that the value of marriage exceeds the value of living alone for all population members in equilibrium, so the correct outside option will be the value of noncooperative marriage.

2.3 Single Agent Welfare

Single agents must produce their own household goods and as a result receive no “subsidy” from a partner in terms of time contributions to production⁴ or money contributions through earnings and nonlabor income. Then the production technology the single individual i faces is

$$K = \tau_i^{\delta_i}(y_i + w_i h_i), \quad (3)$$

where we have used the convention $0^0 = 1$ in eliminating the missing spouse’s time contribution.⁵ Then the single agent has a utility yield of

$$\begin{aligned} V_i^0(S_i) &= \max_{h_i, \tau_i} \alpha_i \ln(\tilde{T}_i - h_i - \tau_i) \\ &\quad + (1 - \alpha_i)[\delta_i \ln \tau_i + \ln(y_i + w_i h_i)], \end{aligned}$$

where $S_i \equiv (\alpha_i, \delta_i, \tilde{T}_i, w_i, y_i)$.

2.4 Marital Sorting

The subpopulation distributions G_1 and G_2 are assumed to exogenously determined. The marriage model equilibrium which matches males and females produces an endogenous joint distribution of S , which we denote by $H(S)$, of which G_1 and G_2 are appropriately defined marginal distributions.

We consider the case of a closed population in which there exists a total of $2N$ individuals, equally divided between males and females. Male i is defined by his vector of characteristics

$$m_i = (\alpha_{1i}, \delta_{1i}, \tilde{T}_{1i}, w_{1i}, y_{1i}),$$

while female j is defined by her characteristics vector

$$f_j = (\alpha_{2j}, \delta_{2j}, \tilde{T}_{2j}, w_{2j}, y_{2j}).$$

Following GS, we consider the simple case in which there exists a marriage market in which individuals from the different subpopulations are matched one-to-one, all individual characteristics are perfectly observable, and the market clears instantaneously. Each male has preferences over possible mates, with the preference ordering of male m_i given by $P(m_i)$.

Similarly, the preference ordering of woman j is given by $P(f_j)$. In each case, the preference ordering amounts to a sequence of potential mates ranked in descending order,

⁴Our specification of household production and utility could lead to “negative” subsidies if the spouse provides less than 1 unit of time to household production. Income externalities could be zero but never negative.

⁵That is, the missing spouse has an associated δ equal to 0 and supplies 0 amounts of time to household production.

and may include ties. In addition, remaining single may dominate being married to certain individuals of the opposite sex. The value of this state we shall denote by f_0 to a male (that is, the “null” female) and m_0 if we are describing the preference ordering of a female. For example, with $N = 5$, we could have

$$P(m_4) = f_3, f_1, f_2, f_5, f_4.$$

That is, male 4’s first choice as a mate is female 3, followed by 1, 2, 5, and 4. The preferences of female 2 might be represented by

$$P(f_2) = m_4, m_1, m_3, m_0. \tag{4}$$

In this case, she prefers male 4 to male 1 to male 3, and would rather live alone than be married to either male 2 or male 5. As soon as we hit the null individual in the preference ordering, the ordering “stops.”

A marriage market is defined by $(M, F; P)$, where

$$P = \{P(m_1), \dots, P(m_N); P(f_1), \dots, P(f_N)\}$$

is the collection of preferences in the population, $M = \{m_1, \dots, m_N\}$, and $F = \{f_1, \dots, f_N\}$. Then we have the following:

Definition 1 *A matching μ is a one-to-one correspondence from the set $M \cup F$ onto itself of order 2 (that is $\mu^2(x) = x$) such that $\mu(m) \in F$ and $\mu(f) \in M$. We refer to $\mu(x)$ as the mate of x .*

The notation $\mu^2(x) = x$ is read as $\mu(\mu(x))$, and just means that the mate of individual x ’s mate is individual x .

Definition 2 *The matching μ is individually rational if each agent is acceptable to his or her mate. That is, a matching is individually rational if it is not blocked by any (individual) agent.*

This is a weak concept, particularly in our application, since matched individuals will almost invariably be better off than unmatched individuals no matter what the quality level of their mate. A stronger notion is one of stability. Say that a matching μ has resulted in $\mu(m_i) = f_j$ and $\mu(f_k) = m_l$, but that male i strictly prefers f_k to f_j and female f_k strictly prefers m_i to m_l . Then the pair (m_i, f_k) can deviate from the matching assignment μ and improve their welfare. Such a match is unstable in the terminology of GS.

Definition 3 *A matching μ is stable if it is not blocked by any individual or any pair or agents.*

The main achievement of GS was to set out an algorithm for finding an equilibrium of the marriage game that was decentralized and constructive in the sense of establishing that at least one stable matching equilibrium exists. They assumed that preferences of agents were public information and a convention regarding the meeting and offering technology. Roth and Sotomayer (1990) devote considerable attention to the design of mechanisms that elicit truthful revelation of preference orderings when preferences are not public information, and also explore alternative meeting and proposal technologies. These important issues will be of less importance to us here given the nature of the application and the econometric and empirical focus of our analysis.

In our application a male individual i is characterized by the vector $m_i = (\alpha_{1i}, \delta_{1i}, \tilde{T}_{1i}, w_{1i}, y_{1i})$. His induced preference ordering over the females f_1, \dots, f_N is determined by R in the following manner. If m_i and f_j are matched, then the household is characterized by

$$S_{i,j} = m_i \cup f_j. \quad (5)$$

Then equilibrium time allocations in the household are given by

$$C_{ij}(R) = R(S_{ij}). \quad (6)$$

Given our assumptions regarding the form of the “payoff” functions to i and j , we can define the value to m_i of being matched with f_j under R as

$$\begin{aligned} V_i(j; R) = & \alpha_{1i} \ln(l_1^*(S_{ij}; R)) + (1 - \alpha_{1i}) \ln(\tau_1^*(S_{ij}; R)^{\delta_{1i}} \tau_2^*(S_{ij}; R)^{\delta_{2j}} \\ & \times (w_{1i} h_1^*(S_{ij}; R) + w_{2j} h_2^*(S_{ij}; R) + y_{1i} + y_{2j})). \end{aligned}$$

Given behavioral mode R , the preference ordering of i is given by

$$P(m_i|R) = f_{(1)}^i(R), f_{(2)}^i(R), \dots, f_{(N)}^i(R),$$

where

$$V_i(f_{(1)}^i(R); R) > V_i(f_{(2)}^i(R); R) > \dots > V_i(f_{(N)}^i(R); R).$$

Given knowledge of m_i , f_j , and R , the preference ordering of all population members is determined. This implies the following.

Definition 4 *A marriage market is defined by $(M, F; R)$.*

An equilibrium assignment is a function of marriage market characteristics. Then the set of stable matchings is determined by the characteristics vectors M and F and the behavioral model R , or $\Theta(M, F; R)$. Now there may exist, and generally do exist, multiple stable assignment equilibria. Among this set of equilibria, attention has focused on the two “extreme” stable matchings, the one that is most beneficial to men and the one most

beneficial to women.⁶ The GS matching algorithm, which they termed “deferred acceptance,” enables one to determine at least these two, of the many possible, equilibria in a straightforward manner. We describe the computation of the male-preferred equilibrium. In a given round,

1. Each male not tentatively matched with a female makes a marriage proposal to the woman he most prefers among the set of women who have not rejected a previous proposal of his. If he prefers the state of being single to any of the women in his choice set, he makes no offer.
2. Each woman (tentatively) accepts the proposal that yields the maximum payoff to her from the set of offers made to her during the round plus the value of the match with the offer she carries over from the previous round (she may reject one or more proposals because the option of remaining single dominates them). Any man whose offer is refused in the period cannot make another marriage proposal to the woman rejecting him in future rounds.
3. The process is repeated until no man makes a marriage proposal to any woman.

