

Labor Supply, Wealth Dynamics, and Marriage Decisions

Maurizio Mazzocco Shintaro Yamaguchi*
UCLA McMaster University

Current Draft, September 2006
Preliminary and Incomplete

Abstract

Evidence collected using the Panel Study of Income Dynamics (PSID) indicates that labor supply, saving, and marital decisions are strongly linked. This paper has two main goals. The first is to develop a model of household behavior that captures the empirical features of labor supply, saving, and marital choices. The second goal is to simulate the model using the PSID. The results indicate that the proposed model can match most of the features displayed by the data. They also suggest that the relationships between labor supply, saving, and marital status decision are important features of household behavior that should be considered by economists and policy makers when designing and implementing policies formulated to change the welfare of household members.

*We are deeply indebted to Lucas Davis for initial collaboration on this project. We are also very grateful to John Kennan, James Walker, and participants at SITE 2005, at the Econometric Society World Congress 2005, and at Queen's University Quantitative Economics Workshop for their helpful comments. Mazzocco thanks the National Science Foundation for Grant SES-0519276. Yamaguchi acknowledges the use of Shared Hierarchical Academic Research Computing Network (SHARCNET:www.sharcnet.ca) computational facilities.

1 Introduction

Studies of saving and labor supply behavior reveal large differences by marital status. Perhaps most striking are the differences in wealth. Young couples in the Panel Study of Income Dynamics (PSID) 1968-1996 have over four times as much wealth on average than unmarried individuals. There are also large differences in labor supply. Married men work an average of 200 more hours per year than single men whereas married women work an average of 200 hours less annually than single women.

Examining transitions in and out of marriage reveals that these differences are the result of gradual transitions that begin many years before changes in marital status. Saving tends to increase prior to marriage and decrease prior to divorce. Labor supply also goes through a transition. For example, prior to marriage women tend to work longer hours than the average single woman and then decrease labor supply gradually until four years after marriage they reach the level of labor supply of the average married woman. Men follow a similar pattern prior to marriage, but then continue to gradually increase labor supply during marriage. Prior to divorce, women tend to increase labor supply whereas men tend to not change labor supply.

This paper examines these transitions using a dynamic model of household behavior. The goal is to provide a framework that can describe the patterns of marital status, labor supply and wealth accumulation observed in micro data. In the model individuals meet one another and decide whether or not to get married. Couples cooperate when making decisions but are unable to commit to future allocations of resources. Divorce occurs when no reallocation of resources within the household can make both individuals better off married than single.

The inability of spouses to commit implies that the relative decision power of each household member varies with changes in the individual outside options. This has two main implications. First, saving and labor supply behavior change over time not only because of variation in interest rates and wages, but also because of changes in the relative decision power of each individual. Second, saving and labor supply decisions affect the individual outside options and therefore their individual decision power.

In this environment households have an additional saving motive that interacts with the life-cycle and precautionary motives for saving. This additional motive for wealth accumulation helps explain some of the observed pattern of saving during marital transitions. Prior to marriage, individuals increase saving in order to become attractive to potential spouses and to secure bargaining power. Assortative matching makes wealthy individuals more likely to get married. Then, married couples increase saving during marriage because of traditional saving motives. These motives are generally mitigated in marriages facing high divorce probabilities. In these households the

spouse with high decision power has additional incentives to consume because at divorce existing laws tend to distribute resources equally or equitably, whereas consumption during marriage is allocated according to the individual decision power.

Labor supply takes on added importance as well. Working increases decision power within the marriage because it increases human capital, and thus the expected future utility of divorce. For unmarried individuals, human capital increases current and future earnings, making them more attractive to potential spouses and thus increasing their decision power in case of marriage.

The observed pattern of labor supply during marital transitions reflects these dynamics. Individuals increase labor supply prior to marriage in order to become attractive to potential spouses and to secure decision power within the marriage. Then, during marriage, spouses choose individual labor supply according to individual wages and decision power. Finally, when a couple faces high divorce probabilities individuals respond by increasing their labor supply in order to prepare for the possibility of becoming single again.

The model is estimated using data from the 1984-1996 waves of the PSID and the 1984, 1989, and 1994 Wealth Supplement Files. The sample is restricted to include the cohort of individuals that are between the ages 22 and 32 in 1984. This restriction is imposed to reduce the heterogeneity of the sample and because most of the changes in marital status occur early in life. The preliminary results indicate that the proposed model can match many of the features observed in the data. We can explain the difference in labor supply and saving behavior between married and single agents, and between men and women. More important, the simulation results replicate most of the variations in labor supply and saving decisions before and after a change in marital status.

Understanding the dynamics in labor supply behavior, saving, and marital status is important for public policy. A broad and extensive literature considers the incentives for labor supply and saving created by social insurance programs and income tax policies.¹ In this literature, the relationships between marriage decisions, labor supply, and wealth accumulation are ignored. In this paper we show that these relationships are empirically significant. If they are ignored, economists and policy makers will draw the wrong conclusion from policy analysis.

This paper is related to the literature on the collective representation of household behavior. Manser and Brown (1980) and McElroy and Horney (1981) are the first papers to characterize the household as a group of agents making joint decisions. In those papers the household decision process is modeled by employing a Nash bargaining solution. Chiappori (1988; 1992) extends their model to allow for any type of efficient decision process. The theoretical model used in

¹Some examples are asset, earnings-based, and family structure-based eligibility rules for social insurance programs, spousal benefit rules for retirement programs, the structure of the income tax schedule and the marriage penalty, and tax incentives for individual and household saving.

the present paper is a generalization of the static collective model introduced by Chiappori to an intertemporal framework without commitment. The static collective model has been extensively tested and estimated. Thomas (1990) is one of the first papers to test the static unitary model against the static collective model. Browning et al. (1994) perform a similar test and estimate the intra-household allocation of resources. Chiappori et al. (2002) analyze theoretically and empirically the impact of the marriage market and divorce legislations on household labor supply using a static collective model. Blundell et al. (2001) develop and estimate a static collective labor supply framework which allows for censoring and nonparticipation in employment. Donni (2004) shows that different aspects of a static collective model can be identified and estimated.

This paper also contributes to a growing literature which attempts to model and estimate the intertemporal aspects of household decisions using a collective formulation. Lundberg et al. (2003) use a collective model with no commitment to explain the consumption-retirement puzzle. Guner and Knowles (2004) simulate a model in which marital formation affects the distribution of wealth in the population. Van der Klaaw and Wolpin (2004) formulate and estimate an efficiency model of retirement and saving decisions of elderly couples. Duflo and Udry (2004) study the resource allocation and insurance within households using data from Côte D'Ivoire. Mazzocco (2004a) analyzes the effect of risk sharing on household decisions employing a full-efficiency model. Mazzocco (2004b) estimates the standard household Euler equations for couples and separately for singles and rejects the standard unitary life-cycle model in favor of the collective formulation.

Many papers have analyzed labor supply decisions by gender and marital status. For instance, Heckman and MaCurdy (1980) estimate a life cycle model of labor supply decisions of married females. Jones, Manuelli, McGrattan (2003) and Olivetti (forthcoming) study the large increase in labor supply of married women in the United States in the second half of the twentieth century. Erosa, Furster, and Restuccia (2005) document gender differences in wage, employment and hours of work during the life cycle. They use a model with fertility decision and human capital accumulation to rationalize the empirical patterns. This is, however, one of the first attempts to estimate a model of labor supply decisions that considers the transitions in and out of marriage.

The paper proceeds as follows. The next section documents the patterns observed in the PSID. Sections 3 and 4 describe the dynamic model of the household and the wage and fertility process. Sections 5 and 6 explain how the model is implemented and present results.

2 Empirical Evidence

This section presents empirical evidence which indicates that labor supply, saving, and marital decisions are related. The discussion is based on data from the Panel Study of Income Dynamics

(PSID). The PSID is well suited to analyze the relationship between labor supply, saving, and marital decisions for two reasons. First, the PSID has gathered individual-level data on labor supply and marital status annually each spring since 1968. Wealth data are also collected at intervals of 5 years starting from 1984. Second, since the PSID is a true panel that follows the same households and their split-offs over time, the dataset can be used to examine the dynamics of household decisions.

Table 1 summarizes marriage, labor supply, saving, and consumption decisions made by unmarried females, married females, unmarried males, and married males. The sample covers the period 1968-1996 and is restricted to include individuals between the ages of 20 and 40. The latino and immigrant samples are excluded as well as individuals who did not complete high school. These restrictions are used to reduce the heterogeneity in the sample and because most of the changes in marital status occur at young ages. The PSID does not include a measure of total consumption. However, it does include measures of food expenditure and other predictors of total consumption. Therefore, following Skinner (1987), linear predictions of total consumption are used as a proxy. The coefficients used to create the fitted values are derived using the Consumer Expenditure Survey (CEX), 1980-1996, using regressors that are common to both datasets. Dollar values are reported in 1984 dollars.

