The Marriage Market, Labor Supply, and Education Choice

Pierre-André Chiappori

Columbia University

Monica Costa Dias

Institute for Fiscal Studies and University of Porto

Costas Meghir

Yale University, National Bureau of Economic Research, Institute for the Study of Labor (IZA), and Institute for Fiscal Studies

We develop an equilibrium life cycle model of education, marriage, labor supply, and consumption in a transferable utility context. Individuals start by choosing their investments in education anticipating returns in the marriage market and the labor market. They then match on the basis of the economic value of marriage and preferences. Equilibrium in the marriage market determines intrahousehold allocation of resources. Following marriage households (married or single) save, supply labor, and consume private and public commodities under uncertainty. Marriage thus has the dual role of providing public goods and offering risk sharing. The model is estimated using the British Household Panel Survey.

This research has greatly benefited from discussions with Jim Heckman, Ana Reynoso, Jeremy Fox, Ilse Lindenlaub, Bob Polak, and Harald Uhlig. We thank six anonymous referees for their detailed and constructive comments, as well as participants in seminars at the University of Chicago, the Becker Friedman Institute conference in honor of Gary Becker and the conference in honor of Martin Browning, Rice University, the University of Washington in St. Louis, and the Human Capital and Economic Opportunity workshop on Human Capital Formation and Family Economics. Monica Costa Dias thanks the Economic and Social Research Council for funding (awards ES/N015304/1, ES/K00624X/1,

[*Journal of Political Economy*, 2018, vol. 126, no. S1]
 0022-3808/2018/126S1-0005$10.00

I. Introduction

A. *Matching on Human Capital*

The present paper explores the intersection of two fundamental Beckerian concepts: human capital and matching.[1] We are now used to considering education as an investment, whereby agents give up present consumption for higher income and consumption tomorrow. Similarly, we routinely think of marriage in terms of a matching game, in which couples create a surplus that is distributed between spouses, according to some endogenous rule that reflects equilibrium constraints. Still, the interaction between these notions remains largely unexplored.[2] In particular, whether individuals, on the marriage market, can be expected to match assortatively on human capital is largely an open question. For instance, in the presence of domestic production, in some cases one may expect negative assortative matching, a point stressed by Becker himself in his seminal 1974 contribution.

Even if household production is disregarded, the analysis of matching on human capital raises challenging questions. Recent work on the dynamics of wages and labor supply has emphasized the importance of productivity shocks, which typically take a multiplicative form. It follows that higher human capital comes with higher expected wages, but also possibly with more wage volatility. In such a context, whether an educated individual, receiving a large but highly uncertain income, will match with a similar spouse or will trade lower spousal expected income for a lower volatility is not clear. While any individual probably prefers a wealthier spouse, even at the cost of higher volatility, how this preference varies with the individual's own income process—the crucial determinant of assortativeness when intracouple transfers are allowed, which is our case—is far from obvious.

We believe that the interaction between educational choices and matching patterns is of crucial importance for analyzing the long-run effects of a given policy. When considering the consequences of, say, a tax reform, standard labor supply models, whether unitary or collective, typically take education and family composition as given. While such assumptions make perfect sense from a short-term perspective, they may severely bias our un-

and ES/M010147/1). Costas Meghir thanks the Cowles Foundation and the Institution for Social and Policy Studies for funding. Pierre-André Chiappori thanks the National Science Foudation (award 1124277) for funding. The usual disclaimer applies.

[1] See Becker (1991, 1994) for comprehensive accounts of his contributions on these themes.

[2] Among notable exceptions are Cole, Mailath, and Postlewaite (2001), Peters and Siow (2002), and Nöldeke and Samuelson (2015), whose contributions are discussed below. See also Hoppe, Moldovanu, and Sela (2009) for an extension to asymmetric information. However, empirical estimations of these effects remain quite rare.

derstanding of the reform's long-term outcome. Education policy, taxation, and welfare programs have a double impact on incentives to invest in human capital. On the one hand, they directly affect the returns from the investment perceived on the labor market. On the other hand, they also influence matching patterns, hence the additional returns reaped on the marriage market—the so-called "marital college premium," whose importance for human capital investment has been emphasized by several recent contributions (Chiappori, Iyigun, and Weiss 2009; Chiappori, Salanié, and Weiss 2017). Added to that is the effect that taxes and welfare have on insurance, which can also affect both marital patterns and investment in human capital. In the long run, these aspects may be of major importance.

The main motivation of the present paper is precisely to provide an explicit framework in which these effects can be conceptually analyzed and empirically quantified. Our model has several original features. Following a Beckerian tradition, we model marriage as a frictionless matching game in a transferable utility (TU) framework with risk-averse agents. Individual utilities have an economic and a noneconomic component. There are two economic gains from marriage: Spouses share a public good and also insure each other against productivity shocks. In addition, marriage provides idiosyncratic, nonmonetary benefits, which are additively separable and education-specific, as in Choo and Siow (2006) and Chiappori et al. (2017). The TU property implies that, once married, households behave as a single decision maker (unitary household). Despite its obvious shortcomings, this property considerably simplifies the analysis of the couple's dynamics of consumption and labor supply.

We consider a three-stage model and assume Pareto efficiency and full commitment. We abstract from issues relating to divorce, and our full commitment assumption precludes renegotiation; these are important questions we wish to address as this research agenda develops. Agents first independently invest in human capital; their decision is driven by their idiosyncratic characteristics—ability, cost of investment (which may, for instance, reflect borrowing constraints), and preferences for marriage—and the expected returns on investment, which is itself determined by the equilibrium prevailing on the relevant markets. In the second stage, individuals match on the marriage market, on the basis of their human capital and their idiosyncratic preferences for marriage. Finally, the last period is divided into T subperiods, during which couples or singles consume private and public goods, save, and supply labor subject to permanent and transitory wage shocks, very much as in standard life cycle models.

As is usual, such a game can be solved backward, starting with the third stage. Owing to the TU assumption, the analysis of the dynamic labor supply model exactly characterizes the total surplus generated by marriage, while it is compatible with any intracouple distribution of surplus. The matching game in the second stage is defined by the joint distribution of human capital and marital preferences among men and women, as resulting from investment during the first stage and the expected surplus generated in the third stage. Crucially, equilibrium conditions on the marriage market fully determine the intrahousehold allocation of the surplus for all possible levels of human capital. In particular, these conditions allow the characterization for each individual of the consequences, in expected utility terms, of the various levels of human capital they may choose to acquire. This "education premium," in turn, determines education decisions in the first stage. In essence, therefore, investments in the first stage are modeled under a rational expectations logic: agents anticipate a given vector of returns to education, and the resulting decisions lead to an equilibrium in the marriage market that is compatible with these expectations.

In this context, the impact of any given policy reform can be considered along several dimensions. Coming back to the tax reform example, the short-term impact can be analyzed from the dynamic labor supply model of the third stage: existing couples (and singles) respond to changes in income tax by adjusting their labor supply and their public and private consumption. From a longer-term perspective, however, matching on the marriage market will also be affected; typically, the respective importance of economic and noneconomic factors will vary, resulting in changes in the level of assortativeness on human capital, therefore, ultimately in inter- and intrahousehold inequality. Finally, the changes affect the returns on investment in human capital both directly (through their impact on after-tax income) and indirectly (by their consequences on the marriage market); they can therefore be expected to propagate to human capital investments. Similar arguments can be made for education policy or welfare reform. Imperfect as it may be, our approach is the first to consider all these aspects in a unified and theoretically consistent framework.

In the next subsection we discuss some of the existing literature. Then in Sections II and III, we present the model and develop its solution. Section IV discusses the identification of the distribution of marital preferences. In Section V we present the data, and in Section VI we present our estimation approach. In Sections VII and VIII we discuss the empirical results and a counterfactual simulation. Section IX presents concluding remarks and discusses future avenues of research.

B. Existing Literature

Our paper is a direct extension of the collective models of Chiappori (1988, 1992) and Blundell, Chiappori, and Meghir (2005), among others. In these models there is no time/dynamic dimension. This restriction is relaxed here. Thus the framework we use is directly related to intertemporal models of labor supply and savings over the life cycle, such as the model of Mazzocco (2007), who uses a collective framework, and Attanasio, Low, and Sánchez-Marcos (2008) and Low, Meghir, and Pistaferri (2010), who focus, respectively, on female and male labor supply. Similarly in a recent paper, Blundell et al. (2016) consider female labor supply over the life cycle in a context in which household composition is changing over the life cycle but exogenously. More closely to this paper, Low et al. (2018) allow for endogenous marriage decisions with limited commitment in a partial equilibrium context with frictions but treating education as exogenous. Goussé, Jacquemet, and Robin (2016) specify an equilibrium model of marriage and labor supply based on search frictions. Their model draws from Shimer and Smith (2000), and the complementarity arises from the production of public goods that depends on the wages of both spouses. Their model does not include savings, and the only source of uncertainty is exogenous divorce. Moreover, it does not allow for endogenous education choices. Finally, precursors of this paper are Chiappori et al. (2009), which specifies a theoretical model of education decisions, the marriage market, and time at home, and Chiappori et al. (2017), which provides an empirical estimation; however, both papers adopt a reduced-form specification in which marital gains are recovered from matching patterns without analyzing actual behavior.[3]