The female preferred stable matching equilibrium is found in the identical way after reversing the roles of two sexes as proposers and responders.

There may well exist other stable matchings besides these two. Given the generality of the preference structure, the size of the individual characteristic space, and the number of individuals in the marriage market in our empirical analysis (877), it is not possible to attempt to enumerate all possible stable matchings. We have computed the predicted marriage assignments using estimates of the state vectors m_i and f_j under the two R that we consider. We found that the same pairs were matched in over 96 percent of the cases in the male-preferred and female-preferred matchings. As a result, we use pairings from the male-preferred equilibria only in all of the empirical work that follows. The reader should bear in mind that other equilibria exist, even if they are not so different in metrics of concern to us in this exercise.

3 Econometrics

We consider estimation of the marriage market equilibrium in sequence. We begin with the issue of the estimation of (M, F) , the distribution of gender types. In this paper we do not treat the difficult censoring issues that arise when not all household members supply time to the labor market or in household production. Then, given that there are no corner solutions in the time allocation decisions with the household, we are able to posit that the entire vector

$$A_k = (h_{1k}, h_{2k}, \tau_{1k}, \tau_{2k}, w_{1k}, w_{2k}, y_{1k}, y_{2k}), \quad k = 1, \dots, N,$$

⁶When there is a unique equilibrium these stable matchings are identical, of course.

is observable by the analyst, where we have constructed the male and female indexing so that in the data male i is married to female i , $i = 1, \dots, N$, in the data. It will be useful to partition this vector into two subvectors,

$$\begin{aligned} A_k^1 &= (h_{1k}, h_{2k}, \tau_{1k}, \tau_{2k}), \\ A_k^2 &= (w_{1k}, w_{2k}, y_{1k}, y_{2k}), \end{aligned}$$

with A_k^1 representing the (endogenous) time allocations of household k and A_k^2 the observable (to the analyst) state variables. The unobservable state variables in household k are $(\alpha_{1k}, \alpha_{2k}, \tilde{T}_{1k}, \tilde{T}_{2k}, \delta_{1k}, \delta_{2k})$. As will become apparent soon, we will require further restrictions on the variability in the unobservable characteristics if we are to be able to nonparametrically identify the model. We will restrict the α_{ik} , $i = 1, 2$, $k = 1, \dots, N$, to have no variation within the population of males and females (individually), so that

$$\begin{aligned} \alpha_{1i} &= \alpha_1, \\ \alpha_{2i} &= \alpha_2, \quad i = 1, \dots, N. \end{aligned}$$

Assume that the values α_1 and α_2 are known, for now. Then denote the remaining unobserved household characteristics by

$$A_k^3 = (\tilde{T}_{1k}, \tilde{T}_{2k}, \delta_{1k}, \delta_{2k}).$$

The data used in the empirical work discussed below are drawn from the Panel Study of Income Dynamics (PSID). In keeping with the static setting of the model, we use data pertaining to household characteristics and time allocation decisions in one year, 2000. We chose this year because information on the time spent in household tasks is widely available for both spouses in that year.

We assume that the PSID is randomly drawn from the population distribution of married households in this year (which is an unlikely situation, admittedly), and that *all households in the population belong to one unified marriage market*. As we shall see below, this assumption is critical if we are to perform meaningful statistical analyses of the PSID data. Within this marriage market, assumed large, we consider the restrictive case in which there exists an equal number of males and females, with the stable match implying all agents are married. The characteristic vectors defining males and females, m and f , have associated distribution functions G_1 and G_2 , respectively. Since we have a random sample of households, we also have a random sample of household members *given* the marriage assignment rule.

Using a random sample of N households from the population marriage market, the first task is to estimate the distribution functions G_1 and G_2 . For household k , we can restate (1) as

$$A_k^1 = R(A_k^2 \cup A_k^3).$$

Proposition 5 *Assume all households in the population behave according to R , and that R is invertible in the sense that there is a unique value of A_k^3 such that*

$$A_k^3 = R^{-1}(A_k^1 \cup A_k^2) \quad (7)$$

for all values of $A_k^1 \cup A_k^2$. Then the distributions G_1 and G_2 are nonparametrically identified and can be consistently estimated.

Proof: Given knowledge and invertibility of R , then R^{-1} is a known function. If A_k^1 and A_k^2 are observed without error, then the vector A_k^3 is observable as well. Since the vectors A_k^1 and A_k^2 are observed for a random sample of households, then A_k^3 is as well. Define the vectors

$$\begin{aligned} X_k &= (A_k^3, w_{1k}, w_{2k}, y_{1k}, y_{2k}), \\ X_k^1 &= (\tilde{T}_{1k}, \delta_{1k}, w_{1k}, y_{1k}), \\ X_k^2 &= (\tilde{T}_{2k}, \delta_{2k}, w_{2k}, y_{2k}). \end{aligned}$$

The vector X_k^1 is an i.i.d. draw from G_1 and X_k^2 is an i.i.d. draw from G_2 . Then define

$$\begin{aligned} \hat{G}_1^N(x) &= N^{-1} \sum_{k=1}^N \chi(X_k^1 \leq x), \\ \hat{G}_2^N(x) &= N^{-1} \sum_{k=1}^N \chi(X_k^2 \leq x). \end{aligned}$$

Since $\{X_1^1, \dots, X_N^1\}$ and $\{X_1^2, \dots, X_N^2\}$ are both random samples from their respective populations, we know that

$$\text{plim}_{N \rightarrow \infty} \hat{G}_i^N(x) = G_i(x), \quad i = 1, 2,$$

by the Glivenko-Cantelli Theorem. ■

The following important implication immediately follows.

Proposition 6 *Let \mathfrak{R} be the set of equilibrium rules that determine time allocations in the household that are invertible in the sense of (7). Then all $R \in \mathfrak{R}$ are equivalent descriptions of sample information.*

Proof: Consider a household k in the sample. We observe four household choices $D_k^1 = (h_{1k}, h_{2k}, \tau_{1k}, \tau_{2k})$ and we have four unobservable characteristics of the spouses. Thus given any $D_k^2 = (w_{1k}, w_{2k}, y_{1k}, y_{2k})$ and any $R \in \mathfrak{R}$, there exists a unique vector of characteristics $(\tilde{T}_{1k}, \tilde{T}_{2k}, \delta_{1k}, \delta_{2k})$ that generate D_k^1 , or

$$D_k^1 = \Gamma(\tilde{T}_{1k}(R), \tilde{T}_{2k}(R), \delta_{1k}(R), \delta_{2k}(R) | D_k^2, R).$$

Then for any two $R, R' \in \mathfrak{R}$, $R \neq R'$,

$$\begin{aligned} & \Gamma(\tilde{T}_{1k}(R), \tilde{T}_{2k}(R), \delta_{1k}(R), \delta_{2k}(R) | D_k^2, R) \\ = & \Gamma(\tilde{T}_{1k}(R'), \tilde{T}_{2k}(R'), \delta_{1k}(R'), \delta_{2k}(R') | D_k^2, R'), \end{aligned}$$

which describes a correspondence between $(\tilde{T}_{1k}, \tilde{T}_{2k}, \delta_{1k}, \delta_{2k})(R)$ and $(\tilde{T}_{1k}, \tilde{T}_{2k}, \delta_{1k}, \delta_{2k})(R')$.