Table 1 displays some features of household behavior that are worth discussing. First, labor supply varies with the marital status and gender of the individual. Conditional on working, unmarried females supply on average almost 200 more hours a year than married females. The annual labor supply of unmarried men is lower than the labor supply of married men by almost 200 hours. Both unmarried and married women supply less hours on the labor market than men. Second, labor force participation of single men is only 2% lower than labor force participation of married men which is equal to 98%. Unmarried women are five percentage points less likely to work than unmarried men. As expected, married women are much less likely to work in the labor market with a participation rate of 65%.

There are also same features of the saving decisions that are worth discussing. Two measures of wealth are reported: total household wealth and total household wealth minus home equity and the value of vehicles. First, married couples save substantially more than unmarried individuals, even after taking into account that married couples are composed of two individuals. Also, unmarried men have on average considerably more in wealth than unmarried women. Table 2 describes the wealth patterns in more detail. This table describes wealth dynamics as individuals enter and exit marriage. The decisions of couples in the period before divorce are significantly different from the decisions of the average couple. In particular, married couples on average save almost \$50,000 whereas couples in the period before divorce save less than \$19,000.

Figures 1-12 describe labor supply as individuals enter and exit marriage relative to two baseline comparison groups: married individuals and single individuals. An index is used where 0 denotes the first year of a transition between marital states. The index $-t$ indicates the t -th year prior to the transition and t indicates the t -th year after the transition. These marriages and divorces occur during different years in the sample for different individuals. In some cases particular observations will not be available for particular individuals. For example, an observation for three years prior to marriage will not be available for an individual who gets married in the second year of the sample. For this reason the number of observations will vary at different points in the index. The baseline comparison groups are weighted to reflect the calendar year composition of labor supply for the transition groups.

Figure 1 describes women's labor supply before and during marriage. During the transition, average annual labor supply falls from 1750 hours to 1200 hours. This dramatic decrease begins many years prior to marriage and continues many years into marriage. The fact that the reduction begins even prior to marriage suggests that the decrease may not be due entirely to child-bearing and household production. Figures 3 and 5 describe similar patterns for labor supply conditional on working and labor force participation. The trend for men is the opposite. While women decrease labor supply during the transition to marriage, men increase labor supply by 300 hours over the baseline level of labor supply for unmarried men. Figures 4 and 6 display labor supply conditional on working and labor force participation for the same sample of men.

Figure 7 describes labor supply of women before and after divorce. Unmarried women work more than married women. Here the differences are even larger than conditional labor supply described in table 1. Three years prior to a divorce, married women tend to work about 100 hours more than married women as a whole. Labor supply increases steadily in the years prior to divorce reaching around 1450 hours annually at the time of divorce, and then increasing an additional 200 hours in the first years after divorce. Figures 9 and 11 describe similar patterns for labor supply conditional on working and labor force participation. These results for women's labor supply closely follow the results in Johnson and Skinner (1986). The evidence before divorce for men, which is presented in Figure 8, suggests that group work about 100 hours less than the average married men. After divorce labor supply declines by about 100 hours, but it never reaches the level of the average single men. Figures 10 and 12 display labor supply conditional on working and labor force participation for the same sample of men.

This section provides evidence that labor supply, wealth, and marriage decisions are related. Traditional studies of labor supply and saving have focused on married individuals or single individuals, but not on the relationship between these variables and marital decisions. The remaining sections are devoted to explaining the evidence on these relationships presented in this section.

3 The Model

In this section a model is developed with two goals. The first goal is to capture the relationships between labor supply, saving, and marital decisions that characterize the evidence presented in the previous section. The second goal is to provide a framework that can be used to evaluate policy interventions aimed at improving the welfare of individual household members.

To that end we develop a model with the following features. Household members cooperate when making decisions. The cooperative decision process depends on the bargaining position of each agent. The individual bargaining position is determined by the best outside option available to the spouses. The best outside option corresponds to the individual value of divorce, since for the majority of individuals this is the best available alternative to being married. The value of divorce is allowed to vary over time with changes in the wage distribution, accumulated human capital, number of children, divorce laws, and marriage market. Any change that improves the best outside option of one spouse will generally increase her decision power. This will modify household decisions because the significance of her preferences in the household decision process increases. One aspect of the model that will help explaining the dynamics in labor supply and wealth before and after a change in marital status is the evolution over time of the surplus generated by the marriage. All the economic variables affect the size of the marital surplus in different periods. To strengthen its dynamics we add a variables that captures the quality of the match. This variable is meant to describe heterogeneity across marriages and over time that is not captured by the economic variables.

To model all this, consider an agent i living for T periods in an environment with uncertainty. In each period and state of nature, agent i can be either single or married. In both cases, the agent enters period t and state ω with a wage $w_t^i(\omega)$, a given amount of human capital $HC_t^i(\omega)$, and a given number of children $n_t(\omega)$. Agent i is also endowed with a given amount of saving, which is individual specific and denoted by $b_t^i(\omega)$ if single and joint with the spouse and denoted by $b_t(\omega)$ if married. If married, agent i is characterized by a given level of match quality, θ . It is assumed that the agent's preferences can be represented using a utility function u^i which depends on private consumption $c_t^i(\omega)$, leisure $l_t^i(\omega)$, and match quality. The utility function depends also on a good that is produced within the household, which is public and denoted by $Q(\omega)$ in a married family and private and denote by $q(\omega)$ in a single family. Two inputs are used in its production: time devoted to household production by the adult members of the households d_t^i , and the number of children n_t .

If agent i is married to agent j , it is assumed that decisions are efficient, except that the two household members cannot commit to future allocations of resources. In each period the two

spouses choose whether to stay married, individual private consumption, labor supply, the time devoted to household production, and the amount to be saved. It is assumed that the couple can only save jointly using a risk-free asset.² If the two spouses decide to divorce, $x\%$ of household saving is allocated to the wife and the rest to the husband, where x is exogenously determined by the divorce law.

If single in period t , agent i meets a potential spouse with probability 1. The two agents then decide whether to marry. If they choose to marry, their decisions correspond to the decisions of a couple entering period t as married. If they decide to stay single, agent i must choose individual consumption, labor supply, the time devoted to household production, and individual saving.

For both single and married agents, it is assumed that the amount of human capital accumulated by agent i in period t , hc_t^i , is an increasing function of labor supply, $f(h_t^i)$, and that human capital depreciates at a rate δ^i . Changes in the number of children are determined according to a well-defined fertility process.

The decisions of agent i will be analyzed using a recursive formulation. In each period t the vector of state variables S_t is composed of agent i 's wage and human capital, the wage and human capital of the spouse or potential spouse, saving, number of children, match quality, marital status, and agent i 's decision power relative to the spouse, which will be denoted by M_t . In any given period and state of nature the relative decision power is defined as the ratio of Pareto weights characterizing the household decision process.

Consider period T and state ω . The value of being single for agent i , $V_T^{0,i}(S_T)$, can be computed by solving the following standard problem:³

$$\begin{aligned} V_T^{0,i}(S_T) &= \max_{c_T^i, l_T^i, h_T^i, d_T^i} u^i(c_T^i, l_T^i, q_T^i) \\ \text{s.t. } c_T^i + P_T n_T^i &= w_T^i h_T^i + R_T b_T^i, \\ q_T^i &= q(n_T^i, d_T^i), \quad l_T^i + h_T^i + d_T^i = \mathcal{T}^i, \end{aligned}$$

where P_T is the cost of one child, \mathcal{T}^i is the amount of time available, R_T is the gross return on the

² In a model with no-commitment, it may be optimal for household members to have individual accounts to improve their outside options. Note, however, that the only accounts that may have an effect on the reservation utilities are the ones that are considered as individual property during a divorce procedure. In the United States the fraction of wealth that is considered individual property during a divorce procedure depends on the state law. There are three different property laws in the United States: common property law, community property law, and equitable property law. Common property law establishes that marital property is divided at divorce according to who has legal title to the property. Only the state of Mississippi has common property law. In the remaining 49 states, all earnings during marriage and all property acquired with those earnings are community property and at divorce are divided equally between the spouses in community property states and equitably in equitable property states, unless the spouses legally agree that certain earnings and assets are separate property. The assumption that household members can only save jointly should therefore be a good approximation of household behavior.

³ The dependence on the state of nature will be suppressed to simplify the notation.

risk-free asset, $q()$ is the household production function, and b_T^i corresponds to individual saving if agent i enters the period as single and to a fraction of the couple's saving otherwise.