Our model is also related to recent developments on the econometrics of matching models under transferable utility (see Chiappori and Salanié [2016] for a recent survey). In particular, the stochastic structure representing idiosyncratic preferences for marriage is directly borrowed from Choo and Siow (2006) and Chiappori et al. (2017). Our framework, however, introduces several innovations. First, agents match on human capital, in contrast to Choo and Siow, where they match on age, and Chiappori et al., where they match on education. Human capital, in our framework, depends on education but also on innate ability. In principle, the latter is not observed by the econometrician. However, observing agents' wage and labor supply dynamics (during the third stage) allows us to recover the joint distribution of education and ability and therefore of human capital; interestingly, these distributions are sufficient to fully characterize the equilibrium. A second difference is that both Choo and Siow (2006) and Chiappori et al. (2017) identify the structural model under consideration from the sole observation of matching patterns. As a result, Choo and Siow's model is exactly identified under strong, parametric assumptions, whereas identification in Chiappori et al.'s model comes from the observation of multiple cohorts together with parametric restrictions on how surplus may change across cohorts. In our case, on the contrary, our structural model of household labor supply allows us to identify preferences and therefore the surplus function. The matching model, therefore, is overidentified and allows us to recover the intracouple allocation of sur-

[3] Theoretical models with prematch investments include Cole et al. (2001), Peters and Siow (2002), and Nöldeke and Samuelson (2015).

plus while generating additional, testable restrictions. Finally, this identification, together with the knowledge of the joint distribution of ability and education, enables us to explicitly model the process of educational choice. As a consequence, we can evaluate the long-term impact of a given policy reform on human capital formation. While the link between intrahousehold allocation and investment in human capital has already been analyzed from a theoretical perspective (see, e.g., Chiappori et al. 2009), our approach is, to the best of our knowledge, the first to explore it empirically through a full-fledged structural model.

II. The Model

A. *Time Structure*

We model the life cycle of a cohort of women $f \in \mathcal{F}$ and men $m \in \mathcal{M}$, so time and age will be used interchangeably and commonly represented by t. The individual's life cycle is split into three stages, indexed 1–3. In stage 1, individuals each draw some ability level θ, a cost of education c, and a vector of marital preferences; then they invest in human capital by choosing one of three educational levels. This investment depends on their innate ability and their cost of education, as well as on the perceived benefits of this investment, including the benefits to be received on the marriage market; the latter, in turn, are directly influenced by marital preferences. The ability of agent i, denoted θ_i, belongs to a finite set of classes, $\Theta = \{\theta^1, \dots, \theta^N\}$. Education costs are continuously distributed on some compact set $\mathcal{C}$, and the agent can choose between a finite number of education levels, $\mathcal{S} = \{S^1, \dots, S^J\}$. At the end of period 1, each agent is thus characterized by human capital (or productivity type) $H(s, \theta)$, which is a summary measure of education and innate ability. The distribution of human capital has a finite support $\mathcal{H}$ of cardinality (at most) $J \times N$. So at this stage the agent belongs to a finite set of classes $\mathcal{H} = \{H^1, \dots, H^{J\times N}\}$ that fully characterize his or her prospects in the marriage and labor markets.

In stage 2, individuals enter the marriage market; the latter is modeled as a frictionless matching process based on the level of human capital and on marital preferences. At the end of stage 2, some individuals are married whereas the others remain single forever.

Stage 3 (the "working life" stage) is divided into T periods; in each period, individuals, whether single or married, observe their (potential) wage and nonlabor income and decide on consumptions and labor supplies. Credit markets are assumed complete, so that during their active life, agents can borrow or save at the same interest rate. Following a collective logic (Chiappori 1988, 1992), decisions made by married couples are assumed Pareto efficient. Moreover, the intrahousehold allocation of private consumption (therefore of welfare) is endogenous and is determined by commitments made at the matching stage. In particular, we do not consider divorce or separation in this model.

B. *Economic Utilities*

The lifetime utility of agent i is the sum of three components. The first is the expected, discounted sum of economic utilities generated during the T periods of i's third stage of life by consumptions and labor supply; the second is the subjective utility of marriage (or singlehood) generated by the agents' marital preferences; and the third is the utility cost of education attendance. In what follows, we consider the following economic utilities at date t of stage 3:

$$u_{it}(Q_t, C_{it}, L_{it}) = \ln(C_{it}Q_t + \alpha_{it}L_{it}Q_t), \tag{1}$$

where L is nonmarket time and C and Q are private and public consumptions, respectively. We take labor supply choices to be discrete: agents choose either to participate in the labor market ($L = 0$) or not to ($L = 1$).

The choices of consumptions, labor supply, and savings are driven by time-varying preferences and income. First, wages at age t are determined by the person's age and human capital, itself a function of education s_i and ability θ_i, and also by an idiosyncratic productivity shock that may have a transitory and a permanent component. Formally,

$$w_{it} = W_G(H_i, t)e_{it}, \tag{2}$$

where w_{it} denotes i's earnings at age t; $G = M, F$ indexes i's gender group; W_G is the aggregate, gender-specific price of human capital class H_i at age t; $H_i = H(s_i, \theta_i)$ is i's human capital; and e_{it} is an idiosyncratic shock. Second, preferences may vary; in practice, the α_{it} are random variables.

Two remarks can be made on these utilities. From an ordinal viewpoint, they belong to Bergstrom and Cornes's generalized quasi-linear family. As a consequence, at any period and for any realization of family income, they satisfy the TU property. For a given couple (m, f), any conditional (on employment and savings) Pareto efficient choice of consumption and public goods maximizes the sum of the spouse's exponential of utilities:[4]

[4] In the static model, one can use exp u_i as a particular cardinalization of i's preferences. Then any (ex post) efficient allocation maximizes some weighted sum of utilities of the form $\exp u_i(Q_t, C_{it}, L_{it}) + \mu \exp \mu_j(Q_t, C_{jt}, L_{jt}) \geq \bar{\mu}_j$ under a budget constraint. Here, the maximand is equal to

$$(C_{it} + \mu C_{jt} + \alpha_{it}L_{it} + \mu\alpha_{jt}L_{jt})Q_t,$$

and the first-order conditions with respect to private consumptions (assuming the latter are positive) give

$$Q_t = \lambda_t = \mu Q_t,$$

$$\exp u_i(Q_t, C_{it}, L_{it}) + \exp u_j(Q_t, C_{jt}, L_{jt}) \\ = (C_{it} + C_{jt} + \alpha_{it}L_{it} + \alpha_{jt}L_{jt})Q_t. \tag{3}$$

Solving this program gives the optimal choice of aggregate household private and public consumptions at each period, conditional on labor supplies and savings. The latter are then determined from a dynamic perspective by maximizing the expected value of the discounted sum (over periods t to the end of life) of utilities.

From a cardinal perspective, the Von Neumann–Morgenstern utilities defined by (1) belong to the ISHARA class, defined by Mazzocco (2007). By a result due to Schulhofer-Wohl (2006), this implies that the TU property also obtains ex ante, in expectations. In particular, there exists a specific cardinalization of each agent's lifetime economic utility such that any household maximizes the sum of lifetime utilities of its members, under an intertemporal household budget constraint. Specifically, we show below the following result. Take a couple (m, f) with respective human capital H_m and H_f, and let V_m and V_f denote their respective lifetime expected utility. There exists a function $\Upsilon(\mathbf{H})$, where $\mathbf{H} = (H_m, H_f)$, such that the set of Pareto efficient allocations is characterized by

$$\exp\left\{\frac{1-\delta}{1-\delta^T} V_m\right\} + \exp\left\{\frac{1-\delta}{1-\delta^T} V_f\right\} = \exp\left\{\frac{1-\delta}{1-\delta^T} \Upsilon(\mathbf{H})\right\}.$$

This function is explicitly derived in Section III. The crucial point now is that the expression

$$\overline{U}_i = \exp\left\{\frac{1-\delta}{1-\delta^T} V_i\right\} \tag{4}$$

is an increasing function of V_i; therefore, in the stage 2 matching game, it is a specific (and convenient) representation of i's (expected) utility. If we define

$$\Gamma(H_m, H_f) = \exp\left\{\frac{1-\delta}{1-\delta^T} \Upsilon(\mathbf{H})\right\},$$

the previous relationship becomes

$$\overline{U}_m + \overline{U}_f = \Gamma(H_m, H_f), \tag{5}$$

which shows that we are in a TU context even ex ante, since the Pareto frontier is, for these utility indices, a straight line with slope −1 for all

where λ_t is the Lagrange multiplier of the budget constraint. It follows that $\mu = 1$, implying that any Pareto efficient solution with positive private consumptions must maximize the sum of exp u_i.

wages and incomes. The function $\Gamma(H_m, H_f)$, when evaluated at the point of marriage, is the economic value generated by marriage. An important consequence is that, throughout the third stage (their working life), couples behave as a single decision maker maximizing the function Γ (or, equivalently, Υ). In particular, a standard, unitary model of dynamic labor supply can be used at that stage.

Alternatively, agents may choose to remain single; then they maximize the discounted sum of expected utility under an intertemporal, individual budget constraint. We denote $V_m^s(H_m)$ and $V_f^s(H_f)$ the respective lifetime economic utility of a single male (female) with human capital H_m (H_f). Note that these expressions, again, are expectations taken over future realizations of the preferences and wage shocks; they are contingent on the information known at the date of marriage, namely, each person's ability and education, as summarized by the person's human capital. In line with the previous notation, we then define

$$\overline{U}_i^s = \exp\left\{\frac{1-\delta}{1-\delta^T} V_i^s(H_i)\right\}. \tag{6}$$

Finally, for any man m with human capital H_m and any woman f with human capital H_f, the difference between the economic value that would be generated by their marriage, $\Gamma(H_m, H_f)$, and the sum of m's and f's respective expected utility as singles is the economic surplus generated by the marriage. Again, it depends only on both spouses' productivity and education and is denoted

$$\Sigma(H_m, H_f) = \Gamma(H_m, H_f) - \overline{U}_m^s - \overline{U}_f^s . \tag{7}$$

Note that all these expressions refer to the same cardinalization of lifetime expected utilities, given by (4).