Consider any distance function

$$\mathbb{Q}(D_k^1, \hat{D}_k^1(\tilde{T}_{1k}, \tilde{T}_{2k}, \delta_{1k}, \delta_{2k} | D_k^2, R)),$$

where \hat{D}_k^1 is the predicted value of the household time allocations given the characteristics $(\tilde{T}_{1k}, \tilde{T}_{2k}, \delta_{1k}, \delta_{2k})$, D_k^2 , and R . But given invertibility

$$\begin{aligned} & (\tilde{T}_{1k}(R), \tilde{T}_{2k}(R), \delta_{1k}(R), \delta_{2k}(R) | D_k^2, R) \\ = & \arg \min \mathbb{Q}(D_k^1, \hat{D}_k^1(\tilde{T}_{1k}, \tilde{T}_{2k}, \delta_{1k}, \delta_{2k} | D_k^2, R)) \end{aligned}$$

and

$$\begin{aligned} \mathbb{Q}(D_k^1, \hat{D}_k^1((\tilde{T}_{1k}(R), \tilde{T}_{2k}(R), \delta_{1k}(R), \delta_{2k}(R) | D_k^2, R) | D_k^2, R)) = 0, \\ \forall R \in \mathfrak{R} \end{aligned}$$

■

Because of the flexible parameterization of spouses in terms of their types, if \mathfrak{R} contains more than one element there are multiple ways to “reparameterize” the data, in essence. The cardinality of \mathfrak{R} depends on assumptions made regarding the functional form of the utility and household production functions and the features of the data. Since the proof is not especially instructive, we simply state the following.

Proposition 7 *For Stone-Geary utility functions and the Cobb-Douglas home good production technology and for a population in which both household members supply time to the market, the Nash equilibrium and the symmetric Nash bargaining behavioral rules both belong to \mathfrak{R} .*

This proposition carries the important implication that it is not possible to determine whether household members (in the general population) operate under Nash equilibrium or Nash bargaining rules of behavior by observing only within household behavior. This “impossibility” result mainly results from the flexible specification of population heterogeneity. Clearly, by restricting the variability of these underlying parameters in the population, it will generally be possible to develop tests pitting the two forms of behavior against one another, but the outcome of such a test will be heavily dependent upon the parametric restrictions adopted.

3.1 Marital Sorting

Flexible specifications of population heterogeneity reduce the analyst's ability to derive distinguishable empirical implications from members of a class of modes of behavior. However, they do provide possibilities for developing tests based on marital sorting patterns.⁷ We explore the construction of such tests in this subsection.

We have assumed that our PSID sample of married individuals is drawn from a large population of married couples. Given the nature of the marriage equilibrium concept we are using, and side-stepping uniqueness issues, we have assumed that the households in our N household sample consist of a subset of husbands and wives who were matched under the GS deferred acceptance algorithm in the marriage market defined over all population members. In our sample of size N , without loss of generality, we index the male and female sample members so that male i is matched with female i under the male-preferred stable match μ , $i = 1, \dots, N$. Then we have the following result.

Proposition 8 *Define a random sample of N households matched under μ by M^N and F^N . Then the set of male-preferred stable matchings in the random sample matches male i with female i , $i = 1, \dots, N$.*

Proof: Let the male-preferred stable matching in the marriage market be given by μ . Begin by considering the case when $N = 1$. The only stable matching in the marriage submarket $A(1)$ is (m_1, f_1) . These are the only two individuals in the sub-market and they are acceptable to one another since they were acceptable to one another in the full marriage market and acceptability is a global property (i.e., independent of the choice set) since the option always exists to remain single. Due to the restricted choice set, there is only one stable matching.

Next consider the male-preferred stable matching for $N = 2$. To show that it must be $\{(m_1, f_1), (m_2, f_2)\}$, assume that the converse is true. For the male-preferred stable matching in the sub-market $A(1, 2)$ to be $\{(m_1, f_2), (m_2, f_1)\}$, one of the following strict preference orderings must hold:

$$\text{Case 1: } \begin{cases} p(m_1) & = & f_2 & f_1 \\ p(m_2) & = & f_1 & f_2 \\ p(f_1) & = & m_2 & m_1 \\ p(f_2) & = & m_1 & m_2 \end{cases}$$

$$\text{Case 2: } \begin{cases} p(m_1) & = & f_2 & f_1 \\ p(m_2) & = & f_2 & f_1 \\ p(f_1) & = & * & * \\ p(f_2) & = & m_1 & m_2 \end{cases}$$

⁷Marital sorting is but one phenomenon that could be used to distinguish between modes of intrahousehold behavior, of course. Others include divorce decisions and investments in marriage-specific capital.

$$\text{Case 3: } \begin{cases} p(m_1) &= & f_1 & f_2 \\ p(m_2) &= & f_1 & f_2 \\ p(f_1) &= & m_2 & m_1 \\ p(f_2) &= & * & * \end{cases}$$

where an ‘*’ indicates that the ordering of this agent’s preferences are irrelevant to the outcome.

Case 1 is not consistent with the stable matching μ in the global marriage market, since both pairs (m_1, f_2) and (m_2, f_1) would block μ . Case 2 is not consistent with μ because the pair (m_1, f_2) block, and Case 3 is not consistent with μ since (m_2, f_1) block. Since there always exists a stable male-preferred matching in $A(1, 2)$ in which all agents are matched, it must be $\{(m_1, f_1), (m_2, f_2)\}$, which is the pairing from μ .

The same argument is extended in a natural, albeit tedious, manner to the male-preferred stable matching in groups of (equilibrium) pairs larger than 2. The set of sub-market pairs not consistent with μ in male-preferred sub-market stable matching contains at least one pair that would block the stable match μ in the complete marriage market. Therefore the only male-preferred stable matching in any sub-market Γ is identical to the μ -pairing in the complete market. ■

Let us be clear what this result does and does not imply. In any sub-market, there may exist more than one stable matching. However, the male-preferred stable matching, which always exists, is the same as the male-preferred stable matching in the complete marriage market. Pairings in other stable matchings in the sub-market need not conform to the μ stable matching with which we are working.

On the positive side, the result gives us something vitally necessary to perform statistical analysis using matched pairs of observations, particularly when the dependent variables, so to speak, are the matches observed within sets of husbands and wives, or functions of those matches. In particular, in some of the statistical analyses we perform, we will work with random subsamples, or partitions, of the “complete” sample of size N . We denote a random subsample by Γ , where $\Gamma \subseteq \{1, \dots, N\}$, where the size of the subsample is given by the cardinality of Γ , denoted $\#\Gamma$. The result contained in the previous proposition clearly applies to any and all subsamples of the original sample of size N .

Definition 9 *A male-preferred stable matching μ has an Independence from Irrelevant Alternatives (IIA) property in the sense that*

$$\begin{aligned} f_i &= \mu(m_i), \quad i = 1, \dots, N \Rightarrow \\ f_i &= \mu_\Gamma(m_i), \quad i \in \Gamma \text{ for all } \Gamma \subseteq \{1, 2, \dots, N\}, \end{aligned}$$

where μ_Γ is the male-preferred stable matching in subsample Γ .

The IIA property of μ is crucial if we are to have a coherent sampling theory.⁸ The sampling theory underlying the statistical analysis is developed as follows. We consider the

⁸The issue being considered here, which is the impact of sampling a complete choice on the assessment

marriage market to consist of N individuals of each gender, with N arbitrarily large. The male-preferring stable matching in this market is given by μ . Sample elements, as is the case in the PSID, are defined as households, and our universe is married households within the complete marriage market, that is, all PSID households are assumed to be drawn from the same population marriage market.⁹

3.2 Choosing Between Alternative R

We look at the ability of either R to predict in-sample matches using three “methods.” The first is purely descriptive, and involves computing the rank order correlation between the predicted marriage partners under the behavioral rules and the actual marriage partners. Since the model does not contain any random elements, if we restrict our attention to the NE and NB rules, one of them should fit perfectly and, unsurprisingly, neither does. The setup we have developed may still be able to produce a perfect correspondence between the observed and observed matches if there exists an $\hat{R} \in \mathfrak{R}$ such that $\mathbb{Q}(\Theta_0, \Theta_{MP}(\hat{M}(\hat{R}), \hat{F}(\hat{R}); \hat{R})) = 0$, where Θ_0 denotes the observed marital sorting pattern and $\Theta_{MP}(\hat{M}(\hat{R}), \hat{F}(\hat{R}); \hat{R})$ denotes the male-preferring GS stable marriage sorting under rule \hat{R} . Since it seems difficult to constructively characterize the set \mathfrak{R} , this does not appear to be a promising avenue to follow.