Under the assumption that married agents make efficient decisions but without commitment, the intertemporal behavior of a couple in period T and state ω can be characterized using a Pareto problem with participation constraints.⁴ This problem can be solved using the following two-step approach. In the first step, optimal consumption and leisure are computed without taking into account the participation constraints and using the relative decision power with which the spouses enter period T , M_T :

$$\max_{c_T^i, l_T^i, h_T^i, d_T^i} u^1(c_T^1, l_T^1, Q_T, \theta_T) + M_T u^2(c_T^2, l_T^2, Q_T, \theta_T) \quad (1)$$

$$s.t. c_T^1 + c_T^2 + P_T n_T = w_T^1 h_T^1 + w_T^2 h_T^2 + R_T b_T, \quad (2)$$

$$Q_T = Q(n_T, d_T^1, d_T^2), \quad l_T^i + h_T^i + d_T^i = \mathcal{T}^i.$$

Let c_T^{i*} and l_T^{i*} , for $i = 1, 2$, be the solution of the couple's problem. Agent i 's value of being married at the current relative decision power M_T can then be computed as follows:

$$V_T^{1,i}(S_T) = u^i(c_T^{i*}, l_T^{i*}, Q_T, \theta_T).$$

In the second step, it is verified whether the individual participation constraints are satisfied, i.e.

$$V_T^{1,i}(S_T) \geq V_T^{0,i}(S_T) \quad \text{for } i = 1, 2.$$

Three possible cases may characterize the couple. First, the participation constraints are satisfied for both agents, which implies that the value of being married for agent i is $V_T^{1,i}(S_T)$. Second, the participation constraints are binding for both agents. In this case the marriage generates no additional surplus to be divided between the spouses and it is optimal to divorce.⁵ Third, only one agent is constrained. Without loss of generality suppose that agent 1's participation constraint is binding. Ligon, Thomas, and Worrall (2002) show that in this case the optimal allocation of resources is such that agent 1 is indifferent between being single and married. This allocation can be determined by choosing individual consumption, individual leisure, and the new relative decision

⁴The model used here builds on the approach developed in the limited-commitment literature. See for example Marcet and Marimon (1992, 1998), Kocherlakota (1996), Attanasio and Rios Rull (2000) and Ligon, Thomas, and Worrall (2002).

⁵ In this model both participation constraints may bind at the same time because of the public good and the possibility of meeting a new spouse if single.

power M'_T according to the following problem:

$$\begin{aligned} \max_{c_T^i, l_T^i, h_T^i, d_T^i, M'_T} \quad & u^1(c_T^1, l_T^1, Q_T, \theta_T) + M'_T u^2(c_T^2, l_T^2, Q_T, \theta_T) \\ \text{s.t.} \quad & c_T^1 + c_T^2 + P_T n_T = w_T^1 h_T^1 + w_T^2 h_T^2 + R_T b_T \\ & Q_T = Q(n_T, d_T^1, d_T^2), \quad l_T^i + h_T^i + d_T^i = \mathcal{T}^i, \\ & u^1(c_T^1, l_T^1, Q_T, \theta_T) = V_T^{0,1}(S_T). \end{aligned}$$

Let c_T^{1**} , l_T^{1**} , c_T^{2**} , l_T^{2**} and M'^* be the solution of this problem. Then if the participation constraint of agent 2 is also satisfied the two agents stay married and the value for agent i is

$$V_T^{1,i}(S_T) = u^i(c_T^{i**}, l_T^{i**}, Q_T, \theta_T).$$

Otherwise they divorce. To summarize, agent i 's value in period T is

$$V_T^i(S_T) = \max \left\{ V_T^{1,i}(S_T), V_T^{0,i}(S_T) \right\}.$$

Given agent i 's value in period T , the decision process in any arbitrary period t can be outlined. The value of being single in period t is

$$\begin{aligned} V_t^{0,i}(S_t) &= \max_{\{c_t^i, l_t^i, h_t^i, d_t^i, b_t^i\}} u^i(c_t^i, l_t^i, Q_t^i) + \beta^i E[V_{t+1}^i(S_{t+1}) | S_t] \\ \text{s.t.} \quad & c_t^i + P_t n_t^i + b_{t+1}^i = w_t^i l_t^i + R_t b_t^i \\ HC_{t+1}^i &= \delta HC_t^i + hc_t^i, \quad w_t^i = g(H_t^i), \quad hc_t^i = f(h_t^i), \quad \forall t, \omega \\ q_t^i &= q(n_t^i, d_t^i), \quad l_t^i + h_t^i + d_t^i = \mathcal{T}^i. \end{aligned}$$

where V_{t+1}^i is the value function of agent i in period $t+1$.

As for period T , the value of agent i in period t if married can be determined using two steps. In the first step, the couple's decisions are determined without considering the participation constraints using the current decision power M_t :

$$\begin{aligned} \max_{c_t^i, l_t^i, h_t^i, b_t^i} \quad & u^1(c_t^1, l_t^1, Q_t, \theta_t) + M_t u^2(c_t^2, l_t^2, Q_t, \theta_t) + E[\beta^1 V_{t+1}^1(S_{t+1}) + M_t \beta^2 V_{t+1}^2(S_{t+1}) | S_t] \\ & c_t^1 + c_t^2 + w_t^1 l_t^1 + w_t^2 l_t^2 + P_t Q_t + b_{t+1} = w_t^1 \mathcal{T}^1 + w_t^2 \mathcal{T}^2 + R_t b_t \\ HC_{t+1}^i &= \delta HC_t^i + hc_t^i, \quad w_t^i = g(HC_t^i), \quad hc_t^i = f(h_t^i), \quad i = 1, 2, \\ Q_t &= Q(n_t, d_t^1, d_t^2), \quad l_t^i + h_t^i + d_t^i = \mathcal{T}^i. \end{aligned} \tag{3}$$

The value of being married with relative decision power M_t can then be computed as follows:

$$V_t^{1,i}(S_t) = u^i(c_t^{i*}, l_t^{i*}, Q_t, \theta_t) + \beta^i E[V_{t+1}^i(S_{t+1}) | S_t] \quad i = 1, 2$$

where c_t^{i*} , l_t^{i*} , and b_t^{i*} is the solution of (3).

In the second step, it is verified whether the participation constraints of both agents are satisfied, i.e.

$$V_t^{1,i}(S_t) \geq V_t^{0,i}(S_t) \quad \text{for } i = 1, 2.$$

If both participation constraints are satisfied, the spouses stay married and agent i 's value of being married is $V_t^{1,i}(S_t)$. If both agents are constrained, there is no feasible renegotiation that makes both agents better off relative to being single and the household dissolves. If only agent i is constrained, the household renegotiates the allocation of resources so that the constrained agent is indifferent between being single and married, i.e. consumption, leisure, and relative decision power are the solution of the following problem:

$$\begin{aligned} \max_{c_t^i, l_t^i, d_t^i, h_t^i, b_t^i, M_t'} & u^1(c_t^1, l_t^1, Q_t, \theta_t) + M_t' u^2(c_t^2, l_t^2, Q_t, \theta_t) + E \left[\beta^1 V_{t+1}^1(S_{t+1}) + M_t' \beta^2 V_{t+1}^2(S_{t+1}) | S_t \right] \\ \text{s.t.} & c_t^1 + c_t^2 + w_t^1 l_t^1 + w_t^2 l_t^2 + P_t Q_t + b_{t+1} = w_t^1 T^1 + w_t^2 T^2 + R_t b_t \\ & H_{t+1}^i = \delta H_t^i + h_t^i, \quad w_t^i = g(H_t^i), \quad h_t^i = f(l_t^i), \quad i = 1, 2, \\ & Q_t = Q(n_t, d_t^1, d_t^2), \quad l_t^i + h_t^i + d_t^i = T^i. \\ & u^i(c_t^i, l_t^i, Q_t, \theta_t) + \beta^i E[V_{t+1}^i(S_{t+1}) | S_t] = V_t^{0,i}(S_t). \end{aligned} \quad (4)$$

Denote with c_t^{1**} , l_t^{1**} , c_t^{2**} , l_t^{2**} , b_t^{**} , and M_t^{**} the solution of (4). Then if the participation constraint of the spouse is also satisfied the two agents stay married and the value function of agent i is

$$V_t^{1,i}(S_t) = u^i(c_t^{i**}, l_t^{i**}, Q_t, \theta_t) + \beta^i E[V_{t+1}^i(S_{t+1}) | S_t] \quad i = 1, 2.$$

Otherwise they divorce. All this implies that agent i 's value in period t is

$$V_t^i(S_t) = \max \left\{ V_t^{1,i}(S_t), V_t^{0,i}(S_t) \right\} \quad \text{for } i = 1, 2.$$

To provide the intuition behind the couple's decision process, observe that a married couple enters period t with a given relative decision power M_t . The individual outside options change over time because of variations in household saving, wages, human capital, and children. This implies that at the given relative decision power M_t , the optimal allocation of resources may be such that one or both agents are better off as single. If this is the case, the couple will try to renegotiate the intra-household allocation, to avoid the dissolution of the household. If both agents are better off as single, then there is no feasible renegotiation and divorce is the only alternative. If only one agent is better off as single, the couple will renegotiate the allocation of resources by increasing the decision

power of the constrained spouse, and by increasing accordingly the amount of resources allocated to her in the current and future periods. Since the household makes efficient decisions without commitment, the optimal renegotiation must generate the smallest deviation from the allocation that is ex-ante efficient. This renegotiation corresponds to the intra-household allocation at which the constrained agent is indifferent between being single or married in period t . If at this allocation the spouse is also better off being married, the couple will remain married with a new relative decision power M'_t .