C. *Marital Preferences*

Our representation of marital preferences follows that of Choo and Siow (2006) and Chiappori et al. (2017). Before investing, agent i draws a vector $\beta_i = (\beta_i^0, \beta_i^H \text{ where } H \in \mathcal{H})$, where β_i^H represents i's subjective satisfaction of being married to a spouse with human capital H and β_i^0 denotes his or her subjective satisfaction of remaining single. We assume that the total gain generated by the marriage of man m with human capital H_m and woman f with human capital H_f is the sum of the economic gain $\Gamma(H_m, H_f)$ defined above and the idiosyncratic preference shocks β:

$$\Gamma_{mf} = \Gamma(H_m, H_f) + \beta_m^{H_f} + \beta_f^{H_m}, \tag{8}$$

and the resulting surplus is

$$\Sigma_{mf} = \Sigma(H_m, H_f) + (\beta_m^{H_f} - \beta_m^0) + (\beta_f^{H_m} - \beta_f^0). \tag{9}$$

Again, the function $\Sigma(H_m, H_f)$ is defined as the expected economic lifetime surplus for a couple with human capital composition (H_m, H_f), over and above what they would each obtain as singles. The remaining part of the expression relates to the noneconomic benefits of marriage.[5] Importantly, it is a restriction of this model (often referred to as the "separability property") that the idiosyncratic preferences of m, as described by the random vector β_m, depend only on the education of m's spouse, not on her identity. In other words, individual, nonpecuniary preferences are over people with different levels of human capital, not over specific persons. This assumption is crucial because it allows us to fully characterize the stochastic distribution of individual utilities at the stable match. Indeed, in a TU model, stability is equivalent to surplus maximization, a linear programming problem, and individual utilities emerge as the solutions to the dual problem. Here, the surplus entails a stochastic component; characterizing the distribution of the dual variables in that framework is an extremely difficult problem, to which the solution is not known except in the separable case. Moreover, the introduction, in the marital gain generated by the couple (m, f), of match-specific terms of the form ε_{mf} would raise specific difficulties in our frictionless framework. For instance, if the ε are assumed independently and identically distributed (iid), then when the number of individuals becomes large, the fraction of singles goes to zero (and their conditional utility tends to infinity), whereas a more complex correlation structure would typically be underidentified. The interested reader is referred to Chiappori and Salanié (2016) for a detailed discussion.[6]

D. Second-Stage Matching Game

At the end of the first stage, agents are each characterized by their marital preferences, but also by their human capital H, a function of their innate ability θ and education s. From the second perspective, the male and female populations are therefore distributed over the space $\mathcal{H}$, which consists of $N \times J$ classes. They then enter a matching game under TU, in which the surplus function for any potential match is given by (9).

[5] It should be stressed that our interpretation of β_i^H as i's subjective utility of being married to a spouse with human capital H is by no means the only possible one. Alternatively, β_i^H could be some unobserved characteristic of i that is identically valued by all spouses with human capital H. The crucial property is that this term enhances total surplus in a way that does not depend on the spouse's identity, but only on her or his human capital.

[6] The separability assumption is fully compatible with a model involving domestic production, even under complementarity in the parents' observable characteristics (e.g., human capital) in the production function; however, it rules out complementarities in the unobserved components. See Chiappori et al. (2017) for a model along these lines.

As usual, a matching is defined by a measure on the product space of male and female characteristics (i.e., $\mathcal{H} \times \mathcal{H}$) and two sets of individual utility levels, (U_m) and (U_f), such that for any pair (m, f) on the support of the measure, that is, for any couple that matches with positive probability,

$$U_m + U_f = \Gamma_{mf}.$$

Intuitively, the pair (U_m, U_f) describes how the total gain g_{mf} generated by the possible marriage of m and f would be divided between the spouses. The matching is stable if (i) no married person would rather be single and (ii) no two individuals would strictly prefer being married to each other to remaining in their current situation. A direct consequence is that for any pair (m, f), it must be the case that[7]

$$U_m + U_f \geq \Gamma_{mf}.$$

Now, a crucial result by Chiappori et al. (2017) is the following:

Theorem 1 (Choo and Siow 2006; Chiappori et al. 2017). If the surplus is given by (9), then there exist $2(NJ)^2$ numbers—$\overline{U}_M(H_m, H_f)$ and $\overline{U}_F(H_m, H_f)$ for $(H_m, H_f) \in \mathcal{H}^2$—such that

1. for any (H_m, H_f),

$$\overline{U}_M(H_m, H_f) + \overline{U}_F(H_m, H_f) = \Gamma(H_m, H_f); \tag{10}$$

2. for any m with human capital H_m married to f with human capital H_f,

$$\begin{aligned} U_m &= \overline{U}_M(H_m, H_f) + \beta_m^{H_f}, \\ U_f &= \overline{U}_F(H_m, H_f) + \beta_f^{H_m}. \end{aligned} \tag{11}$$

Proof. See Chiappori et al. (2017).

In words, the utility of any man m at the stable matching is the sum of a deterministic component, which depends only on his and his spouse's human capital, and of m's idiosyncratic net preference for marrying a spouse with that human capital; the same type of result obtains for women. For notational consistency, if i remains single, we consider the class of his spouse to be 0.

Note that the characterization of utilities provided by (11) refers to a specific cardinalization of individual utilities, defined by (U_m, U_f); technically, this is the particular cardinalization that exhibits the TU property. Obviously, it can equivalently be translated into the initial cardinalization; in that case, the total, expected utility of person i is

[7] If this inequality was violated for some couple (m, f), one could conclude that m and f are not matched (then an equality would obtain) but should be matched (since each of them could be made better off than their current situation), a violation of stability.

$$V_i = \frac{1-\delta^T}{1-\delta}\ln(U_i) = \frac{1-\delta^T}{1-\delta}\ln\left(\overline{U}_{g_i}\left(H_m, H_f\right) + \beta_i^{H_j}\right), \tag{12}$$

where g_i is the gender of i and H_j denotes the human capital of i's spouse.

An immediate corollary is the following:

Corollary 1.

1. For any man m with human capital H_m, m's spouse at the stable matching has human capital H_f if and only if the following inequalities hold for all $H \in \mathcal{H} \cup \{0\}$:
$$\overline{U}_M(H_m, H_f) + \beta_m^{H_f} \geq \overline{U}_M(H_m, H) + \beta_m^H.$$
2. For any woman f with human capital H_f, f's spouse at the stable matching has human capital H_m if and only if the following inequalities hold for all $H \in \mathcal{H} \cup \{0\}$:
$$\overline{U}_F(H_m, H_f) + \beta_f^{H_m} \geq \overline{U}_F(H, H_f) + \beta_f^H.$$
3. The utility of a man m with human capital H_m and preferences β_m is
$$A_M(H_m) = \max_{H_f \in \mathcal{H} \cup \{0\}}\left(\overline{U}_M\left(H_m, H_f\right) + \beta_m^{H_f}\right), \tag{13}$$
and the utility of a female agent f with human capital H_f and preferences β_m is
$$A_F(H_f) = \max_{H_m \in \mathcal{H} \cup \{0\}}\left(\overline{U}_F\left(H_m, H_f\right) + \beta_f^{H_m}\right). \tag{14}$$

The main implication of this result is that marital choices in stage 2 can be modeled as individual, discrete choice problems, in which the thresholds $\overline{U}_M(H_m, H_f)$ and $\overline{U}_F(H_m, H_f)$ can be identified using standard techniques. Note, however, that these parameters are not independent, since they have to satisfy the restrictions (10); we will return to this point later on. Also, note that these ex ante expected utilities depend only on the individual's stock of human capital.

E. *First Stage: The Education Choice*

In the first stage of life, individuals decide on the level of educational investment. We assume there are three choices, corresponding to three classes in $\mathcal{S}$: statutory schooling, high school, and college. Each level of education s is associated with a cost $c_s(X, v_s)$, where X are observable characteristics and v_s is an unobservable cost.

Defining human capital as a function of schooling and ability $H(s, \theta)$, education choice is defined as follows: for man m,

$$s_m = \arg\max_{s \in \mathcal{S}}\{A_M(H(s, \theta_m)) - c_s(X_m, v_{ms})\}; \tag{15}$$

for woman f,

$$s_f = \arg\max_{s \in \mathcal{S}} \{A_F(H(s, \theta_f)) - c_s(X_f, v_{fs})\}, \tag{16}$$

where A_M and A_F are defined in equations (13) and (14) for males and females, respectively, and where the subscript s indexes schooling level s. Individuals are assumed to know their ability at that point, but this may not be observable by the econometrician. Education choice takes into account both the returns in the labor market and the returns in the marriage market, which are embedded in the value functions for each choice.

Finally, the structure of that stage is a simultaneous move game: agents each choose their education independently, but the payoffs they will receive depend on the human capital distribution on both sides of the market, which results from other players' investment. This, potentially, raises existence and uniqueness issues that will be discussed below.

III. Solving the Model

It is instructive to outline the solution of the problem. As is standard in dynamic models of the life cycle, the model is solved working backward from the end of life. We therefore start with the last period of the third stage. As mentioned before, the TU property implies that any married couple behaves as a single decision maker maximizing the sum of the spouses' (exponential of) utilities: the Pareto weights associated with our original logarithmic cardinalization of utilities, which determine the intrahousehold allocation of welfare, do not affect aggregate household consumption, savings, and individual labor supply decisions. Singles maximize their own utility. Both maximizations are subject to an intertemporal budget constraint.