To bring randomness into the model, we allow for measurement error in wages. In particular, we assume that the distribution of the measurement error is known (more on this assumption below), and that the logarithm of observed wages is related to the logarithm of true wages by

$$\ln \tilde{w}_{sk} = \ln w_{sk} + \varepsilon_{sk}, \quad s = 1, 2,$$

where w_{sk} is the true wage of spouse s in household k , \tilde{w}_{sk} is the reported wage of spouse s in household k , and the measurement error ε_{sk} is an independently and identically distributed across households and spouses within households. In order to generate “true” wages based on the observed wage rates, it is necessary for us to make a functional form assumption regarding the distribution of ε_{sk} , and, as is common, we assume normality. One of the principle reasons we have chosen to add measurement error in wages is the availability of

of choice probabilities, is reminiscent of the analysis conducted by McFadden (1978). Using a multinomial logit structure, he demonstrated that consistent estimators of choice probability parameters could be formed using data on restricted choice sets. As is the case here, the primary motivation for sampling large choice sets was computational tractability.

⁹Assume that individuals were drawn from two separate marriage markets, with no information as to the market membership of any sample household. Say that the male-preferred stable matching in marriage market i is given by μ^i . Then if the first 5 households drawn in any sample are all from market 1, say, the male-preferred stable matching for that sub-population would be consistent with μ^1 . But say a 6th household is drawn, and that household is from market 2. Then, the male-preferred stable matching for the subpopulation consisting of the 6 households cannot be compared to either μ^1 or μ^2 since there are different group members, and there is no well-defined correspondence between the three male-preferred stable matchings.

high quality estimates of the measurement error variance in the logarithm of wages in the PSID. Using a special validation survey performed in the 1980s that involved administering the standard PSID survey instrument to a group of workers at a large factory in the Detroit area, Bound et al (1994) were able to get reasonably precise estimates of measurement error in wage reports by comparing subject responses with payroll records. In line with estimates of the variance of ε_s they obtained (see their Table 3), we set $\sigma_s^2 = .13$ for both husbands and wives. Note that since we are working in wage levels, we have

$$\tilde{w}_{sk} = w_{sk} \exp(\varepsilon_{sk}).$$

Since ε_{sk} is distributed as a mean 0 normal with variance 0.13, the measurement error in wages has a lognormal distribution with mean 1.067 and variance 0.158.

Before describing the implications of measurement error for equilibrium marriage patterns, we briefly consider its effect on estimation of the distribution of characteristics of husbands and wives in existing households. In principle, the wage rate w_{sk} is observed, and therefore is not a function of the behavioral rule. Knowledge of the w_{sk} and Y_{sk} , along with the values of the time allocation decisions, allows us to determine the values of $\tilde{T}_{sk}(R)$ and $\delta_{sk}(R)$. We can write

$$A_k^3(R) = R^{-1}(A_k^1, y_{1k}, y_{2k}, w_{1k}, w_{2k}), \quad (8)$$

where we recall that $A_k^3(R) \equiv (\tilde{T}_{1k}(R), \tilde{T}_{2k}(R), \delta_{1k}(R), \delta_{2k}(R))$. Under the measurement error assumptions, the true wage of spouse 1 is $w_{1k} = \tilde{w}_{1k} \exp(-\varepsilon_{1k})$ and of spouse 2 is $w_{2k} = \tilde{w}_{2k} \exp(-\varepsilon_{2k})$. Then we rewrite (8) as

$$A_k^3(R) = R^{-1}(A_k^1, y_{1k}, y_{2k}, \tilde{w}_{1k} \exp(-\varepsilon_{1k}), \tilde{w}_{2k} \exp(-\varepsilon_{2k})).$$

Then define

$$\begin{aligned} X_k(\varepsilon_k) &= ((A_k^3(R)(A_k^1, y_{1k}, y_{2k}, \tilde{w}_{1k}, \tilde{w}_{2k}, \varepsilon_{1k}, \varepsilon_{2k}), \tilde{w}_{1k}, \tilde{w}_{2k}, \varepsilon_{1k}, \varepsilon_{2k}, y_{1k}, y_{2k}) \\ X_k^1(\varepsilon_k) &= ((\tilde{T}_{1k}, \delta_{1k})(A_k^1, y_{1k}, y_{2k}, \tilde{w}_{1k}, \tilde{w}_{2k}, \varepsilon_{1k}, \varepsilon_{2k}), \tilde{w}_{1k}, \varepsilon_{1k}, y_{1k}), \\ X_k^2(\varepsilon_k) &= ((\tilde{T}_{2k}, \delta_{2k})(A_k^1, y_{1k}, y_{2k}, w_{1k}^*, w_{2k}^*, \varepsilon_{1k}, \varepsilon_{2k}), \tilde{w}_{2k}, \varepsilon_{2k}, y_{2k}). \end{aligned}$$

Under the measurement error assumption, we have redefined the vector X_k^s to include the measured wage of spouse s as opposed to the actual wage. We think of X_k^s as being conditional on the measurement error draws of both the spouses, ε_{1k} and ε_{2k} . The estimator of the unconditional distribution of the characteristics $(\tilde{T}_{sk}, \delta_{sk}, w_{sk}^*, y_{sk})$ for household k is then given by

$$\begin{aligned} \hat{G}_1^N(x) &= N^{-1} \int \cdots \int \sum_{k=1}^N \chi(X_k^1(\varepsilon_{1k}, \varepsilon_{2k}) \leq x) d\Phi\left(\frac{\varepsilon_{11}}{\sigma_1}\right) d\Phi\left(\frac{\varepsilon_{21}}{\sigma_2}\right) \cdots d\Phi\left(\frac{\varepsilon_{1N}}{\sigma_1}\right) d\Phi\left(\frac{\varepsilon_{2N}}{\sigma_2}\right), \\ \hat{G}_2^N(x) &= N^{-1} \int \cdots \int \sum_{k=1}^N \chi(X_k^2(\varepsilon_{1k}, \varepsilon_{2k}) \leq x) d\Phi\left(\frac{\varepsilon_{11}}{\sigma_1}\right) d\Phi\left(\frac{\varepsilon_{21}}{\sigma_2}\right) \cdots d\Phi\left(\frac{\varepsilon_{1N}}{\sigma_1}\right) d\Phi\left(\frac{\varepsilon_{2N}}{\sigma_2}\right). \end{aligned}$$

As was the case without measurement error, the distribution function estimators are consistent by the Cantelli-Glivenko theorem. While the marginal distributions of the estimators of y_s remain step functions, integrating over the measurement error distributions results in smooth estimators of the marginal distributions of the true wage and the unobserved individual characteristics \tilde{T}_s and δ_s .

3.3 Computation of α_s

To this point we have assumed that the preference weight on leisure varies only by gender (i.e., all individuals of the same gender share the same value of α_s) and we have treated it as known. The four first order conditions uniquely determine the four unobserved characteristics of the husband and wife conditional on a behavioral rule R and α_1 and α_2 . We determine values of α_s after adopting a particular normalization.

To stress the dependence of the implied values of the time endowments in the household on the preference weights α_1 and α_2 , write the implied time endowment for individual of gender s in household i as

$$\tilde{T}_{si}(R; \alpha). \quad (9)$$

There are 168 hours in a week. We define the values of $\hat{\alpha}_s$ as those that result in the *average* time endowment in the sample being equal to 168, or

$$168 = N^{-1} \sum_{i=1}^N \tilde{T}_{1i}(R; \hat{\alpha}) \quad (10)$$

$$168 = N^{-1} \sum_{i=1}^N \tilde{T}_{2i}(R; \hat{\alpha}). \quad (11)$$

The use of the average is admittedly somewhat arbitrary, and an argument could be made for using the median, for example, instead. Nonetheless, given the parameterization of the model adopted, some such normalization is required if we are to “estimate” the two values α_1 and α_2 .