One last point should be discussed. As mentioned, any individual entering period t as single draws a spouse and then decides whether to get married. The marriage decision requires an initial value for M_t . This M_t determines the initial distribution of resources between spouses and therefore affects future decisions. The outside options at the time of household formation enable one to determine a set of relative weights that are feasible. There are at least three possible approaches that can be used to select the one of the feasible weights. First, M_t can be inferred from the allocation of resources in the first year of marriage. This method is demanding in terms of data because a different M_t must be identified and estimated separately for each married household. An alternative approach is to draw M_t randomly for each single agent. The main weakness of this approach is that it generates heterogeneity in individual behavior that is not explained by economic variables. An third solution is to assign to the potential spouse the relative decision power that corresponds to the Nash-Bargaining solution in period t . The advantage of this approach is that the initial intra-household allocation of resources satisfies the symmetry condition that characterizes Nash Bargaining. Hence we do not introduce the additional heterogeneity that is produced by the first approach. Because of this and the difficulties in identifying a household specific M_t , in the rest of the paper we will use Nash Bargaining to determine the initial value of M_t . After the first period of marriage, the bargaining position is determined using the method described above.

4 Assumptions on Preferences, Human Capital, Uncertainty, and Household Production

The estimation of the proposed model requires assumptions on preferences, human capital, household production, and the uncertainty that characterizes the environment. The next four subsections outline these restrictions.

4.1 Preferences and Human Capital

Empirical evidence presented in Attanasio and Weber (1995) and Meghir and Weber (1996) suggests that individual preferences are not separable in consumption and leisure. To consider this non-

separability in the estimation, the individual utility function is assumed to have the following form:

$$u^i(c^i, l^i, Q, \theta) = \frac{[(c^i)^{\sigma_i} (l^i)^{1-\sigma_i}]^{1-\gamma_i}}{(1-\gamma_i)} + \frac{(\alpha_i Q + \xi)^{1-\gamma_i}}{(1-\gamma_i)} + \frac{(\theta + \xi)^{1-\gamma_i}}{(1-\gamma_i)},$$

with $\gamma_i > 0$, $0 < \sigma_i < 1$, $\alpha_i > 0$, $\xi > 0$. The parameter γ_i captures the intertemporal aspects of individual preferences. In particular, $-1/\gamma_i$ is agent i 's intertemporal elasticity of substitution, which measures the willingness to substitute the composite good $\bar{C} = (c^i)^{\sigma_i} (T - h^i)^{1-\sigma_i}$ between different dates. The parameter σ_i captures the intraperiod features of individual preferences and it measures in each period the fraction of expenditure assigned to agent i which is allocated to private consumption.

The preferences for public consumption are assumed to be strongly separable from consumption and leisure. The functional form has been chosen to enable one to compare the utility provided by children with the utility provided by the composite good \bar{C} . In particular, $\alpha_i Q + \xi$ represents the amount of composite good \bar{C} required to provide the same level of utility as Q units of the public good, where ξ is a constant that measures the utility associated with a household that produces no public consumption.

The preferences for match quality have been similarly chosen. The variable θ is set to zero for singles. Couples with good match quality, $\theta > 0$, experience an increase in utility relative to singles that corresponds to the difference between the utility associated with ξ units of \bar{C} and $\theta + \xi$ units. The drop in utility for couples with bad match quality, $\theta < 0$, is computed in a similar way.

Another important component of the model is human capital. In the simulation it is assumed that human capital corresponds to labor market experience. In each period an individual gains one year of experience if she works full-time and half year if she works part-time. The depreciation rate is assumed to be zero.

4.2 Wage Process and Match Quality

In the model, agents face three sources of uncertainty: wages, match quality, fertility, and marriage market. This subsection discusses uncertainty over wages and match quality. The following subsection analyzes fertility. The marriage market uncertainty is discussed in section 5.

The estimation of the model requires a distribution of wages conditional on individual characteristics, for both workers and non-workers. The wage process from which individual wages are drawn is derived from the estimation of a standard Heckman selection model. The processes for women and men are estimated separately. The estimated coefficients provide the distribution of next year's wages conditional on individual characteristics. The mean of the distribution is the fitted value of the wage equation evaluated at a particular point of the state space, and the variance

is the estimated variance of the error term in the wage equation.

For workers, wages should be reasonably persistent. This is accomplished by allowing the conditional distribution of wages to depend on wages in the previous period interacted with labor force participation during the previous period and on experience. This allows wages for workers to be highly correlated over time. This also implies that individuals who did not work in the previous period have a wage distribution that depends on labor market experience only. In the estimation the wage variable will be discretized using a three-point grid. Thus, the dependence of the current wage on the lagged one is modeled by adding three dummies to the wage equation. The first one is equal to one if the lagged wage is higher than the 66-th percentile of the empirical wage distribution. The second one is equal to one if the lagged wage is between the 33-rd and the 66-th percentiles. The third one is equal to one if the wage is below the 33-rd percentile.

The selection equation depends on lagged wage interacted with labor force participation, lagged labor force participation, experience, the number of children in the household, marital status, a dummy equal to one if the worker owns a house, and age. The last four variables are assumed to enter only the selection equation. This restriction reflects the common intuition that children, marital status, and age affect the marginal productivity of household production and the marginal utility of leisure and thus labor force participation, but have less effect on productivity in the labor market.

The estimates of the wage process are presented in table 3. Several features are worth mentioning. The average wage of a woman that did not work in the previous period draws on average a wage of 3.9 1984-dollars, whereas a men with similar characteristics draws on average a wage of 5.2 1984-dollars. All the other variables have a similar effect on the average wage of women and men. I will therefore only describe the wage process of women. If a woman worked in the previous year with a wage lower than the first percentile, her average wage this period will be 30 cents lower than a woman that did not work. If her lagged wage was between the 33-rd and 66-th percentile, her current wage will on average be equal to 5.2 1984-dollars. If the wage was higher than the 66-th percentile, her average wage will be 7.9 dollars. An additional year of labor market experience increases wages by 1%. For workers the interpretation of this parameter is different from the standard estimate of the returns to experience because the specification includes lagged wages. In our specification, the experience coefficient for working individuals allows for increments in wages between periods that vary with experience. When the sample is used to estimate a standard earnings regression, the results are consistent with previous results in the literature.⁶ Table 4 presents these estimates.

The second form of uncertainty in the model is uncertainty over the quality of the match. Match

⁶See Heckman, Lochner and Todd (2003) and Mincer (1974).

quality captures household heterogeneity that is not explicitly modeled. The corresponding process is modeled to capture two features. First, couples learn gradually about the quality of the match. Second, match quality should be reasonably persistent. These aspects of marriage are introduced in the framework using the following statistical model. When two unmarried individuals meet for the first time, they draw a starting value for match quality and a trend for this variable. The trend can be either upward or downward. In each period after marriage, a new match quality is drawn conditional on the match quality in the previous period using a Markov transition matrix that depends on the trend characterizing the couple.⁷ The value of match quality at the time of household formation is drawn from the stationary distribution that corresponds to the transition matrices. In the current version of the model we use 4 levels of match quality. The upward and downward trend is modeled using two transition matrices with the following form:

$$\begin{bmatrix} 1-a & a & 0 & 0 \\ b & 1-a-b & a & 0 \\ 0 & b & 1-a-b & a \\ 0 & 0 & b & 1-b \end{bmatrix}.$$

where $a > b$ if there is an upward trend and $a < b$ in the case of a downward trend. In each period there a probability p that the couple switches from an upward trend to a downward trend in match quality. The choice of the shape for the transition matrices is arbitrary but it enables us to capture in a simple way the idea that spouses learn gradually about their match quality.

4.3 Fertility

Children should be considered in a realistic model of the household for at least two reasons. First, children are one of the main reasons for the existence of marriages. Second, the presence of children affects other aspects of household behavior. However, since the main goal of this research project is to develop and analyze a unified model of labor supply, wealth accumulation, and marital decisions, we do not explicitly model fertility choices. Instead, following Brien, Lillard, and Stern (forthcoming) we assume that fertility choices can be characterized using a statistical process that matches the data. There are two possible ways to construct the fertility process. First, it may be assumed that the probability of having a child depends on individual decisions about saving, labor force participation, marriage, and children. This enables one to construct a fertility process that reflects the main features of the data: married women are more likely to have children; wealthier women have fewer children; women who participate in the labor market are less likely to

⁷We use this statistical model instead of a standard learning model because it captures the main insight behind the learning model with fewer state variables.

have children. This treatment of fertility has one limitation: women will take into account that their labor supply and saving choices will affect the probability of having children. An alternative approach is to assume that fertility depends only on marital status and number of children. In this case women will try to marry to increase their chances of having a child, which is realistic, but women will not consider the fertility process when choosing saving and labor supply. We experimented with both specifications and decided to use the second one to avoid the limitation mentioned above. This approach for characterizing fertility significantly simplifies the estimation without making the unrealistic assumption that fertility has no effect on labor supply, saving, and marital choices. The generalization of the current model to an environment in which individuals are also allowed to choose when to have children is important, but it is left for future research.

The statistical process for fertility is estimated using a standard probit. The sample employed to estimate the probit includes married and unmarried women. An observation is a woman/year and the dependent variable is a dummy variable that takes on the value of one if the number of children in the household increases in the year following the current year.