A. *Employment, Consumption, and Savings during the Working Life*

We start with the labor supply and consumption decisions. The form of preferences allows us to easily derive consumptions from savings and labor supply choices; savings are then chosen to satisfy the conditional (on labor supply) intertemporal optimality condition; optimal labor supply is then the solution to a discrete choice problem.

1. General Solution to the Couple's Problem in Period t

In appendix A (apps. A–D are available online) we derive the solution to the last period of life, T. Many of the properties of that last period, such as the separability of the Pareto weights in the individual value function, carry over to the general solution for any of the earlier periods. Here we show the form of the solution for an earlier period, $t < T$.

Consumptions.—Each period/age t sees the arrival of new information on each spouse's preferences for working and productivity, $\boldsymbol{\alpha}_t = (\alpha_{mt}, \alpha_{ft})$ and $\mathbf{e}_t = (e_{mt}, e_{ft})$. Choice is also conditional on the other circumstances faced by the couple, namely, savings carried over from the previous period, K_{t-1}, and the spouses' human capital, $\mathbf{H} = (H_m, H_f)$. Given the information set $(\boldsymbol{\alpha}_t, \mathbf{e}_t, K_{t-1}, \mathbf{H})$, we first consider the couple's consumption decisions conditional on savings and employment, K_t and $\mathbf{L}_t = (L_{mt}, L_{ft})$. For the within-period problem of resource allocation to private consumption (C) and public good (Q), we can use the exponential cardinalization of individual preferences. The couple thus solves

$$\max_{Q_t, C_t} Q_t\left(C_t + \alpha_{mt}L_{mt} + \alpha_{ft}L_{ft}\right)$$

under the budget constraint

$$w_{mt} + w_{ft} + y_t^C + RK_{t-1} = K_t + C_t + w_{mt}L_{mt} + w_{ft}L_{ft} + pQ_t.$$

Here $w_{mt} + w_{ft}$ is the couple's total ("potential") labor income in period t, and y_t^C is the couple's nonlabor income. Note that the latter may depend on individual labor supplies and earnings, which allows for means-tested benefits and taxes as well as benefits that depend on participation, such as unemployment insurance or earned income tax credits. Wages are as defined in equation (2) and considered net of income taxes. Finally, R is the risk-free interest rate at which savings accumulate over periods, $C_t = C_{mt} + C_{ft}$ is total expenditure in the private consumption of spouses, and pQ_t is total expenditure in the public good.

Conditional on savings and labor supply, the solutions for public and private consumptions are

$$\begin{aligned} Q_t(K_t, \mathbf{L}_t) &= [y_t^C + RK_{t-1} - K_t + w_{mt}(1 - L_{mt}) + w_{ft}(1 - L_{ft}) \\ &\quad + (\alpha_{mt}L_{mt} + \alpha_{ft}L_{ft})]/2p, \\ C_t(K_t, \mathbf{L}_t) &= y_t^C + RK_{t-1} - K_t + w_{mt}(1 - L_{mt}) + w_{ft}(1 - L_{ft}) \\ &\quad - pQ_t(K_t, \mathbf{L}_t) \\ &= pQ_t(K_t, \mathbf{L}_t) - (\alpha_{mt}L_{mt} + \alpha_{ft}L_{ft}), \end{aligned}$$

where consumptions are written as functions of $(K_t, \mathbf{L}_t)$ to highlight the fact that they are conditional solutions.

Efficient risk sharing conditional on savings and employment.—We now consider the intrahousehold allocation of resources during period t from an ex ante perspective—that is, before the realization of the shocks. Here, efficiency relates to sharing the wage and preference risks. In this context, it requires the maximization of a weighted sum of expected utilities, using the initial, logarithmic cardinalization, which reflects preferences toward risk. If μ denotes the wife's Pareto weight corresponding to that cardinalization, the standard efficiency condition imposes that the ratio of marginal utilities of private consumption be constant (and equal to the Pareto weight) for all periods and all realizations of the random shocks (see, e.g., Townsend 1994):

$$\frac{\partial u_{mt}(Q_t, C_{mt}, L_{mt})}{\partial C_{mt}} = \mu \frac{\partial u_{ft}(Q_t, C_{ft}, L_{ft})}{\partial C_{ft}}.$$

Note that the Pareto weight μ is a price endogenously determined in the marriage market. Thus, it depends only on the information available then, namely, the human capital of both spouses (H_m, H_f). Moreover, it remains constant over the couple's working life—a direct implication of efficiency under full commitment. Efficient risk sharing then yields private consumptions as follows:

$$C_{mt} = \frac{1}{1+\mu} pQ_t - \alpha_{mt} L_{mt},$$

$$C_{ft} = \frac{\mu}{1+\mu} pQ_t - \alpha_{ft} L_{ft}.$$

Therefore, the conditional (on employment and savings) instantaneous indirect utilities are

$$v_{mt} = 2 \ln Q_t(K_t, \mathbf{L}_t) + \ln p + \ln \frac{1}{1+\mu}, \tag{17}$$

$$v_{ft} = 2 \ln Q_t(K_t, \mathbf{L}_t) + \ln p + \ln \frac{\mu}{1+\mu}. \tag{18}$$

Note that Q_t is also a function of the entire state space, including the wage and preference shocks, savings, and human capital, $(\mathbf{e}_t, \boldsymbol{\alpha}_t, K_{T-1}, \mathbf{H})$. We therefore write $v_{it}(K_t, \mathbf{L}_t; \mathbf{e}_t, \boldsymbol{\alpha}_t, K_{t-1}, \mathbf{H}, \mu)$.

Expected value functions.—Appendix A shows that, for period T,

$$E_{T|T-1} V_{mT}(\mathbf{e}_T, \boldsymbol{\alpha}_T, K_{T-1}, \mathbf{H}, \mu)$$
$$= I_T(\mathbf{e}_{T-1}, \boldsymbol{\alpha}_{T-1}, K_{T-1}, \mathbf{H}) + \ln \frac{1}{1+\mu},$$
$$E_{T|T-1} V_{fT}(\mathbf{e}_T, \boldsymbol{\alpha}_T, K_{T-1}, \mathbf{H}, \mu)$$
$$= I_T(\mathbf{e}_{T-1}, \boldsymbol{\alpha}_{T-1}, K_{T-1}, \mathbf{H}) + \ln \frac{\mu}{1+\mu},$$

where

$$I_T(\mathbf{e}_{T-1}, \boldsymbol{\alpha}_{T-1}, K_{T-1}, \mathbf{H}) = E_{T|T-1} \max_{L_T, K_T} [2 \ln Q_T(\mathbf{L}_T, K_T) + \ln p | \mathbf{e}_{T-1}, \boldsymbol{\alpha}_{T-1}],$$

where expectations are taken over the (education-specific) distribution of $(\mathbf{e}_t, \boldsymbol{\alpha}_t)$ conditional on their realization at $t-1$. Note that here $K_T = 0$ since bequests are not being considered. Given the conditional instantaneous indirect utilities in (17)–(18), it is easy to show by recursion that the additive separability of the Pareto weight carries over to earlier periods:

$$E_{t|t-1} V_{mt}(\mathbf{e}_t, \boldsymbol{\alpha}_t, K_{t-1}, \mathbf{H}, \mu) = I_t(\mathbf{e}_{t-1}, \boldsymbol{\alpha}_{t-1}, K_{t-1}, \mathbf{H}) + \ln\left(\frac{1}{1+\mu}\right) \sum_{\tau=t}^{T} \delta^{\tau-t},$$

$$E_{t|t-1} V_{ft}(\mathbf{e}_t, \boldsymbol{\alpha}_t, K_{t-1}, \mathbf{H}, \mu) = I_t(\mathbf{e}_{t-1}, \boldsymbol{\alpha}_{t-1}, K_{t-1}, \mathbf{H}) + \ln\left(\frac{\mu}{1+\mu}\right) \sum_{\tau=t}^{T} \delta^{\tau-t},$$

where δ is the discount factor. The common term in the individual value functions, I_t, is defined recursively by

$$I_t(\mathbf{e}_{t-1}, \boldsymbol{\alpha}_{t-1}, K_{t-1}, \mathbf{H}) = E_{t|t-1} \max_{L_t, K_T} [2 \ln Q_t(\mathbf{L}_t, K_t) + \ln p + \delta I(\mathbf{e}_t, \boldsymbol{\alpha}_t, K_t, \mathbf{H}) | \mathbf{e}_{t-1}, \boldsymbol{\alpha}_{t-1}],$$

where expectations are taken over the (education-specific) distribution of $(\mathbf{e}_t, \boldsymbol{\alpha}_t)$ conditional on $(\mathbf{e}_{t-1}, \boldsymbol{\alpha}_{t-1})$. A crucial feature of the above expressions is that the Pareto weight μ affects individual welfare but drops out of the aggregate value function I, reflecting the TU property. This then implies that the intertemporal optimality condition for savings (Euler equation) is the same for both spouses. For any choice of labor supplies (including the optimal one), conditional optimal savings $(K_t^*(\mathbf{L}_t))$ satisfy

$$2 \frac{\partial \ln Q_t(K_t, \mathbf{L}_t)}{\partial K_t} + \delta \frac{\partial I_{t+1}(\mathbf{e}_t, \boldsymbol{\alpha}_t, K_t, \mathbf{H})}{\partial K_t} = 0.$$

Finally, the optimal choices of labor supplies are defined by

$$\left(L_{mt}^*, L_{ft}^*\right) = \arg\max_{\mathbf{L}_t \in \{0,1\}^2} \{2 \ln Q_{it}(K_t^*(\mathbf{L}_t), \mathbf{L}_t) + \ln p + \delta I_{t+1}(\mathbf{e}_t, \boldsymbol{\alpha}_t, K_t^*(\mathbf{L}_t), \mathbf{H})\}.$$

The single's problem is a close replica of the couple's problem, just simpler, and its solution can be derived using the same approach as briefly discussed in appendix B.