3.4 Assessing the Relative Performance of the Two Behavioral Assumptions

We now turn to the predictive part of the exercise, and describe the three measures of fit we consider.

3.4.1 Rank Order Correlation

The most straightforward comparison of the predictive abilities of the two R we consider uses a rank correlation metric. For this comparison we assume that wages are correctly measured. As a result, there is no randomness in the model that is consistent with the

rank order correlation being less than one for one of the two R , if the true state of the world is, in fact, either NE or NB . That is why this measure cannot be used as the basis for constructing a formal statistical analysis comparing the two R . Since neither fits the observed match pattern perfectly (in fact, far from it), both can be rejected as the true state of the world if the model is correctly specified and all individual characteristics are measured without error.

3.4.2 Match Prediction using Sample Subsets

Given an N married household sample, there exist $2N$ measurement errors associated with all of the measured wage rates. Given independence of these shocks across households as well as across spouses, it is conceptually straightforward to express the probability that a given observed pattern of sorts was generated under any of our alternative behavioral models R . To simplify notation, let

$$\Theta_{MP}(M, F, R|\varepsilon) \tag{12}$$

denote the marital sorting pattern given measured characteristics M and F , behavioral rule R , and measurement errors ε . The observed marital sorting pattern is given by Θ^0 . Then over the sample of size N , the probability that the observed marriage pattern is generated by R is

$$p^N(R) = \int \cdots \int \chi[\Theta^0 = \Theta_{MP}(M, F, R|\varepsilon)] d\Phi\left(\frac{\varepsilon_{11}}{\sigma}\right) d\Phi\left(\frac{\varepsilon_{21}}{\sigma}\right) \cdots d\Phi\left(\frac{\varepsilon_{1N}}{\sigma}\right) d\Phi\left(\frac{\varepsilon_{2N}}{\sigma}\right), \tag{13}$$

where we have restricted the standard deviation of measurement errors to be the same across genders (i.e., $\sigma_1 = \sigma_2 = \sigma$). It is not immediately apparent that a given Θ^0 can be generated by *any* draw of ε given (M, F, R) . In this case, $p^N(R) = 0$ and no further consideration of the rule R is warranted.

In computing $p^N(R)$ we face a computational problem stemming from the fact that there is no closed form expression for the integral in (13). We adopt a Monte Carlo integration approach, in which we take $2N$ independent draws from a mean-zero normal distribution with standard deviation σ over M replications. Our estimate of $p^N(R)$ is then given by the proportion of the M replications that resulted in the observed distribution of marital sorts. More formally, let the m^{th} draw of the $2N$ measurement errors be denoted ε^m . Then

$$\hat{p}_M^N(R) = M^{-1} \sum_{m=1}^M \chi[\Theta^0 = \Theta_{MP}(M, F, R|\varepsilon^m)]. \tag{14}$$

Consistency of the Monte Carlo integration estimator in this case requires M grow indefinitely large, or

$$\text{plim}_{M \rightarrow \infty} \hat{p}_M^N(R) = p^N(R)$$

Computation of this quantity is conceptually and numerically straightforward. However, the size of M required to adequately approximate $p^N(R)$ will depend critically on the size of the married population in the sample. For example, say M is set at 10000. If $N = 10$, we may expect to observe a nontrivial number of correspondences between the predicted matches under R and the observed marriage sorts if R is indeed the correct behavioral rule. However, even if households behave according to R , we would expect the likelihood that a sample of M draws yields the observed sorts to be arbitrarily close to 0 if N is equal to 10 million. We circumvent this problem by subsampling our group of 877 households in the following manner.

From the original sample of N households, randomly select J groups of size n . Let the j^{th} grouping of households selected n at a time be denoted by C_j^n . The groups are selected, with replacement, from the size N sample subject to the condition that no household appears more than once within any size n group. For example, for $N = 200$, $n = 2$, the first group selected, denoted by C_1^2 , might be composed of households $\{5, 173\}$. The second group defined can be $C_2^2 = \{5, 140\}$, for example, but cannot be $\{5, 5\}$. We want to preclude replication of households in the same “choice set” because this would violate the strict preference orderings over alternatives that we have assumed in defining marriage market equilibrium.

For each of the J groups, we then take M replications of $2n$ independent draws from a mean-zero normal distribution with standard deviation σ . Denote the m^{th} draw of the vector ε in group j by $\varepsilon^m(j)$. Since the subsamples are randomly drawn, we think of the proportion of correct picks in subsample j as being an estimate of the probability of correct sorting predictions in a randomly selected set of n households from the marriage market under behavioral rule R . First, denote the estimate of this probability in subgroup j by

$$\hat{p}_M^n(j; R) = M^{-1} \sum_{m=1}^M \chi[\Theta_j^0 = \Theta_{MP}(M(j), F(j), R|\varepsilon^m(j))],$$

where Θ_j^0 is the observed marital sorting pattern in subgroup j , and $M(j)$ and $F(j)$ are the characteristics of men and women in subgroup C_j^n . Then define the estimator of the probability of correctly predicting the actual marriage outcomes in a random sample containing n spousal pairs by

$$\hat{p}_M^n(R) = J^{-1} \sum_{j=1}^J \hat{p}_M^n(j; R).$$

Clearly $\text{plim}_{M \rightarrow \infty} \hat{p}_M^n(R) \neq \text{plim}_{M \rightarrow \infty} \hat{p}_M^{n'}(R)$ for arbitrary choices of n and n' , even as $N \rightarrow \infty$. Neither is it possible to explicitly characterize the relationship between these two quantities. As a result, we compute $\hat{p}_M^n(R)$ for four different values of n , 2 through 5. For each value of n , we set $J = 1000$. While the values of $\hat{p}_M^n(R)$ vary greatly across n , the relationship between $\hat{p}_M^n(NE)$ and $\hat{p}_M^n(NB)$ displays a great deal of regularity over the four values of n .

Since the subsets are randomly selected, and because we have a large number of them, we consider the distribution of the $\hat{p}_M^n(j; R)$, $j = 1, \dots, J$, as representing the sampling distribution of the proportion of correct predictions in a groups of size n . Note that there are two sources of randomness, given the complete sample information, that generate the dispersion in the $\hat{p}_M^n(j; R)$. The first comes from the random composition of the groups, and the second from the measurement error shocks.

We first find the distribution of each of the group size n sampling probabilities for the J groups under each of the two rules. From these we compute the sample average and sample variance, which are consistent estimators of the corresponding population qualities given our random (sub)sampling assumption. From the sampling distribution of the differences,

$$\hat{d}_M^n(j) \equiv \hat{p}_M^n(j; NE) - \hat{p}_M^n(j; NB),$$

we can consistently estimate the average difference in predictive ability under the two decision rules as well as the variance of the average difference. Since J is relatively large for any of the n we consider here, and since the C_n^j are considered i.i.d. draws from the population of sub-marriage markets of size n , we invoke the central limit theorem to determine whether there exists a statistically significant difference in the predictive ability of the two R . The sample mean of the differences by $\hat{d}_M^n = J^{-1} \sum_j \hat{d}_M^n(j)$, and the (estimated) standard error of the difference by $\widehat{s.e.}(\hat{d}_M^n) = J^{-1} \sqrt{\sum_j (\hat{d}_M^n(j) - \hat{d}_M^n)^2}$. Then under the null of no difference in predictive ability, $\hat{d}_M^n / \widehat{s.e.}(\hat{d}_M^n)$ is approximately distributed as a standard normal random variable. If the absolute value of $\hat{d}_M^n / \widehat{s.e.}(\hat{d}_M^n)$ is sufficiently large to cast doubt on the validity of the null, the evidence will favor the model that provides the best correspondence to the observed marital sorts.