Table 5 reports the estimates from the fertility probit. The most important predictor for fertility is marital status. Everything else equal, a married woman has a probit score that is higher by .753 standard deviations. Evaluated at the mean, this implies that a married woman has a 7.6 percentage point higher probability of having a child during a given year. This effect is large relative to the overall fertility rate in the sample of 10.1%. The current number of children in the household is also predictive. Relative to households without children, household with one child are more likely to give birth to a second child, whereas families with two or more children are less likely to give birth to an additional child.

Children also enter the model through the budget constraints. The cost of children is estimated separately from the central model using the following strategy. First, a measure of household expenditure is created that includes the key components of expenditure related to children. The CEX is used rather than the PSID because it provides more detailed expenditure information. The expenditure measure includes food at home, child care, boy's clothing and shoes, girl's clothing and shoes, and infant clothing. We also experimented with a broader measure of expenditure, but it yielded a smaller estimate of the cost of children because households without children spend more on other components of expenditure like food consumed in restaurants. Second, expenditure is regressed on a vector of household characteristics including the number of children. Table 6 presents the estimates of the cost of children. They suggest that the cost of children is a concave function, with the first child costing more than subsequent children.

4.4 Household Production

In the estimation of the model we make two assumptions about household productions. First, we assume that the production function is linear in the inputs, i.e.

$$Q(n_t, d_t^1, d_t^2) = \eta_0 n_t + \eta_1 d_t^1 + \eta_2 d_t^2.$$

Note that we will not be able to separately identified the preferences for public consumption from the production function. We can only identify their product $\alpha_i \eta_j$. Since in this paper we are only interested in understanding the effect of children and domestic labor on the marriage surplus, this is not a concern.

As a second assumption we impose the restriction that domestic labor is not directly chosen by the individuals, but it is a linear function of marital status, number of children, and labor market variables. This function is estimated separately from the main model using the information available in the PSID for domestic labor. During the 1984 to 1996 waves, the following question was asked to household heads and spouses: “About how much time do you spend on housework in an average week? I mean time spent cooking, cleaning, and doing other work around the house.” The answer to this question provides the amount of time that each household member dedicates to domestic production. We then employ this information with a standard OLS regression to characterize the quantity of domestic labor supply as a function of marital status, number of children, and labor market variables. The estimated coefficients, which are reported in table 7, are used in the estimation to impute domestic labor supply. This specification enables us to capture the effect of domestic labor on household behavior, without the computational burden associated with determining the corresponding choices for each point of the state space. In particular, we can take into account that families with children not only pay additional financial costs, but have also to spend additional time in domestic production. The long term goal is to allow the agents to choose the time devoted to household production, but this is left for future research.

5 Estimation Issues

This section discusses the estimation of model. The model developed in this paper is estimated using the Simulated Method of Moments (SMM) and data from the 1984-1996 waves of the PSID. In 1997 the survey was redesigned for biennial data collection. This explains why data gathered starting in 1997 are not included.⁸

⁸In 1990, 2,000 latino households were added to the sample. This latino sample was dropped after 1995 and replaced with a sample of 441 immigrant families in 1997. We exclude both the latino and immigrant samples.

The parameters of the model are estimated by a simulated minimum distance estimator, which is closely related to the simulated method of moments, the efficient method of moments, and indirect inference (see Carrasco and Florens, 2002, for a survey.) The model is simulated to generate an artificial data set of labor supply, marital status, consumption, and wealth paths. This simulated data is used to compute moments such as the divorce hazard rate, the percentage of married individuals, the average labor supply of married men. These simulated moments are then compared with the corresponding moment computed using the actual data. The objective of the estimation is to search for structural parameters that minimize a weighted sum of the distances between the simulated and data moments. In the estimation we use the inverse of the covariance matrix of the data moments as a weighting matrix. The covariance matrix is computed using a standard bootstrap method with 10000 bootstraps.

The implementation of the SMM requires the choice of a set of auxiliary moments. In the current version of the paper we use the following 31 moments: the percentage of married individuals without children, the percentage of married individuals with children; the divorce hazard for couples without children; the divorce hazard for couples with children; the average labor supply of women t periods from the time of marriage, with $t = -4, -3, -2, -1, 0, 1, 2, 3, 4$; the average labor supply of men t periods from the time of marriage, with $t = -4, -3, -2, -1, 0, 1, 2, 3, 4$; the average labor supply of women t periods from the time of divorce, with $t = -4, -3, -2, -1, 0, 1, 2, 3, 4$.

We will now describe the main aspects of the simulation that is required to compute the simulated moments. In case of divorce, in the model household wealth must be divided between the two spouses. To determine the fraction that is allocated to the wife, we use divorce settlements from the National Longitudinal Study of the High School Class of 1972 (NLS-72), Fifth Follow-up (1986). The sample ($n=1685$) includes all first marriages that ended in legal divorce prior to 1986. The average percentage of wealth allocated to the wife is 0.496 with a standard deviation of 0.177, where household wealth is the total net value of all property including the house value and the value of other real estate. In the simulation we, therefore, assume that wealth is divided equally between the two spouses.

In the simulation one has to decide which parent receives custody of the children in case of divorce. According to the United States Census, 58.1% of mothers are the custodial parent. Data collected from the US Census also indicates that only 54.9% of fathers have either joint custody or visitation rights. The standard arrangement in case of visitation rights is that the non-custodial parents can spend 2 weekend days every two weeks with their children. This corresponds to around the 15% of time for 54.9% of fathers. To be consistent with these facts, it is assumed that mothers maintain custody of their children and that children spend 85% of their time with the mother and the remaining 15% with the father. In the simulation, we implement this by multiplying the

number of children in the production function of a divorce mother and father by, respectively, 0.85 and 0.15. To simplify the simulation, it is also assumed that in case of remarriage of one of the two parents, the children spend 100% of their time with the mother and that the father does not pay any child cost. This is equivalent to assuming that the utility provided by children to divorced fathers in case of remarriage is equivalent to the utility provided by the quantity of composite good \bar{C} that divorced fathers can purchase because they do not have to pay the cost of children.

The state space has been chosen to reflect the features of the sample that will be used in the simulation. We start with a description of the grid for the three sources of uncertainty. First, the continuous distribution of wages obtained using the Heckman-selection model is approximated using the method proposed in Kennan (2004). It shows that the best discrete approximation \hat{F} to a given distribution F using the fixed support points $\{w_i\}_{i=1}^n$ is given by

$$\hat{F}(w_i) = \frac{F(w_i) + F(w_{i+1})}{2},$$

where n denotes the number of support points and i indexes each point. In the simulation we set n equal to 3 and we use as grid points the wages that corresponds to the 1/6th, 3/6th, and 5/6th points in the empirical distribution. The corresponding grid for women is \$2.93, \$4.70, and \$7.91. The grid for men is \$3.96, \$6.60, and \$10.72. Second, the grid for children is composed of 3 points, starting with zero children and ending with 2 children. The number of children varies according to the fertility process illustrated in a previous section. Third, it is assumed that match quality is drawn from a eight-point distribution, which implies that the evolution of this feature of household behavior is characterized by a eight by eight Markov Transition Matrix.

The decision variables labor supply, saving, and relative decision power, M are also discretized. The grid for labor supply is chosen to reflect no labor force participation, part-time, and full-time work. In particular, a three-point distribution is used with values 0, 1400, 2400. Wealth is described using an equally-spaced 13-point grid varying between -\$4,000 and \$44,000 for singles and between -\$8,000 and \$88,000 for married individuals. This range is chosen to reflect the distribution of total wealth minus home equity and the value of vehicles in the PSID, where about 5.2% of households report a wealth level below our range and 10.7% above. We have experimented with using a larger range for wealth, particularly on the high end and we have found that few households choose to accumulate such high levels of wealth. The grid for relative decision power includes 21 points .01, .05, .10, ..., .90, .95, and .99. We have tested the robustness of the simulation with respect to changing the number of grid points for the Pareto weight and we have found that it is important to include a reasonably fine grid. With a grid that is too coarse, there may be mutually beneficial marriages that do not occur because the grid does not contain any points within the range of Pareto weights for which the marriage is sustainable.

The grid for experience requires a separate discussion. This grid can be constructed in two different ways. A first possibility is to choose a number of points in the grid that is equal to the number of periods in the simulation. Then if an individual works in a given year, her experience increases by one. Since this approach is computationally demanding, we adopt a different strategy that generates similar results. The grid for experience is described using a two-point distribution, 0 and 25 years. Experience then increases according to the following law of motion. If an agent works over 500 hours in a given period, she has a 1 in 25 chance of increasing to the high experience state. Thus, the expected increase in experience for such an individual is equivalent to the expected increase in the first approach. This mechanism allows us to capture the effect of working on human capital accumulation without using such a fine grid that the simulation becomes excessively time consuming. The robustness of these results to richer specifications of experience has been tested and the results do not change substantially.