2. The First Period after Marriage

The Markov processes for $(\mathbf{e}_t, \boldsymbol{\alpha}_t)$ start at date $t = 1$, and initial savings are set to zero. So the functions I_1 and I_1^S do not depend on past values of the shock or on past investment, but only on human capital; we denote them, respectively, by $\Upsilon(\mathbf{H})$ and $\Upsilon^S(H_i)$. It follows that the expected economic utility, at marriage, of each spouse is given by

$$V_m(\mathbf{H}, \mu) = \Upsilon(\mathbf{H}) + \left(\sum_{\tau=0}^{T-t} \delta^\tau\right) \ln\left(\frac{1}{1+\mu}\right) \tag{19}$$

and

$$V_f(\mathbf{H}, \mu) = \Upsilon(\mathbf{H}) + \left(\sum_{\tau=0}^{T-t} \delta^\tau\right) \ln\left(\frac{\mu}{1+\mu}\right), \tag{20}$$

which depends on the spouses' respective levels of human capital and on the Pareto weight μ that results from the matching game in the earlier life cycle stage 2. For singles, expected lifetime utility is simply

$$V^\varsigma(H_i) = \Upsilon^\varsigma(H_i).$$

B. Matching

We now move to the second stage, that is, the matching game. Remember that marriage decisions are made before preferences and productivity shocks $(\boldsymbol{\alpha}, \mathbf{e})$ are realized and that we assume full commitment. We first compute the expected utility of each spouse, conditional on the Pareto weight μ. We then show that the model can be reinterpreted as a matching model under TU; finally, we compute the equilibrium match and the corresponding Pareto weights.

Formal derivation.—Consider a match between a man with human capital H_m and a woman with human capital H_f. The spouses' expected, economic lifetime utilities are given by (19)–(20). However, an alternative cardinalization, already introduced in (4), turns out to be more convenient here. Specifically, define $\bar{U}_i$ by

$$\overline{U}_i = \exp\left\{\frac{1-\delta}{1-\delta^T} V_i\right\}. \tag{21}$$

Then if $\mathbf{H} = (H_m, H_f)$,

$$\overline{U}_m \exp\left\{-\Upsilon(\mathbf{H}) \frac{1-\delta}{1-\delta^T}\right\} = \frac{1}{1+\mu},$$

$$\overline{U}_f \exp\left\{-\Upsilon(\mathbf{H}) \frac{1-\delta}{1-\delta^T}\right\} = \frac{\mu}{1+\mu}.$$

Therefore,

$$\overline{U}_m + \overline{U}_f = \exp\left\{\frac{1-\delta}{1-\delta^T}\Upsilon(\mathbf{H})\right\} = \Gamma(\mathbf{H}),$$

which expresses that the individual economic utilities add up to the marital gain $\Gamma(\mathbf{H})$. Finally, we can add the idiosyncratic shocks to both sides of this equation; we finally have that, for any married couple $\mathbf{H} = (H_m, H_f)$,

$$\overline{U}_m + \beta_m^{H_f} + \overline{U}_f + \beta_f^{H_m} = \Gamma(\mathbf{H}) + \beta_m^{H_f} + \beta_f^{H_m} = g_{mf}. \tag{22}$$

The matching game, therefore, has a transferable utility structure: if the utility of person i is represented by the particular cardinal representation $(\bar{U}_i)$, then the Pareto frontier is a straight line with slope −1.

In particular, whether matching will be assortative on human capital or not depends on the supermodularity of function Γ, given iid shocks (β_m, β_f). One can easily check that the sign of the second derivative $\partial^2\Gamma/\partial H_m \partial H_f$ is indeterminate (and can be either positive or negative depending on the parameters), so this needs to be investigated empirically.[8]

Clearly, one can equivalently use any of the two cardinalizations described before to study marital sorting, because under TU matching patterns are driven by ordinal preferences; remember, though, that the Pareto weight μ refers to the initial cardinalization (V_m, V_f). This Pareto weight μ is match-specific; as such, it might in principle depend on the spouses' stocks of human capital, but also on their marital preferences. However, the following result, which is a direct corollary of theorem 1, states that this cannot be the case:

Corollary 2. At the stable match, consider two couples (m, f) and (m', f') such that $H_m = H_{m'}$ and $H_f = H_{f'}$. Then the Pareto weight is the same in both couples.

Proof. From (11) in theorem 1, we have that

$$U_m = \overline{U}_m + \beta_m^{H_f} = \overline{U}_M(H_m, H_f) + \beta_m^{H_f},$$
$$U_f = \overline{U}_f + \beta_f^{H_m} = \overline{U}_F(H_m, H_f) + \beta_f^{H_m}.$$

It follows that

$$\overline{U}_m = \overline{U}_M(H_m, H_f), \quad \overline{U}_f = \overline{U}_F(H_m, H_f).$$

[8] A result due to Graham (2011) states that, with this particular iid stochastic structure, for any two levels H and $\overline{H}$ of human capital, the total number of "assortative couples" (i.e., H and H or $\overline{H}$ and $\overline{H}$) will exceed what would be expected under purely random matching if and only if the deterministic function Γ is supermodular for H and $\overline{H}$, i.e.,

$$\Gamma(H,H) + \Gamma(\overline{H},\overline{H}) \geq \Gamma(\overline{H},H) + \Gamma(H,\overline{H}).$$

That is what is meant by assortative matching.

Since

$$\overline{U}_i = \exp\left\{\frac{1-\delta}{1-\delta^T} V_i(\mathbf{H}, \mu)\right\},$$

we conclude that μ depends only on (H_m, H_f). QED

C. The First Stage in the Life Cycle: Education Choice

The solution to the matching problem allows us to construct the expected value of marriage for males and females, conditional on each of the three education levels. At this point the stochastic structure is provided by the realization of random marital preferences and the costs of education, which can include exogenous shifters. Given this, the education choice is described in equations (15)–(16).

1. Existence: The Nöldeke-Samuelson Approach

As mentioned earlier, the first stage can be modeled as a normal form noncooperative game, where each player's payoff depends on the other players' decisions. As such, existence has to be demonstrated, and neither uniqueness nor efficiency is guaranteed. We now discuss these issues.

Start with existence. The central idea, due to Cole et al. (2001) and Nöldeke and Samuelson (2015), is to consider what we shall call an *auxiliary game*, defined as follows. Assume that stages 1 (investment) and 2 (matching), instead of taking place sequentially, are simultaneous. That is, consider the following two-stage game:

- At stage 1, agents match (on the basis of their idiosyncratic characteristics, namely, ability, education costs, and marital preferences) and choose their education. In particular, matched pairs jointly (and efficiently) choose their respective investments in human capital.
- Stage 2 (the "working life" stage) is identical to stage 3 of the initial game; that is, it is divided into T periods, during which individuals, whether single or married, observe their (potential) wage and nonlabor income and decide on consumptions and labor supplies.

Again, the auxiliary game can be solved by backward induction. The behavior of a given couple is described as before; in particular, and using the same cardinalizations as in Section III.B, the function Γ defined by (22) still characterizes the total surplus generated by a given match.

We now introduce the following lemma:

LEMMA 1. The auxiliary game admits a stable matching.

Proof. The auxiliary game is a matching game, in which agents are each defined by their initial characteristics (ability, cost of education, and marital preferences), and the surplus function is

$$\tilde{\Gamma}(\theta, \beta, v, X_m, X_f) = \max_{s_m, s_f} \Gamma(H_m(s_m, \theta_m), H_f(s_f, \theta_f)) + \beta_m^{H_f} + \beta_f^{H_m} - c_s(X_m, v_{ms}) - c_s(X_f, v_{fs}).$$

In this context, the existence of a stable matching is equivalent to the existence of a measure on the product space of characteristics, with given marginals, that maximizes total surplus over the product space. This is a standard optimal transportation problem. From the earlier assumptions the surplus is continuous in the random preferences β and random costs of education v, whereas ability θ can take only a finite number of values (here two); this guarantees the existence of a maximum for the surplus, and existence of a stable matching then stems from standard results in optimal transportation (see Villani 2009; Chiappori, McCann, and Neisheim 2010). QED

The main result is then as follows:

Proposition 1. The first-stage game has a Nash equilibrium.

Proof. The proof is based on a central result by Nöldeke and Samuelson (2015, propositions 1 and 2), which states that any stable matching of the auxiliary game can be implemented as a Nash equilibrium of the initial game. This applies in our case because the strategy spaces and the set of abilities are discrete, whereas the surplus is continuous in the other characteristics, which guarantees the existence of a maximum.[9] From the previous lemma, the auxiliary game admits a stable matching, which completes the proof. QED

In other words, existence of an equilibrium is never an issue in our model, in sharp contrast with other contexts.[10]

2. Efficiency and Uniqueness

The implications of the previous result go beyond existence. The auxiliary game is a standard matching game under TU; stability, in this context,

[9] An equilibrium in our initial game corresponds to Nöldeke and Samuelson's notion of an ex post equilibrium, whereas a stable match of the auxiliary game is an ex ante equilibrium. Note that since costs and preferences are continuously distributed, the equilibrium is in pure strategies for almost all agents: if an agent is indifferent between two strategies, any agent with a close enough investment cost will strictly prefer one of the two; the logic is similar to Harsanyi's "purification" argument.