The prediction metric we use here is the same as that utilized in maximum score estimation, and a recent application of this estimator to the bilateral matching problem (with transferable utility) is considered in Fox (2006). In his model there exist free parameters, which are not present in our analysis of the marriage market equilibrium. If we allowed there to exist free parameters that characterized the marriage market and that do not appear in the payoffs of household members under a given rule R , a maximum score estimator could be implemented using the entire sample of matches rather than the subsamples we use here. But the main objective our exercise is model selection, as it were, in a tightly specified model of household behavior and marriage market characteristics in which no free parameters appear in our prediction metric. This is what distinguishes the approach here from that of Fox.¹⁰

¹⁰An implementation of the Fox-type estimator in our context could be the following. Assume a parametric form for the distribution of individual characteristics, with the parameter vector characterizing these two marginal distributions given by Λ . Then, assuming a rule R , find the set $\hat{\Lambda}$ of values that maximize the number of correct predictions. This one step estimator and the associated value of the objective function could be used to compare fit under various values of R . The downside to its use is the necessity of selecting a parametric form for the marginal distributions of male and female characteristics.

3.4.3 The Assortative Mating Metric

It is common to characterize matching equilibria in terms of the association of observed characteristics among spouses. In our application, it is most natural to focus on wage rates, since these are observed for all individuals, albeit with measurement error.¹¹ As was discussed above, it is infeasible to compare the two rules in terms of their explicit matching predictions when using the entire sample. However, it is possible to compare their ability to generate stable match patterns of association between observable spousal characteristics across all matches in the sample with what is observed in the data. We will use the simple (zero-order) correlation between spousal wages to characterize assortative mating.

Denote the wage correlation by under the male-preferred stable matching R by $\rho(R)$, and in the data by ρ_D . Now the value ρ_D should not be thought of as the zero order correlation from the actual spousal wage rates, since by assumption they are measured with error. Thus ρ_D is the correlation between the measured wages in the data over the 877 cases.

To generate the correlation of *measured* wages under the model we proceed as follows. Under a given R , generate $2N$ measurement error draws, one for each spouse in the total sample. Denote one of these measurement error vectors by $\varepsilon(m)$. We then compute the true wage for spouse s in household k from $w_{sk} = \tilde{w}_{sk} \exp(-\varepsilon_{sk})$, and back out all of the implied characteristics of all sample members based on the “true” wages, nonlabor incomes, and time allocation decisions given R and the measurement error vector $\varepsilon(m)$. We then apply the GS algorithm to obtain the male-preferred stable matching under R and $\varepsilon(m)$, and based on these matches, we compute the correlation in *measured wages* between the spouses. Denote this correlation by $\rho_m(R)$.

We repeat this procedure for M draws of the measurement error vector. The comparison we wish to make between the observed and predicted level of assortative mating in wages is based on a fixed sample of size N from the marriage market. Then the sampling distribution of $\rho_m(R)$ we are interested in treats *only* the measurement errors as the source of randomness. From the empirical distribution of $\{\rho_m(R)\}_{m=1}^M$, we can construct Monte Carlo confidence intervals in the standard way. A confidence interval that contains the “true” correlation of measured spousal wages with probability ν has a lower limit $\hat{F}_R^{-1}(\nu/2)$ and an upper limit of $\hat{F}_R^{-1}(1 - \nu/2)$, where \hat{F}_R denotes the Monte Carlo distribution of measured spousal wage correlations under behavioral rule R . We use independent measurement error shocks under the two behavioral rules, and then compare the ν confidence intervals on two, interrelated dimensions. We begin by examining the degree to which the confidence intervals intersect. If there exists a large amount of overlap, then clearly this metric is not very useful in not in distinguishing between the two hypotheses. Secondly, we see which, if any, of the confidence intervals includes the actual correlation

¹¹Nonlabor income levels are also observed for the spouses, but these are close to zero for most individuals. Moreover, it is difficult to assign household nonlabor income to individual spouses in many cases, a problem that does not exist with respect to wages.

observed in the data. If one does, we say that this behavioral rule is consistent with the pattern of assortative mating (on wages) observed in the data.

This metric has the advantage that it can be computed using all sample observations simultaneously, rather than small subsets of the sample, as was the case for the proportion of correct predictions metric described above. As we shall see, both the *NE* and *NB* rules perform respectably using this metric, which is not so clearly the case under the more “direct” prediction metric.

4 Empirical Results

The empirical work is performed using a sample of married couples taken from the PSID. The data contain information on household characteristics in 2000 that were collected in the 2000 and 2001 survey years. To be included in the sample, the household must have been headed by a married couple, both of whom were between the ages of 25 through 49, inclusive. All information on time allocations within the household must have been available for both spouses; this consists of the average amount of time spent in the labor market per week in 2000 as well as average hours spent in housework per week. Because household production activities change so markedly when young children are present, we excluded all households in which there was a child less than six years of age.

We also excluded any household in which one of the spouses made more than \$150 an hour or reported more than 80 hours of market work per week. We also required that the household not receive more than \$1000 per week in nonlabor income. A few households reported negative total income for the year, and these were excluded.

The (almost) final selection criterion imposed was that both spouses spend time in the labor market and in home production. This, of course, is a substantive restriction that is imposed so that we can invert four first order conditions for each household to obtain four values of the unobserved characteristics of the spouses (two for each spouse). Approximately 18 percent of the sample was eliminated by insisting that both spouses report supplying time to the market in the previous year. Some spouses were reported to have supplied zero time to household production; for these individuals we assumed that the actual amount of time spent in housework was 1 hour per week.¹² During the process of estimation we found that data from 9 households in our “final” sample produced problematic values for the four unobserved household characteristics. We excluded these from all further analyses. The total sample size with which we work is $N = 877$.

Table 1 contains descriptive statistics for our sample. We think of the decision period as the week. The unit of time is the week, and all monetary units are expressed in terms of

¹²It would be interesting to look at the distribution of responses to these housework questions as a function of the identify of the respondent. We hazard the conjecture that, conditional on observable characteristics, respondents are likely to over-emphasize their contributions to the household workload while under-emphasizing the spouse's.

current (year 2000) dollars. For now we focus only on the means and variances of variables taken directly from the data.

The average wage of husbands is about 40 percent greater than the average wage of wives. They work about 20 percent more hours per week in the labor market than their wives do, on average, while their wives supply about twice as much time in housework. It is interesting to note that the average total time spent in the labor market and performing household tasks is essentially identical for husbands and their wives. This appears to be a common finding in empirical studies of labor supply and housework.

The average nonlabor income per household is 120 dollars per week, with a large standard deviation. Nonlabor income of less than 100 dollars is reported by two-thirds of the households in the final sample. Recall that we have excluded households in which nonlabor income exceeded 1000 dollars per week.

The first task performed was to back out the implied values of $(\tilde{T}_{si}, \delta_{si}, \alpha_s)$, $s = 1, 2$, $i = 1, \dots, N$, under *NE* and *NB*. The means and standard deviations of these characteristics are presented in Table 1. We see that the preference weights on leisure are far greater under *NB*. This is to be expected since cooperative behavior will lead to a greater supply of time to the market and household tasks for a given set of household characteristics. Thus, to be consistent with the same observed time allocations, the leisure weights under Nash bargaining must be greater than those computed under Nash equilibrium. The normalization of the mean time endowments results in this value being equal to 168 for both sexes and under either behavioral mode. The large standard deviation of \tilde{T}_s indicates substantial heterogeneity in this characteristic in the sample.