The model is simulated for 45 periods. In the first 25 periods individuals make decisions about marriage, labor supply, consumption, and saving. The remaining 20 years represents the retirement period and individuals can only choose consumption and saving. The rate of return on saving is allowed to change over time. In particular, for 1982-2004 the 20-year municipal bond rate is used as the rate of return on saving.⁹ For 2005-2009 the interest rate is assumed to remain unchanged at the 2004 level.

We solve the problem using backward induction. Consider an arbitrary period. Each agent enters the period as either single or married. If the agent is single, she draws a potential spouse from the distribution of available spouses. For each agent, we evaluate first the level of utility associated with being single. Afterwards, we compute the level of utility associated with being married to the current spouse if already married or to the potential spouse if single at the given relative decision power, from this point in time forward. The level of utility conditional on marital status is computed by checking all possible alternatives for consumption, labor supply, and saving, and selecting the choice that yields the highest level of utility. At the current relative decision power, each married or potential couple can be in one of the following three regimes: (i) both agents prefer being single; (ii) both agents prefer being married; (iii) one agent prefers being single and the other is better off as married. In the first two cases the marital choice is straightforward. In the third case, the couple renegotiates the current allocation of resource to make the spouse that prefers being single just indifferent between the outside option and staying in the household. This goal is achieved by shifting relative decision power, and accordingly the allocation, until this indifference condition is satisfied. If at new allocation both agents prefer being married, the couple stay married or get married. Otherwise the marriage does not generate any additional surplus and

⁹The rates are obtained from Bloomberg.

the agents will divorce or stay single.

It is worth discussing in more detail the mechanism by which potential spouses are drawn. Individuals are characterized by experience, wage, wealth, number of children, relative decision power, and match quality. For potential spouses experience and wage are drawn from the empirical distribution obtained from the PSID data and then discretized. Wealth is drawn using a similar approach, but each individual can only draw a potential spouse with a wealth level that is one point below, one point above, or at the same point in the wealth grid. This restriction is imposed to capture the fact that people search for their spouses in similar circles. With regard to children, it is assumed that single men draw only women with no children. We make this assumption for two reasons. First, men in the age range considered in our simulation marry mostly women with no children. Second, in our model men derive utility from children independently of whether the children were conceived during or before the marriage. This implies that, if we do not make this assumption, single men search for single women with children to increase their utility after marriage, which is unrealistic. As discussed in the theory section, the initial relative decision power, M , is the one that corresponds to the Nash Bargaining solution to avoid generating heterogeneity that is not explained by economic variables. Finally match quality is drawn from the stationary distribution derived from the transition matrix.

The solution of the model is characterized by policy functions. For every state of the world the policy functions return the optimal choices for marital status, consumption, labor supply and saving. In addition, for couples the policy function includes relative decision power. The policy functions are used to simulate the model for the group of individuals available in the 1984 wave of the PSID that satisfy the selection restrictions mentioned above. For each individual in the 1984 wave we match her wage, experience, wealth, marital status, and number of children to the point on the grid that most closely approximates each of the characteristics observed in the data. For married individuals there are two state variables that are not observed: relative decision power and match quality. It is assumed that individuals that are married in the 1984 wave have the relative decision power that corresponds to the Nash Bargaining solution. After the first period of the simulation, relative decision power is optimally computed. For married individuals in 1984, it is also assumed that their marriage is characterized by the second highest level of match quality. Since the couple is still married this is the best prediction we can make. The model is then simulated for 1984-2029. Because we want these simulated individuals to follow a rich set of paths, we simulate multiple possible paths for each individual. This is done by sampling with replacement from the 1984 sample in the PSID 50,000 times.

It is worth describing how individuals are followed in the simulation. In the simulation there are three types of agents. First, individuals that are in the 1984 wave and are tracked by the PSID.

These are individuals in households that belong to the original 1968 PSID sample and that are tracked by the PSID even if there is a split-off. There are two main reasons for a split-off: a child in one of the original families forms her own household; a couple in the original sample or one of their children divorces. We will refer to these individuals as tracked agents. Given the age range used in this paper, all our tracked agents are children of families in the original sample. Second, individuals that are in the 1984 wave, but are not tracked by the PSID. If there is a split-off, these agents vanish from the PSID. Third, agents that appear in the simulation as potential or actual spouses of the tracked agents. To illustrate the different treatment of these types of individuals in the simulation, consider the case of a track agent who is single in the initial period. In this period this individual meets a potential spouse. If the couple decides not to marry, then the potential spouse is dropped from the simulation. If the couple marries, then the individual will remain in the sample until the couple gets a divorce. After a divorce this individual is no longer considered in the simulation. For individuals who are initially married, only one of the two individuals is the tracked agent, the other individual is dropped from the simulation if the household ever decides to divorce.

In the version of the model that we are currently estimating there are 14 parameters: the risk aversion parameter γ , the parameter σ , and the discount factor β for women and men; the three parameters of the household production function; the normalization parameter ξ in the utility function; the parameters a and b that determines the upward and downward trends in match quality; a parameter that measures the probability of switching from a positive to a negative trend; a parameter that determine how disperse the match quality values are. We estimate six of them. The risk aversion parameters are set equal to 2.5 for men and to 4.5 for women using the results in Mazzocco (2004c), in which the risk aversion parameter is estimated separately for men and women. The discount factor is assumed to be constant and equal to .98. We will experiment with alternative values for β . The parameter ξ is set equal to the average consumption of the composite good for singles using their average budget share as σ . The parameters a and b are set equal to 0.05 and 0.0375. The probability of switching is set equal to 0.05.

The next section presents the simulated results. They are used to replicate the tables and figures reported in section 2. For consistency with the PSID data, the simulated tables and figures are constructed using only tracked or married individuals.

6 Results

The first part of table 8 reports preliminary estimates of coefficients. The second part presents the corresponding simulated results and compares them to the data. The model does a reasonable

job of matching the pattern of labor supply observed in the data. First, married men work more than unmarried men whereas unmarried women work more than married women. This behavior is explained by four aspects of the model: differences in the wage process between women and men; differences in domestic labor between women and men and between married and single individuals; the presence of children in the household; human capital accumulation. Women work less than men because they have on average lower wage offers. This feature of the wage process also explains why married men work more than single men and single women supply more market hours than married women. Since the two spouses cooperate, it is optimal for the husband to specialize in market production and for the wife in domestic production. This pattern is strengthened by the requirement that married women have to supply more hours to household production than single women, whereas married men have to supply less hours. Children increase the degree of specialization within the household for two reasons. First, both parents devote more time to household production, but the increase for mothers is much larger. Everything else equal, this implies that it is more costly for the wife to supply labor on the market. Second, since in our model mothers receive custody of their children in case of divorce, the presence of children transfers bargaining power from the husband to the wife. This leads married men to increase and married women to decrease labor supply. Human capital accumulation has a similar effect on labor supply as children. Intra-household specialization reduces the wife's labor supply and increases the husband's. Thus, married women accumulate less human capital relative to married men. Therefore the couple's efficient allocation of resources requires an additional increase in the husband's labor supply and an additional reduction in the wife's labor supply.

We are not doing as well in capturing the main features of labor force participation in the data. Our simulated results indicates that married women participate less than single women and that the fraction of single women with a positive amount of labor hours is lower than the fraction of single men. But in our simulation, single men participate slightly more than married men which is contrary to what we observe in the data. The introduction of labor force participation moments to be matched should solve this problem.

The simulation replicates only some of the features observed in the PSID for wealth and consumption. Table 8 shows that saving of single men is greater than saving of single women, but saving of couples is only slightly larger twice saving of unmarried men. The simulated data for consumption show a similar pattern. Consumption of single men is greater than consumption of single women and consumption of couples is less than twice consumption of single men. However, we overpredict consumption for each group. This high level of consumption and low level of saving is probably explained by the lack of durable goods in the model and by the fact that we are not matching any consumption and wealth moments.

Figure 13 describes the labor supply pattern for women as they enter marriage. As before, an index is used where 0 denotes the first year of a transition between marital states. In the PSID, women's labor supply decreases starting three years prior to marriage and continues to gradually decrease during the first five years of marriage. Moreover, at the time of marriage women's labor supply is between the average labor supply of married women and the average labor supply of single women. The simulated data display a similar pattern. The average labor supply of women at the time of marriage is above the average labor supply of married women and below the average labor supply of single women. Then in the first few years after marriage, average labor supply decreases gradually toward the labor supply of the average married woman. A notable difference between the PSID data and the simulation is that there is a sharp decline of labor supply before marriage in the simulation that we do not observe in the data. The gradual decline of labor supply for married women after marriage is explained by three features of our model: the presence of children; the differential accumulation of human capital between wife and husband; the absence of commitment. During these years married women tend to work less and less as they increase the time devoted to domestic labor and increase their bargaining power in the marriage. This causes married women to accumulate less human capital relative to married men. For the household it is therefore optimal to further increase specialization in the household, reducing the wife's labor supply and increasing the husband's labor supply even more. The lack of commitment produces a more gradual decline in labor supply of wives. To see this note that married women know that with some probability the marriage will end, in which case the accumulated human capital will be important. Conditional on children, the probability of divorce depend on the level of match quality. In our model the marriages that survive are the ones with an upward trend in match quality. Or equivalently, marriages where the spouses learn over time that the quality of their match is high. In this marriages the spouses increase gradually the degree of specialization, which explain the gradual decline in labor supply of women. Figures 14 and 15 show the simulation results for labor supply conditional on working and labor force participation.