[10] Some papers (e.g., Doval [2016], although in a nontransferable utility framework) consider models in which matching opportunities arrive sequentially; then existence is no longer guaranteed, because an agent may turn down a potential match in order to wait for a better one yet to come. See Atakan (2006) and Lauermann (2013) for related results in the TU case. However, these phenomena cannot appear in our framework.

is equivalent to surplus maximization. It follows that the corresponding equilibrium of our game is efficient, in the (strong) sense that it maximizes aggregate surplus.

Uniqueness is a more difficult issue. Note, first, that the stable match of the auxiliary game is "generically" unique, in the sense that a maximization problem has "in general" a unique solution: while it is always possible to construct situations in which the maximum is reached for different solutions, such cases are in general not robust to small perturbations.[11] This, however, does not imply uniqueness of the Nash equilibrium in the initial game. Indeed, Cole et al. (2001) and Nöldeke and Samuelson (2015) provide examples of "coordination failures," whereby an alternative, Nash equilibrium of the game involves all agents investing in a globally suboptimal but individually rational way.[12] One intuition is that an agent's optimal investment is a (typically increasing) function of the other agents' education. In a context in which the other agents underinvest, a person's best response may well be to underinvest as well, and these best responses may sometimes form a Nash equilibrium. Note, however, that such a scenario is likely when the only benefit from education is perceived on the matching market, much less so in our case, where even singles gain from the investment because of the labor market return.

While the study of coordination failures is an interesting topic, we will not pursue it in the present context. From a theoretical perspective, in the presence of (potentially) multiple equilibria, a natural solution is to use an equilibrium refinement concept. In our case, a natural criterion is Pareto efficiency because, in contrast to most games, one of the equilibria is always Pareto efficient in a strong sense (surplus maximization); in particular, it Pareto dominates any other possible equilibrium. In what follows, we shall therefore concentrate on the ("generically" unique) stable matching of the auxiliary game as the relevant Nash equilibrium. This conceptual choice, however, raises empirical issues that are discussed later on.

IV. Identification of the Distribution of Marital Preferences

The model as presented now requires a distributional assumption on marital preferences for identification of the Pareto weights. However, this can be relaxed if we are willing to allow preferences for marriage to depend

[11] A precise definition of the "genericity" concept invoked in this—admitedly vague—statement would require transversality arguments in functional spaces, which would be well beyond the scope of this paper. See Chiappori et al. (2010) for a detailed analysis.

[12] Nöldeke and Samuelson (2015) provide a set of conditions that are sufficient for uniqueness of ex post equilibria. These conditions, however, are quite restrictive and cannot be expected to hold in our context.

on exogenous variables that do not affect the economic surplus from marriage.

To do this we still assume that marriage generates a surplus, which is the sum of an "economic" component, reflecting the gains arising in marriage from both risk sharing and the presence of a public good, and a nonmonetary term reflecting individual, idiosyncratic preferences for marriage. The economic part is, as before, a deterministic function of the spouses' respective levels of human capital; its distribution between husband and wife is endogenous and is determined by the equilibrium conditions on the marriage market. Regarding the nonmonetary part, however, we assume that the nonmonetary benefit of agent i $(= m, f)$ is the sum of a systematic effect, which depends on some of i's observable characteristics (but not on his spouse's), and of an idiosyncratic term; as before, we assume that the idiosyncratic term, modeled as a random shock, depends only on the human capital of i's spouse. Equation (9) is thus replaced with

$$\Sigma_{mf} = \Sigma\left(H_m, H_f\right) + \left(X_m a^{H_m,H_f} + \beta_m^{H_f} - \beta_m^0\right) + \left(X_f b^{H_m,H_f} + \beta_f^{H_m} - \beta_f^0\right), \tag{23}$$

where X_i is a vector of observable characteristics of agent i. For instance, X_i may include the education levels of i's parents, a possible interpretation being that an individual's preferences for the spouse's human capital are directly affected by the individual's family background. Many alternative interpretations are possible; the crucial assumption here is simply that the marital surplus depends on both X_m and X_f but not on their interaction. Also, note that the coefficients a and b may depend on both spouses' human capital.

In such a setting, one can, under standard, full support assumptions, identify the vectors of parameters a^{H_m,H_f}, b^{H_m,H_f}, and the distribution of $\beta_m^{H_f} - \beta_m^0$ and $\beta_f^{H_m} - \beta_f^0$ (up to the standard normalizations). To see why, note that theorem 1 and corollary 1 can be extended in the following way:

Theorem 2. If the surplus is given by (23), then there exist $2(NJ)^2$ numbers—$\overline{U}_M(H_m, H_f)$ and $\overline{U}_F(H_m, H_f)$ for $(H_m, H_f) \in \mathcal{H}^2$—such that

1. for any (H_m, H_f),
$$\overline{U}_M\left(H_m, H_f\right) + \overline{U}_F\left(H_m, H_f\right) = \Gamma\left(H_m, H_f\right);$$
2. for any m with human capital H_m married to f with human capital H_f,
$$U_m = \overline{U}_M\left(H_m, H_f\right) + X_m a^{H_m,H_f} + \beta_m^{H_f},$$
$$U_f = \overline{U}_F\left(H_m, H_f\right) + X_f b^{H_m,H_f} + \beta_f^{H_m},$$
with the normalization $a^{H_m,0} = b^{0,H_f} = 0$.

Proof. Assume that m and m' have the same human capital H_m, and their respective partners f and f' have the same human capital H_f. Stability requires that

$$U_m + U_f = \Gamma\left(H_m, H_f\right) + X_m a^{H_m, H_f} + \beta_m^{H_f} + X_f b^{H_m, H_f} + \beta_f^{H_m}, \tag{24}$$

$$U_m + U_{f'} \geq \Gamma\left(H_m, H_f\right) + X_m a^{H_m, H_f} + \beta_m^{H_f} + X_{f'} b^{H_m, H_f} + \beta_{f'}^{H_m}, \tag{25}$$

$$U_{m'} + U_{f'} = \Gamma\left(H_m, H_f\right) + X_{m'} a^{H_m H_f} + \beta_{m'}^{H_f} + X_{f'} b^{H_m H_f} + \beta_{f'}^{H_m}, \tag{26}$$

$$U_{m'} + U_f \geq \Gamma\left(H_m, H_f\right) + X_{m'} a^{H_m H_f} + \beta_{m'}^{H_f} + X_f b^{H_m H_f} + \beta_f^{H_m}. \tag{27}$$

Subtracting (24) from (25) and (27) from (26) gives

$$U_{f'} - U_f \geq (X_{f'} - X_f) b^{H_m, H_f} + \beta_{f'}^{H_m} - \beta_f^{H_m} \geq U_{f'} - U_f; \tag{28}$$

hence

$$U_{f'} - U_f = (X_{f'} - X_f) b^{H_m, H_f} + \beta_{f'}^{H_m} - \beta_f^{H_m}.$$

It follows that the difference $U_f - X_f b^{H_m, H_f} - \beta_f^{H_m}$ does not depend on f, that is,

$$U_f - X_f b^{H_m, H_f} - \beta_f^{H_m} = \bar{U}_F(H_m, H_f).$$

The proof for m is identical. QED

As before, an immediate consequence is the following:

Corollary 3.

1. For any man m with human capital H_m, m's spouse at the stable matching has human capital H_f if and only if the following inequalities hold for all $H \in \mathcal{H} \cup \{0\}$:

$$\overline{U}_M(H_m, H_f) + X_m a^{H_m, H_f} + \beta_m^{H_f} \geq \overline{U}_M(H_m, H) + X_m a^{H_m, H} + \beta_m^{H}.$$

2. For any woman f with human capital H_f, f's spouse at the stable matching has human capital H_m if and only if the following inequalities hold for all $H \in \mathcal{H} \cup \{0\}$:

$$\overline{U}_F(H_m, H_f) + X_f b^{H_m, H_f} + \beta_f^{H_m} \geq \overline{U}_F(H, H_f) + X_f b^{H, H_f} + \beta_f^{H}.$$

3. The ex ante expected utility of a man m with human capital H_m is

$$A_M(H_m, X_m) = \mathbb{E}\left[\max_{H_f \in \mathcal{H} \cup \{0\}} \left(\overline{U}_M\left(H_m, H_f\right) + X_m a^{H_m, H_f} + \beta_m^{H_f}\right)\right], \tag{29}$$

and the ex ante expected utility of a woman f with human capital H_f is

$$A_F(H_f, X_f) = \mathbb{E}\left[\max_{H_m \in \mathcal{H} \cup \{0\}} \left(\overline{U}_F\left(H_m, H_f\right) + X_f b^{H_m,H_f} + \beta_f^{H_m}\right)\right], \qquad (30)$$

where the expectation is over the distribution of unobserved preferences for spouses' types, β_m and β_f for men and women, respectively.

It follows that the marital choice of any male m (female f) with human capital H_m (H_f) boils down to a standard, multinomial choice discrete model; the standard identification results apply. However, in this version of the paper we rely on an extreme value distribution for individual utilities and not on covariates.

Beyond this, there are other important aspects of identification because both education and marriage are endogenous in our model. A key identifying assumption is that marriage does not cause changes in wages. In other words, any correlation of wages and marital status is attributed to composition effects. However, education does cause changes in wages, and it is likely that the ability compositions of the various education groups differ: labor market ability is known when educational choices are made in our model. To control for the endogeneity of education we allow the costs of education to depend on residual parental income, when the child was 16, after removing the effects of parental background (see below). The key idea is that children need to be at least in part financed by their parents, and if the latter suffer an adverse liquidity shock, this may inhibit educational attainment.