The average value of efficiency in household production varies across the genders and the modes of behavior. For the same reason that *NB* led to higher imputed preference weights on leisure, it also leads to lower values of the household production elasticities for both sexes. For both sexes, the average value of the Cobb-Douglas parameter under *NB* is about one-third of its average value under *NE*.

There are large changes in the means of α_s and δ_s when moving from *NE* to *NB*, and in the standard deviations of \tilde{T}_s and δ_s . Nevertheless, as Figures 1 and 2 and Table 2 illustrate, the imputed values of $(\tilde{T}_{si}, \delta_{si}, \alpha_s)$ computed under *NB* are linear transformations of the values computed under *NE*.¹³ In spite of this extreme dependence of the parameter values computed under the two behavioral rules, the preference orderings and resulting marital sorts can be very different, as we shall see below.

It may be of some interest to investigate the gains to cooperative behavior and “rational” marriage sorts starting from the noncooperative baseline. We perform an experiment that utilizes our parameter estimates under *NE* and first computes the welfare gains to existing households if they switched their behavior to *NB*. We then look at the change in welfare that would result if all households continued to behave noncooperatively, but were

¹³The small deviations from this claim that we see in Table 2 result from numerical inaccuracies involved in performing the inversion of the first order conditions in the Nash bargaining case.

matched according to the GS algorithm.

Table 3 and Figure 3 contain the results of this exercise. By definition, when existing households switch to Nash bargaining, there is a welfare gain for all husbands and wives in the sample. However, the welfare gains are small, raising the utility levels of husbands by less than 1 percent and those of wives by 1.2 percent. While these small gains are specific to the cardinal utility function we have assumed, they do line up with similar analyses involving divorced parents that are reported in Del Boca and Flinn (1994). They also point to the fact that cooperative behavior may not be “efficient” if implementing cooperative outcomes is more costly than simply employing best-response strategies (Del Boca and Flinn (2006)).

Percentage gains in welfare are also small, on average, when individuals are resorted using the GS algorithm (they are assumed to behave noncooperatively both in the baseline case and after being resorted). Unlike the switch from *NE* to *NB* behavior for fixed households, in this experiment there will be winners and losers. However, we find that average welfare increases for husbands are identical to those recorded in the first experiment, while they are about one-half as great for wives. Roughly speaking, the scope for welfare improvements is about as great for marital reshuffling as it is for moves to cooperative behavior. In neither case are they large given the cardinal utility measures employed here.

We have now reached the main focus of the empirical analysis. Which behavioral assumption is most consistent with the observed patterns of marital sorts? The short answer is that there is no clear cut winner, though we will conclude that the evidence presented here is slightly more supportative of one of them.

We begin with the simple rank order correlation between actual marriages and the predicted ones, in the case in which wages are assumed to be measured without error. Table 4 contains the rank order correlations between observed, *NE*, and *NB* equilibrium sorts. We notice that even though there exists a linear mapping between unobserved parameters characterizing individuals computed under *NE* and *NB*, there is only a rank order correlation of 0.028 of the marital sorts under these two models. While the correlation between the observed sorts and that predicted under *NE* is only 0.015, the correlation between observed sorts and those predicted under *NB* is a relatively strong and “perverse” -0.063. Thus neither model provides a good fit to the data, though there is no statistical basis for us to form a formal measure of fit in this deterministic world.

The second prediction metric yields more interesting results. We computed the average proportions of correct predictions, and the standard deviations, across 1000 sub-marriage markets of size $n = 2, \dots, 5$. The results are presented in Table 5. The baseline we use to assess the success of the models is the probability that the individuals would be correctly matched by a purely random mechanism, such as flipping coins. For example, when $n = 2$, by flipping a fair coin we will correctly match male 1 with female 1 50 percent of the time. For a sub-market with n individuals, the probability of correctly matching all the individuals using a random assignment mechanism is $1/n!$

We see that there are no noticeable differences between the average correct predictions

of NE and NB . The last column of the table reports the difference in this proportion and the standard deviation of the difference. The standard deviation is often close to an order of magnitude larger than the difference, which is always positive. That is, NE seems to slightly outperform NB for the size n we have examined, but in no case does it remotely appear as though the difference is “important.”

As we increase n , the predictive ability of both models falls, naturally enough. It is interesting to note, however, that each perform increasingly better relative to the random matching baseline as sub-market size increases.

The most conclusive results we obtained were for the assortative mating metric; the results of this exercise are reported in Table 6. Recall that the correlation of measured wages between spouses from the data was 0.285. As discussed in the preceding section, we drew two samples of measurement error draws to evaluate the level of assortative mating under NE and NB , so that the wage correlations would be independent across the two evaluations. The first line under Sample 1 and Sample 2 reports the average of the correlations between observed wages allowing for the measurement error shocks, and we see that the mean and standard deviation of the distribution of correlation coefficients is very similar in the two samples of draws. The second line in each panel reports the mean correlation in measured wages under the male-preferred stable matches for the NE and NB rules, as well as the standard deviation. The last two columns of the table report the upper and lower bound of the Monte Carlo confidence intervals corresponding to the 0.05 probability level.

We first note that the wage correlation generated under either R is reasonably close to that observed in the data, even though this metric was never directly used in obtaining estimating the model. The average correlation in wages under the predicted sorting from NB is especially close to the correlation from the data. The average wage correlation implied by the model in this case is 0.243, which differs from the data only by 0.042. The average wage correlation under NE is 0.174, instead. Given that each has a Monte Carlo standard error of 0.022, we might conclude that the NB correlation is appreciably closer to the sample value than is the wage correlation generated under NE .

This can be more formally stated using the Monte Carlo confidence intervals. First note that, while the confidence intervals do overlap, the set of values of the correlation coefficient that belong to one confidence interval but not the other is of significant size. The behavioral rules are clearly more distinguishable under this metric than they were over the prediction metric just discussed.

We see that at the 0.05 probability level, neither wage correlation confidence interval contains the observed value of 0.285, though it is not far from the upper bound of the NB confidence interval, 0.272. If we use a probability level of 0.01, the sample value is included in the NB confidence interval, but not within the one associated with NE . On the basis of this evidence, and the lack of strong support in favor of one or the other under the other two metrics, our conclusion is that time allocations and marital sorting patterns are more consistent with the hypothesis that all households make decisions consistent with NB than

with NE .

5 Conclusion

In this paper we have attempted to make the point that there is no general nonparametric test to distinguish between modes of household behavior when individual heterogeneity in unobservable and observable characteristics is not introduced in severely restrictive ways. Using a flexible specification means that within household behavior is not useful in distinguishing between competing modes of behavior, which is the negative conclusion we draw. The good news is that this heterogeneity does produce interesting implications regarding the assignments of husbands to wives in equilibrium, and that these can be exploited in investigating the mode of behavior followed by population members. Using the Gale-Shapley bilateral concept of stable matchings, we developed two metrics with which the competing hypotheses of Nash bargaining and Nash equilibrium could be compared. Under the first, which measured the ability of each behavioral hypothesis to generate predictions consistent with the observed matches in distinct, small marriage sub-markets, there was no distinguishable difference between the two. Under the second, which measured the ability of each behavioral rule to generate spousal wage correlations consistent with those found in the data, the Nash bargaining hypothesis performed distinctly better than that of Nash equilibrium.

The general methodological point we stress is reminiscent of the general problem of model over-fitting. We adopted a model framework that was capable of perfectly fitting the data (i.e., the mapping from the data space to the parameter space was 1 to 1) under an entire class of behavioral rules \mathfrak{R} . In order to “test” one specification against another, some restrictions have to be imposed on the parameterization to make the mapping no longer 1 to 1, and to raise the possibility that one of the elements of \mathfrak{R} “fits” better than another. Of course, the test results we obtain in the end are a function of sample realizations and the restrictions we have placed on the parametric specification of individual utilities and the household production technology. It is, of course, seldom possible to claim that one parameterization should be preferred over another on theoretical grounds.