Figure 16 describes the labor supply behavior for men as they enter marriage and may be compared to corresponding figures that describes the pattern for men in the PSID. In both the simulation and the data, the overall trend in labor supply is the opposite of the trend for women. In the data, men tend to increase labor supply two years prior to marriage and then continue to increase their labor supply throughout marriage. In the simulation, men start increasing their labor supply throughout at the time of marriage. As for women, these changes in labor supply are a result of children, human capital accumulation, and lack of commitment.

Figure 17 describes women's labor supply before and after divorce. Both this figure and the analogous figure from the PSID describe a steady increase in labor supply beginning several years

prior to divorce. In the model this behavior is explained by two types of married couples: couples with a downward trend in match quality; couples in which the wife experiences a high wage offer and therefore increases her labor supply. To provide the intuition underlying this result, consider first a couple with a downward trend in match quality. In this marriages, the probability of divorce increases over time. It is therefore optimal for the spouses to reduce the degree of specialization within the household, which explains the gradual increase in the wife's labor supply. Consider now a couple in which the wife receives a good wage offer. In this case, she will increase her labor supply and reduce the time spend on domestic labor. As a consequence, the marriage produced a lower amount of public good, which increases the probability of divorce. As a consequence, it is optimal for the pair to reduce the degree of specialization, which increases even more the probability of divorce. Currently we cannot determine which one of the two types of marriages explains the pattern in the data. When we will have the final estimates we will attempt to separate the effect of the two types of couples. Figures 18 and 19 show the simulation results for labor supply conditional on working and labor force participation.

Figure 20 describes the labor supply pattern for men before and after divorce. This figures may be compared to Figure 8 that describes the same pattern in the PSID. In the PSID, men work less before divorce than the average married men. In the simulation labor supply displays a similar pattern. After divorce in the PSID men decrease gradually their labor supply, whereas in the simulation there is sudden drop after divorce and an increase afterward. The labor supply behavior of men before divorce can be rationalized using the results for women before divorce. It is optimal for couples with a high probability of divorce to reduce the degree of specialization.

When we will have the final estimates we will attempt to determine how important different features of the model are in explaining the empirical patterns.

7 Conclusions

This paper proposes a unified model of labor supply, saving, and marital decisions. The model is estimated using the 1984-1996 waves of the PSID. The results indicates that the model captures some of the features of labor supply, wealth, and marriage decision displayed by the PSID data.

References

- [1] Attanasio, Orazio P. and José Victor Rios Rull: “Consumption Smoothing and Extended Families,” Manuscript, (July 2000).
- [2] Attanasio, Orazio P. and Guglielmo Weber: “Is Consumption Growth Consistent with Intertemporal Optimization? Evidence from the Consumer Expenditure Survey,” *Journal of Political Economy* 103, no. 6 (1995): 1121-1157.
- [3] Blundell, R., P. A. Chiappori, T. Magnac, and C. Meghir: “Collective Labor Supply: Heterogeneity and Nonparticipation,” Institute for Fiscal Studies Working Paper (2001).
- [4] Brien, Michael J., Lee A. Lillard, and Steve Stern: “Cohabitation, Marriage, and Divorce in a Model of Match Quality,” Manuscript, Department of Economics, University of Virginia (2004).
- [5] Browning, Martin, Francois Bourguignon, Pierre-Andre Chiappori, Valerie Lechene: “Income and Outcomes: A Structural Model of Intrahousehold Allocation,” *Journal of Political Economy*, Vol. 102, No. 6. (Dec., 1994), pp. 1067-1096.
- [6] Chiappori, Pierre-Andre’: “Rational Household Labor Supply,” *Econometrica*, Vol. 56, No. 1. (Jan., 1988), pp. 63-90.
- [7] Chiappori, Pierre-Andre’: “Collective Labor Supply and Welfare,” *Journal of Political Economy*, Vol. 100, No. 3, (1992), pp. 437-467.
- [8] Chiappori, Pierre-Andre’, Bernard Fortin and Guy Lacroix: “Household Labor Supply, Sharing Rule and the Marriage Market,” *Journal of Political Economy*, Vol. 100, No. 3, (2002), pp. 437-467.
- [9] Donni, Olivier: “A Collective Model of Household Behavior with Private and Public Goods: Theory and Some Evidence from U.S. Data,” DELTA Working Papers, (2004).
- [10] Duflo, Esther and Christopher Udry: “Intrahousehold Resource Allocation in Côte D’Ivoire: Social Norms, Separate Accounts and Consumption Choices,” NBER Working Paper, (May, 2004).
- [11] Erosa Andrés, Luisa Fuster, and Diego Restuccia: “A Quantitative Theory of the Gender Gap in Wages,” Manuscript, September 2005.

- [12] Guner, Nezih and John Knowles: "Marital Instability and the Distribution of Wealth," Working Paper, Department of Economics, University of Pennsylvania (2003).
- [13] Heckman, James J., Lance J. Lochner, and Petra Todd: "Fifty Years of Mincer Earnings Regressions," NBER Working Paper Number 9732, (2003).
- [14] Heckman, J. James and Thomas E. MaCurdy: "A Life Cycle Model of Female Labour Supply," *The Review of Economic Studies*, Vol. 47, No. 1, Econometrics Issue. (Jan., 1980), pp. 47-74.
- [15] Johnson, William R. and Jonathan Skinner: "Labor Supply and Marital Separation," *The American Economic Review*, Vol. 76, No. 3. (Jun., 1986), pp. 455-469.
- [16] Kennan, John: "A Note on Approximating Distribution Functions," Working Paper, Department of Economics, University of Wisconsin (2004).
- [17] Kocherlakota, Narayana R.: "Implications of Efficient Risk Sharing without Commitment," *Review of Economic Studies*, 63, no. 4 (October 1996): 595-609.
- [18] Ligon Ethan, Jonathan P. Thomas and Tim Worrall: "Informal Insurance Arrangements with Limited Commitment: Theory and Evidence from Village Economies," *Review of Economic Studies*, Vol. 69, No. (2002), pp. 209-244.
- [19] Lundberg, Shelly, Richard Startz, and Steven Stillman: "The Retirement-Consumption Puzzle: a Marital Bargaining Approach," *Journal of Public Economics*, (May 2003), pp. 1199-1218.
- [20] Manser, Marilyn and Murray Brown: "Marriage and Household Decision-Making: A Bargaining Analysis" *International Economic Review*, Vol. 21, No. 1. (Feb., 1980), pp. 31-44.
- [21] Jones, Larry E., Rodolfo E. Manuelli, and Ellen R. McGrattan: "Why are Married Women Working so Much?" Staff Report 317, Federal Reserve Bank of Minneapolis.
- [22] Marcet, A. and R. Marimon: "Communication, Commitment and Growth," *Journal of Economic Theory*, 58, no. 2 (1992): 219-249.
- [23] Marcet, A. and R. Marimon: "Recursive Contracts," Manuscript, Universitat Pompeu Fabra (1998).
- [24] Mazzocco, Maurizio: "Savings, Risk Sharing and Preferences for Risk," *American Economic Review*, Vol. 94, No. 4, (Sep. 2004a), pp. 1169-1182.
- [25] Mazzocco, Maurizio: "Intertemporal Behavior and Household Structure," Manuscript, Department of Economics, University of Wisconsin (2004b).

- [26] Mazzocco, Maurizio: “Individual Rather Than Household Euler Equations: Identification and Estimation of Individual Preferences Using Household Data,” Manuscript, Department of Economics, University of Wisconsin (2004c).
- [27] Mazzocco, Maurizio and Pierre-André Chiappori: “Labor Supply Decisions and Commitment Within The Household: The Effect of The Irish Divorce Law,” Working Paper, Department of Economics, University of Wisconsin (2004).
- [28] McElroy, Marjorie B. and Mary Jean Horney: “Nash-Bargained Household Decisions: Toward a Generalization of the Theory of Demand,” *International Economic Review*, Vol. 22, No. 2. (Jun., 1981), pp. 333-349.
- [29] Meghir Costas and Guglielmo Weber: “Intertemporal Nonseparability or Borrowing Restrictions? A Disaggregate Analysis Using a U.S. Consumption Panel,” *Econometrica*, 64, no. 5 (September 1996): 1151-1181.
- [30] Mincer, Jacob: *Schooling, Experience, and Earnings*, Columbia University Press: New York, (2004).
- [31] Olivetti, Claudia: “Changes in Women’s Aggregate Hours of Work: The Role of Returns to Experience,” *Review of Economic Dynamics*, forthcoming.
- [32] Skinner, Jonathan: “A Superior Measure of Consumption from the Panel Study of Income Dynamics,” *Economics Letters*, Vol. 23, No. 2, (1987), pp. 213-216.
- [33] Tauchen, George: “Finite State Markov-chain Approximations to Univariate and Vector Autoregressions,” *Economics Letters*, Vol. 20, No. 2, (1986), pp. 177-181
- [34] Thomas, Duncan: “Intra-Household Resource Allocation: An Inferential Approach,” *Journal of Human Resources*, Vol. 25, no. 4. (Autumn 1990), pp. 635-664.
- [35] van der Klaauw, Wilbert and Kenneth Wolpin: ”Social Security, Pensions and the Savings and Retirement Behavior of Households,” Working Paper, Working Paper, University of North Carolina (2004).
- [36] Weiss, Yoram and Robert J. Willis: “Transfers among Divorced Couples: Evidence and Interpretation,” *Journal of Labor Economics*, Vol. 11, No. 4. (Oct., 1993), pp. 629-679.