V. Data

Estimation uses the 18 annual waves (1991–2008) of the British Household Panel Survey, which includes interviews with all household members over 16 and follows them even when they start their own household.

We select two subsamples drawn from the original members of the panel and those added later. The main sample comprises longitudinal information for individuals born between 1951 and 1971 between the ages of 25 and 50.[13] To this sample we add information on the spouses they marry during the observation window. To avoid underestimating marriage rates, those who are not observed past age 30 are dropped from the sample. Overall, the final data set contains information on education, employment, earnings, and family demographics for 4,295 families, 3,046 of which are couples and 629 and 620 are single women and men, respectively. Of these, over 60 percent are observed for at least 5 years. We exclude Northern Ireland as it is surveyed only in the booster samples, during the 2000s. In total,

[13] For couples, we take the reference year of birth to be that of the wife.

the sample size is just over of 41,000 observations.

In the resulting sample, singles are defined as individuals who are never observed married or cohabiting. For the rest, who are classified as married, we use only observations during their first observed marriage.

In estimating educational participation we use parental income observed when the child is 16.[14] This information is available only for individuals who are observed living with their parents at that age. So for this part of the model we use an additional smaller sample of individuals, born between 1973 and 1985, containing parental income information when the young respondent is aged between 16 and 18 and completed education by the age of 23. This sample includes 1,245 individuals, 636 of whom are women.

In the empirical analysis, employment is defined as working at least 5 hours per week. Earnings are measured on a weekly basis. We use the central 5–98 percent of the distribution of pretax real earnings for employees only. Since our model does not deal with macroeconomic growth and fluctuations, we subtract aggregate earnings growth from earnings. Finally, we consider three education levels, corresponding to secondary education (leaving school at 16), high school diploma, and university (college) degree.

VI. Empirical Specification and Estimation

A. *Outline of Estimation*

We estimate the model in three steps. In a first step we estimate the age profiles of weekly earnings (interchangeably referred to as wages) by gender and education. This is done outside the model, on the basis of the control function approach to allow for endogenous selection into work and for the endogeneity of education. For the former we use policy variation in out-of-work income as an instrument. For the latter we use the residual from a regression of parental income when the person was 16 on his or her family background characteristics.

The next two steps take these wage profiles as given and are performed within the model. Since we assume that preferences for work are drawn after the matching stage, we can separately estimate the life cycle model after marriage, exploiting the TU structure (which implies that life cycle labor supply, the public good, and household consumption do not depend on the Pareto weights). Given estimates for preferences and the distribution of unobserved ability, we can then estimate the economic value of marriage for each type of match (by ability and education—which define human capital) and for singles. In a final stage, taking these values as given, we can estimate the preferences that drive marriage, the parame-

[14] We also need family background, since we actually use the residual parental income as we explain in the estimation section.

ters driving the costs of education, and the implied Pareto weights. We now provide details on this procedure as well as our specification.

B. *Earnings Process*

The earnings w_{it} of individual i vary by gender (g), education (s), ability (θ), and age (t). We thus estimate the following earnings equation:

$$\ln w_{it} = \ln W(\theta_i, s, g) + \delta_1^{gs} t_i + \delta_2^{gs} t_i^2 + \delta_3^{gs} t_i^3 + e_{it} + \epsilon_{it}, \quad (31)$$

$$e_{it} = \rho^{gs} e_{it-1} + \xi_{it}, \quad (32)$$

where e is the productivity shock, assumed to follow an AR(1) process with normal innovations ξ_{it} whose variance is $\sigma^2_{\xi,gs}$; ϵ_{it} is an iid shock with variance $\sigma^2_{\epsilon,gs}$ that we interpret as measurement error; and $W(\theta_i, s, g)$ is the market wage faced by an individual of ability type θ_i, gender g, and schooling s. Ability is assumed to follow a distribution with two points of support. While this can be viewed as an approximation from the econometric point of view, it also simplifies the marital matching problem by defining a finite number of individual types.[15]

In a first step we estimate the education and gender-specific age profiles using a control function approach as in Heckman (1979) to allow for endogenous selection into employment and for the endogeneity of education. For this purpose we use a reduced-form binary choice model for employment driven by an index $z_1'\beta_E$. The education reduced form is taken to be an unordered discrete choice among three levels (secondary, high school, and university). This choice is driven by two separate indices, $z_2'\beta_{HS}$ and $z_2'\beta_U$, as in a random utility model with three alternatives; z_1 are the instruments for employment and z_2 are the instruments for education choice. We then use the regression

$$\ln \tilde{w}_{it} = \delta_0^{gs} + \delta_1^{gs} t + \delta_2^{gs} t^2 + \delta_3^{gs} t^3 + \lambda_E(z_1'\beta_E) + \lambda_{ed}(z_2'\beta_{HS}, z_2'\beta_U) + v_{it}, \quad (33)$$

where $\ln \tilde{w}_{it}$ are detrended wages, $\lambda_E(z_1'\beta_E)$ is a control function for employment, and $\lambda_{ed}(z_2'\beta_{HS}, z_2'\beta_U)$ is a control function to account for the endogeneity of education.[16]

We use a probit for employment to estimate the index $z_1'\beta_E$, and we then approximate $\lambda_E(z_1'\beta_E)$ by a quadratic function in the Mills ratio evaluated at the estimated index. As an instrument we use the predicted resid-

[15] Three education groups and two ability types, giving us six classes of individuals for the matching game.

[16] We make the simplifying assumption that we can control for selection into employment by this simple control function without accounting for the dependence of the employment and the education reduced form. This is possible to relax if we estimate the model all in one step, but this is highly time-consuming.

ual out-of-work income. This is a residual from a regression of predicted out-of-work income (based on the welfare system at the relevant time) on household demographic composition and marital status (which are used in the calculation of welfare benefits). Since we have many years of data, we are identifying the impact of out-of-work income on the basis of how it changes for different demographic groups over time. The probit also includes time and age dummies.[17]

For education we use a multinomial logit model for the three levels of education we consider (secondary, high school, and university) to estimate the two indices. We then approximate $\lambda_{ed}(z_2'\beta_{HS}, z_2'\beta_U)$ by a quadratic function in the probability of attending high school and the probability of attending college. As instruments for education we use residual parental income and its square.[18] This is a residual of a regression of parental income when the person was 16 on a set of family background variables. What is left is assumed to reflect a liquidity shock at the time the individuals are making education choices (see Blundell et al. 2016). If individuals are to an extent liquidity constrained when making education decisions, this residual will affect educational outcomes.[19] We assume that the liquidity shock does not affect either wages or preferences in later life, but only the costs of education.

Having estimated the education and gender-specific age profiles, we then take these as known and proceed to estimate the remaining components that determine postmarital behavior, namely, the stochastic process of wages characterized by $\sigma^2_{\xi,gs}$ and ρ^{gs} and the variance of the measurement error σ^2_ϵ, as well as the preferences of leisure and the distributions of unobserved ability and unobserved preferences. This is described in the next section.

C. *Estimating Preferences and the Distribution of Wages*

As already noted, and repeated here for convenience, the period utility function for the household, consisting of a man m and a woman f with education s_m and s_f, respectively, can be written as

[17] Out-of-work income varies over time because of policy changes and over types of individuals. It thus provides important exogenous variation for identifying selection into work. However, our structural model does not account for the tax and welfare system, something that we are intending to do in future. In estimating the age-education profiles outside the model, we are thus able to use policy-induced information that we could not use if we estimated the entire model in one step.

[18] Ideally we would need two instruments or assume that education choices are ordered. Our estimates are almost identical if we assume that choices are ordered.

[19] Family background includes the education of both parents (five levels each), number of siblings and sibling order (dummies for no siblings, three or more siblings, and whether respondent is the first child), books in childhood home (three levels), and whether lived with both parents when aged 16.

$$Q_t\left(C_t + \alpha_{mt}^{Ms_m\gamma} L_{mt} + \alpha_{ft}^{Fs_f\gamma} L_{ft}\right).$$

An important feature, which we exploit in estimating the model, is that household utility and therefore public good consumption and labor supply do not depend on the Pareto weights. To capture the way labor supply varies over the life cycle without having to control explicitly for the presence of children and other important taste shifters over the life cycle, we specify the α parameters to be a polynomial in age (t):

$$\alpha_{it}^{gs\gamma} = \alpha_0^{gs\gamma} + \alpha_1^{gs\gamma} t + \alpha_2^{gs\gamma} t^2 + \alpha_3^{gs\gamma} t^3 + \eta_i + u_{it}, \tag{34}$$

where the parameters $(\alpha_\ell^{gs\gamma}, \ell = 0, \dots, 3)$ are specific to gender (g), education (s), and marital status $(\gamma = 1$ for married and 0 for single). In other words, in our model, preferences for singles and married individuals can differ. The variable η represents unobserved heterogeneity in preferences for working, accounting for persistent differences in labor supply across individuals that are not fully explained by differences in earnings capacity. Individuals draw preferences for work after the matching stage from a distribution that depends on education and has two points of support (although this can easily be relaxed since individuals do not match on this). This assumption allows us to take marital sorting as exogenous for labor supply and to estimate the model for the postmarital choices separately. Finally, u is an iid normal shock, drawn each period.