Given this inherent arbitrariness, we have moved the test to a different playing field - one that is “out of sample,” so to speak. The richness of the specification of individual heterogeneity leads to zero power in testing one element of \mathfrak{R} against another using only time allocation data, but has the potential to produce the implication of very different marital sorts - an empirical phenomenon that is not used in backing out the individual characteristics. In this application, we believe that we have generated some evidence that households behave in a cooperative manner. Those advocating the “sharing rule” approach to the analysis of household allocation decisions posit efficient allocations as a fundamental identification condition. In this paper, we think we have provided some evidence to support their assumption, though of course the best way to specify the manner in which surplus is

distributed between the spouses remains an open question.

Table 1
Means and (Standard Deviations) of Individual Characteristics
N = 877

<i>Characteristic</i>	<i>Husband</i>		<i>Wife</i>	
	<i>NE</i>	<i>NB</i>	<i>NE</i>	<i>NB</i>
α	0.563	0.715	0.467	0.655
\tilde{T}	168.000 (58.637)	168.000 (50.532)	168.000 (70.130)	168.000 (57.139)
δ	0.101 (0.097)	0.027 (0.031)	0.139 (0.109)	0.045 (0.037)
w	21.522 (13.655)		15.206 (9.434)	
h	45.707 (8.421)		38.202 (10.569)	
τ	7.853 (6.878)		15.323 (9.672)	
Y	120.455 (183.175)			

Table 2
Correlation Between Imputed Parameters

<i>Nash Equilibrium</i>	<i>Nash Bargaining</i>			
	\tilde{T}_1	\tilde{T}_2	δ_1	δ_2
\tilde{T}_1	1.000	-0.172	-0.137	0.070
\tilde{T}_2	-0.175	1.000	0.160	-0.255
δ_1	-0.166	0.141	0.993	0.097
δ_2	0.066	-0.256	0.108	0.998

Table 3
Changes in Average Welfare Values from
NE Baseline Behavior and Observed Matches
(Proportionate Gain from Baseline)

	Husbands	Wives
Baseline	6.103	6.396
NB Behavior	6.159 (0.009)	6.473 (0.012)
NE Marriage	6.158 (0.009)	6.431 (0.005)

Table 4
Correlations Between Marriage Sorts

	<i>Actual</i>	<i>Nash Equilibrium</i>	<i>Nash Bargaining</i>
<i>Actual</i>	1.000	0.015	-0.063
<i>Nash Equilibrium</i>		1.000	0.028
<i>Nash Bargaining</i>			1.000

Table 5
Proportion of Correct Predictions
(Standard Deviation)
J = 1000

<i>Group Size</i>	<i>“Random”</i>	<i>Nash Equilibrium</i>	<i>Nash Bargaining</i>	<i>Difference</i>
2	0.500	0.607 (0.410)	0.602 (0.390)	0.016 (0.147)
3	0.167	0.287 (0.346)	0.291 (0.333)	0.006 (0.135)
4	0.042	0.111 (0.211)	0.131 (0.220)	0.010 (0.098)
5	0.008	0.034 (0.106)	0.028 (0.080)	0.009 (0.060)

Table 6
Spousal Wage Correlations
M = 200

<i>Wage Correlation</i>	<i>Mean</i>	<i>St. Dev.</i>	<u><i>MCCI</i></u> (0.05)	<u><i>MCCI</i></u> (0.05)
Sample 1				
Data	0.284	0.011		
Nash Equilibrium	0.174	0.022	0.135	0.220
Sample 2				
Data	0.285	0.014		
Nash Bargaining	0.243	0.022	0.196	0.272

References

- [1] Aiyagari, S. R., J. Greenwood, and N. Guner (2000) “The State of the Union” *Journal of Political Economy* 109, 213-244.
- [2] Becker, G. S. (1991) *A Treatise on the Family*. Cambridge, MA: Harvard University Press.
- [3] Bound, J., C. Brown, G. Duncan, and W. Rodgers (1994). “Evidence on the Validity of Cross-Sectional and Longitudinal Labor Market Data.” *Journal of Labor Economics* 12, 345-368.
- [4] Browning, M. and P-A Chiappori. (1998) “Efficient Household Allocation: A Characterization and Tests. *Econometrica* 66, 1241-1278.
- [5] Browning, M., F. Bourguignon, P-A Chiappori, and V. Lachene (1994). “Incomes and Outcomes: A Structural Model of Intra-Household Allocation.” *Journal of Political Economy* 102, 1067-1096.
- [6] Browning, M, P-A Chiappori, and Y. Weiss (2003) “A Simple Matching Model of the Marriage Market.” Manuscript, University of Chicago.
- [7] Chen, Z. and F. Woolley. (2001) “A Cournot-Nash Model of Family Decision-Making.” *Economic Journal* 111, 722-748.
- [8] Chiappori, P-A. (1988) “Rational Household Labor Supply.” *Econometrica* 56, 63-90.
- [9] Chiappori, P-A. (1992) “Collective Labor Supply and Welfare.” *Journal of Political Economy* 100, 437-467.
- [10] Choo, E. and A. Siow (2006) “Who Marries Whom and Why.” *Journal of Political Economy* 114, 175-201.
- [11] Dagsvik, J., A. Flaatten, and H. Brunborg (2001). “A Behavioral Two-Sex Marriage Model.” *Mathematical Population Studies* 9, 97-121.
- [12] Del Boca, D. and C. Flinn (1994). “The Determination of Welfare in Nonintact Families.” Chapter 5 in *The Measurement of Household Welfare* (Blundell, Preston, and Walker, eds.). Cambridge: Cambridge University Press.
- [13] Del Boca, D. and C. Flinn (2006) “Modes of Spousal Interaction and the Labor Market Environment.” Manuscript, NYU.
- [14] Dey, M. and C. Flinn (2005). “Household Search and Health Insurance Coverage.” Manuscript, NYU.

- [15] Fernandez, R., N. Guner, and J. Knowles (2005). "Love and Money: A Theoretical and Empirical Analysis of Household Sorting and Inequality." *Quarterly Journal of Economics* 120, 273-344.
- [16] Fox, J. (2006) "Estimating Matching Games with Transfers." Manuscript, University of Chicago.
- [17] Gale, D. and L. Shapley. "College Admissions and the Stability of Marriage." *The American Mathematical Monthly*, 69 (1962): 9-15.
- [18] Greenwood, J., N. Guner, and J. Knowles (2003). "More on Marriage, Fertility, and the Distribution of Income." *International Economic Review* 44, 827-862.
- [19] Iyigun, M. and R. Walsh (2004). "Building the Family Nest: A Collective Household Model with Competing Pre-Marital Investments and Spousal Matching." Manuscript, University of Colorado.
- [20] Manser, M. and M. Brown (1980). "Marriage and Household Decision-Making: A Bargaining Analysis." *International Economic Review* 21, 31-44.
- [21] McElroy, M. (1990). "The Empirical Content of Nash-bargained Household Behavior." *Journal of Human Resources* 25, 559-583.
- [22] McElroy, M. and M. Horney (1981). "Nash-bargained Decisions: Towards a Generalization of the Theory of Demand." *International Economic Review* 22, 333-349.
- [23] McFadden, D. (1978). "Modelling the Choice of Residential Location," in *Spatial Interaction Theory and Planning Models* (A. Karlqvist, L. Lundqvist, F. Snikars, and J. Weibull, eds.). Amsterdam: North Holland.
- [24] Pollack, R. (1990) "Two-Sex Demographic Models" *Journal of Political Economy* 98, 399-420.
- [25] Roth, A. and M. Sotomayor (1990). *Two-sided Matching; A Study in Game-theoretic Modeling and Analysis*. Cambridge: Cambridge University Press.
- [26] Seitz, S. (1999) "Employment and the Sex Ratio in a Two-Sided Model of Marriage." Manuscript, Queens University.