A. Figures

Figure 1: PSID 1968-1996, Women's Labor Supply Before and During Marriage.




Figure 2: PSID 1968-1996, Men's Labor Supply Before and During Marriage.




Figure 3: PSID 1968-1996, Women's Labor Supply Conditional on Working Before and During Marriage.




Figure 4: PSID 1968-1996, Men's Labor Supply Conditional on Working Before and During Marriage.




Figure 5: PSID 1968-1996, Women's LFP Before and During Marriage.




Figure 6: PSID 1968-1996, Men's LFP Before and During Marriage.




Figure 7: PSID 1968-1996, Women’s Labor Supply Before and After Divorce.




Figure 8: PSID 1968-1996, Men’s Labor Supply Before and After Divorce.




Figure 9: PSID 1968-1996, Women’s Labor Supply Conditional on Working Before and After Divorce.




Figure 10: PSID 1968-1996, Men’s Labor Supply Conditional on Working Before and After Divorce.




Figure 11: PSID 1968-1996, Women's LFP Before and After Divorce.




Figure 12: PSID 1968-1996, Men's LFP Before and After Divorce.




Figure 13: Simulation, Women's Labor Supply Before and During Marriage.




Figure 14: Simulation, Women's Labor Supply Conditional on Working Before and During Marriage.



Figure 15: Simulation, Women’s Labor Supply Conditional on Working Before and During Marriage.




Figure 16: Simulation, Men’s Labor Supply Before and During Marriage.




Figure 17: Simulation, Women's Labor Supply Before and After Divorce.




Figure 18: Simulation, Women's Labor Supply Conditional on Working Before and After Divorce.




Figure 19: Simulation, Women's LFP Before and After Divorce.





Figure 20: Simulation, Men's Labor Supply Before and After Divorce.



B. Tables

Table 1: PSID Descriptive Statistics

	Mean	SD
Marital Status and Transitions:		
percentage married	.79	(.41)
marriage hazard	.189	(.392)
divorce hazard	.030	(.169)
Labor Force Participation:		
unmarried women	.91	(.26)
unmarried men	.96	(.18)
married women	.65	(.39)
married men	.98	(.15)
Annual Hours Worked if Working:		
unmarried women	1861	(669)
unmarried men	2095	(711)
married women	1660	(736)
married men	2312	(622)
After-Tax Hourly Wage:		
unmarried women	6.21	(4.12)
unmarried men	7.14	(4.59)
married women	5.68	(4.57)
married men	8.36	(5.77)
Wealth:		
unmarried women	20,415	(43,130)
unmarried men	31,252	(91,767)
married couples	89,831	(274,182)
Wealth Excluding Home Equity and Vehicles:		
unmarried women	8,180	(28,597)
unmarried men	14,149	(78,329)
married couples	49,732	(257,622)
Consumption:		
unmarried women	6,522	(2,698)
unmarried men	7,927	(3,283)
married couples	10,847	(3,246)

Note: The PSID is a longitudinal study of a representative sample of U.S. individuals. The sample is from the 1968-1996 waves and the 1984, 1989 and 1994 Wealth Supplement Files. The Latino and Immigrant Samples have been excluded along with the 1968 low-income Census oversample. The sample is restricted to include only household heads and wives, not sons, daughters or other household members, unless they have started their own household. After these exclusions there are 29,594 total individual-year observations, or, about 2000 individual-year observations per year. Each survey wave records extensive individual-level information on labor supply and wages. Annual hours worked is constructed as the typical number of hours worked at each of up to three jobs multiplied by the number of weeks during the year worked at each job. Wages are calculated as annual individual labor market income divided by total annual hours worked. After-tax wages are adjusted for federal and state income tax using the NBER's TAXSIM using household income, year, state of residence, marital status, and number of children.

Table 2: Average Wealth Levels in the PSID.

married couples	\$49,732
unmarried individuals	\$11,161
married couples the year before divorce	\$18,471
individuals the year before marriage	\$11,499
just married couples	\$19,798
just divorced individuals	\$8,385

Note: The transitions are described over the last year. For example, “married couples just before divorce” are currently married but will not be married next year and “just married couples” are currently married but were not married last year. The wealth measure excludes wealth in housing and cars.

Table 3: Wage Process Estimates

	Coefficient	SE
Women		
high wage \times lagged labor force participation	.716	(.084)
medium wage \times lagged labor force participation	.289	(.083)
low wage \times lagged labor force participation	-.089	(.076)
experience	.010	(.0045)
experience squared	-.0004	(.0002)
constant	1.351	(.097)
variance of wage equation	.258	(.004)
log likelihood = -1380.2		
Men		
high wage \times lagged labor force participation	.623	(.043)
medium wage \times lagged labor force participation	.280	(.042)
low wage \times lagged labor force participation	-.064	(.064)
experience	.012	(.003)
experience squared	-.0003	(.0001)
constant	1.636	(.046)
std. dev. of wage equation	.256	(.002)
log likelihood = -1164.7		

Note: The dependent variable is log hourly wage deflated to reflect year 1984 prices. Lagged wage is also in logs.

Table 4: Returns to Education and Experience in the PSID, 1982-1996

	HLT 2003	Women		Men	
		OLS	FE	OLS	FE
education	.129 (.0002)	.168 (.007)		.146 (.006)	
experience	.130 (.0001)	.107 (.010)	.103 (.009)	.097 (.008)	.109 (.007)
experience squared	-.0023 (.000)	-.0011 (.000)	-.0017 (.000)	-.0020 (.000)	-.0025 (.000)
constant	6.89 (.003)	5.63 (.124)	8.04 (.062)	7.12 (.097)	9.05 (.051)

Note: The dependent variable is annual labor earnings in logs. The first column reports coefficients from the 1990 census for white males reported by Heckman, Lochner and Todd (2003). The next four columns report coefficients from the PSID. Following Heckman, Lochner and Todd (2003), experience is calculated in 1982 as age minus education minus 6. Experience is then updated for the years 1983-1996; increasing by one in years in which the individual works 500 or more hours.

Table 5: Estimates of Fertility Process (Probit).

dependent variable	
one child currently at home	.076 (.029)
two children currently at home	-.722 (.028)
married	.755 (.056)
constant	-1.73 (.055)
n	11180

Note: The dependent variable is a dummy equal to one if during the current year the household gives birth to a child.

Table 6: Annual Cost of Children, CEX 1980-1996

dependent variable	
cost for first child	486.0 (16.2)
cost for second child	370.8 (17.0)
cost for additional children	287.3 (9.98)

Note: The dependent variable is a measure of annual expenditure that includes food at home, child care, boy's clothing and shoes, girl's clothing and shoes, and infant clothing. The sample includes households with the head with ages 20 to 40. All dollar values have been deflated to reflect year 1984 dollars. Coefficients for after tax income, age of the household head, education of the household head, race and marital status are omitted. The regression includes 29,961 observations and the R^2 is .36.

Table 7: Household Production Regression, OLS

single woman	-8.7 (16.9)
married woman	230.0 (15.2)
married man	-30.4 (14.0)
one child	29.7 (13.3)
two or more children	68.1 (11.3)
one child \times woman	194.0 (19.1)
two or more children \times woman	370.9 (16.3)
working 500+ hours annually	-238.6 (16.5)
wage if working	-0.001 (.01)
labor supply	-0.135 (.006)
constant	924.8 (18.3)

Note: The dependent variable is reported annual hours spent on household production. The variable is derived from answers to the question, "About how much time do you [or does your spouse] spend on housework in an average week? I mean time spent cooking, cleaning, and doing other work around the house.". The sample is the same as described in the descriptive statistics. The excluded category is single men. The R^2 is .22.

Table 8: Estimated Coefficients and Simulation Results.

Coefficients		
	Data	Simulation
woman's σ	0.45	
man's σ	0.41	
kids' α	2.19	
wife's α	0.0	
husband's α	0.0	
match quality dispersion	0.1	
Annual Hours at Work:		
unmarried women	1,824	1,643
unmarried men	2,094	2,089
married women	1,534	1,479
married men	2,281	2,318
Labor Force Participation:		
unmarried women	.88	0.97
unmarried men	.94	0.999
married women	.71	0.87
married men	.97	0.99
Consumption:		
unmarried women	6,522	12,156
unmarried men	7,927	14,119
married couples	10,847	25,773
Wealth:		
unmarried women	8,180	9,198
unmarried men	14,149	10,228
married couples	49,732	21,276