In general, identification of preferences requires some variables to affect labor supply only through wages. Various strategies are followed in the literature. For example, Blundell et al. (2016) identify their labor supply model by using tax reforms that affect the return to work but not preferences (see also Blundell, Duncan, and Meghir 1998). In this simpler model we do not use this source of exogenous variation, although in principle one could extend our model to allow for taxes and welfare benefits and thus exploit policy changes. This is beyond the scope of the paper, but it is certainly part of our future research program. Here the identification problem is resolved because of the very tight specification of the utility function.

We set the annual discount factor δ to 0.98 and the annual interest rate to 0.015, implying that agents have some degree of impatience.[20] All other parameters are estimated using the method of simulated moments (McFadden 1989; Pakes and Pollard 1989). In our model there are 36 possible types of marital matches and individuals may also be single.[21] For each possible match we simulate wages and labor supply for the entire life cycle and construct several simulated moments that we then match to the equivalent data moments. This provides us with estimates for the joint distribu-

[20] The same discount rate and interest rate were used in Blundell et al. (2016).
[21] Each person can have one of two ability levels and one of three education levels.

tion of unobserved ability in all couple types ($\Pr(\theta_m, \theta_f | s_m, s_f)$ for married couples and $\Pr(\theta_i | s_i)$, $i = m, f$, for singles), estimates of the α parameters, the distribution of preference heterogeneity (η), and the stochastic process of wages (given the preestimated age profiles for each of the three education groups). Critical to our strategy is the fact that we have estimates of the age-education profiles for men and women from the previous step. By constructing wage series that are censored whenever an individual decides not to work (on the basis of the model) and matching the resulting moments to those observed, we control for selection into work when estimating the stochastic process of wages.

From these estimates we can recover the marital sorting patterns by ability and education, as well as the unconditional distribution of ability for men and women. Given this, the marriage market outcomes—$\Pr(H|H_m)$ and $\Pr(H|H_f)$ for all $H \in \mathcal{H} \cup \{0\}$ and each $H_m \in \mathcal{H}$ and $H_f \in \mathcal{H}$—can be recovered by applying a simple conditional probability rule:

$$\begin{aligned}\Pr(H|H_i) &\equiv \Pr(S, \theta | S_i, \theta_i)\\ &= \frac{\Pr(S, S_i, \theta, \theta_i)}{\Pr(S_i, \theta_i)}\\ &= \frac{\Pr(\theta, \theta_i | S, S_i)\Pr(S, S_i)}{\sum_{s \in \mathcal{S} \cup \{0\}} \Pr(\theta_i | s, S_i)\Pr(s, S_i)}\end{aligned}$$

for $i = m, f$ and where $H = H(\theta, S)$. All the quantities in the third line of the expression are either directly observed in the data ($\Pr(s, S_i)$ for all $s \in \mathcal{S} \cup \{0\}$) or estimated from this estimation stage ($\Pr(\theta, \theta_i | S, S_i)$ and $\Pr(\theta_i | s, S_i)$).

Heuristically, identification works as follows: the autocovariance structure of wage growth identifies the stochastic process of wages. The cross-sectional dispersion of wages and their serial dependence that is not explainable by the stochastic process identify the distribution of unobserved heterogeneity in earnings. The age profiles of participation (for each education and gender), given the already estimated age profiles of wages, identify the age effects on labor supply. Finally, since unobserved heterogeneity induces persistence in employment choices, the degrees of individual labor market attachment over 5 years or more, as well as changes with education, identify the distribution of unobserved preference heterogeneity, given the functional forms we choose.[22]

[22] The 328 moments include the means, variances, and several quantiles of the earnings distribution, and the regression coefficients of employment on a quadratic polynomial in age and moments describing the individual-level persistency of employment, measured by the proportion of years working among those observed for at least 5 years, all by education, gender, and marital status. For couples, it also includes quantiles of the joint distribution of earnings. A full list of data and simulated moments together with the diagnostics of the quality of fit can be found in app. D. Appendix C presents the estimated parameters.

D. *Preferences for Marital Sorting and Education*

In our model individuals choose education at a first stage in life and then enter the marriage market. We allow for three levels of educational attainment: secondary (statutory schooling), high school (corresponding to A-levels or equivalent), and university, corresponding to 3-year degrees or above. We interchangeably use the term college for this group. At the point at which they make the education and the matching decisions, ability of all individuals and their preferences for partners are observable by all. Preferences for work are not known.

As discussed earlier, this choice process can lead to many equilibria, one of which is efficient (see Nöldeke and Samuelson 2015). This equilibrium is equivalent to one in which individuals choose education level and type of partner at the same time. We assume that the data are characterized by that equilibrium, and we thus estimate preferences for type of partner and the determinants of education choice in one step. What follows is a discrete choice problem for men and women, respectively, where each chooses one option out of all possible combinations of education and types of spouse. Since we are assuming that the observed patterns correspond to the efficient equilibrium, we can then back out the Pareto weights, which are the prices that decentralize this market.

The value for a woman f with human capital $H_f = H_F(\theta_f, s_f)$ marrying a man m with human capital $H_m = H_M(\theta_m, s_m)$ is the sum of an economic component and a random preference component for a type of spouse (defined by education and ability). In the earlier steps we have estimated the parameters that allow us to compute the economic component for all possible matches defined by the ability and education of each member of the couple, up to the Pareto weight, which we will identify in this step. We can also compute the economic value of being single for all types.

This utility can also be interpreted as the value of choosing both the level of education and type of partner, given own ability, if we net out the costs of education. We define these costs for individual i to be

$$c_i^{gs} = \iota_0^{gs} + \iota_1^{gs} y_i^p + \kappa_i^{gs},$$

where the parameters ι^{gs} are gender and education-specific. We include the residual parental income at 16, y_i^p (described earlier), as a determinant of the costs of education.

The utility for female f with ability θ_f of choosing type of male partner $H = H_M(\theta, s)$ and of own education s_f is given by

$$\begin{aligned}\widetilde{U}_F(\theta_f, s_f, H) = {} & \overline{U}_F(H, H_F(\theta_f, s_f)) + \varphi_F(H_F(\theta_f, s_f))\mathbf{1}(H = 0) \\ & + \tilde{\varphi}_F(\theta_f, |s - s_f|)\mathbf{1}(H_F(\theta_f, s_f) \neq 0) - c_f^{s_f F} + \beta_f^H,\end{aligned}$$

where $\overline{U}_F(H, H_F)$ corresponds to the economic value of marriage and β_f^H is a random preference component for a type of spouse H. To this we have added extra components of marital preferences. Specifically, $\varphi_F(H_F(\theta_f, s_f))$ is a set of six fixed coefficients (one for each type) measuring the noneconomic utility component for remaining single, and $\tilde{\varphi}_F(\theta_f, |s - s_f|)$ is a set of coefficients capturing the (dis)taste for disparity in the educational attainment of spouses. To preempt, these coefficients proved to be important for fitting the sorting patterns in the data. Otherwise the simpler model predicted too little sorting.

The optimal schooling and partner choice is obtained by

$$\left(s_f^*, H^*\right) = \arg\max_{s_f, H}\left(\tilde{U}_F\left(\theta_f, s_f, H\right)\ \forall\ H, s_f\right).$$

Assuming that β_f^H follows an extreme value I distribution, the probability of any observed choice given κ is given by the multinomial logit with 21 alternatives to choose from.[23] To obtain the probabilities that need to be matched with those observed in the data, we integrate out κ, which is assumed to be normally distributed, thus relaxing the distributional assumption and in particular the independence of irrelevant alternatives.

Recall that $\overline{U}_F$ is

$$\overline{U}(H, H_F) = \exp\left\{\frac{1-\delta}{1-\delta^T} V_F(H, H_F, \mu(H, H_F))\right\},$$

where

$$V_F(H, H_F, \mu(H, H_F)) = \Upsilon(H, H_F) + \frac{1-\delta^T}{1-\delta}\ln\frac{\mu(H, H_F)}{1+\mu(H, H_F)},$$

where the μ are such that each individual has a well-defined utility value.[24] The key point is that $\Upsilon(H_M, H_F)$ can be computed on the basis of the estimates from the life cycle estimation stage for each pair of (H_M, H_F). We treat the Pareto weights as unknown parameters, along with the various preference parameters, in the estimation problem.

The 36 Pareto weights, for each possible matched pair of (H_M, H_F), can be fully identified just by using the observed female choices (whom they marry and what education level they choose). However, in equilibrium, we can also use the male choices to identify the same Pareto weights, which provides a set of overidentifying restrictions. This level of overidentification originates from the fact that we can estimate the economic value of marriage from the life cycle problem, which generalizes the Choo and Siow (2006) approach. Using the male choices as well as the female ones

[23] Three levels of own education and a partner with one of two levels of ability and one of three levels of education or remain single: $3 \times (3 \times 2 + 1)$.

[24] Recall that individual utility is in logs, and hence the resulting argument after intrahousehold allocations has to be positive.

is also necessary for estimating the preferences for being single and for marrying a different type of spouse than oneself for both genders.

We obtain the estimates by the method of moments estimator, using simultaneously the male and the female choices, where we match the observed choice probabilities to the equivalent ones implied by the model. In doing so we also minimize the distance between the predicted marital patterns based on the male choices and those based on the female ones, thus finding the parameters that are most consistent with equilibrium. The extent to which the resulting predicted patterns differ from each other is a diagnostic for whether the model can rationalize the observed pattern as an equilibrium in the marriage market.

VII. Results

All estimates relating to the earnings equations, the results on the distribution of ability, as well as the preference parameters determining labor supply choices are presented in appendix C, since they are not of a central interest in themselves. We also present details on the overall model fit.

A. *The Parameters for Marital Sorting and Education Choice*

Table 1 presents the preference parameters governing marital sorting. Preferences for remaining single increase with education. Higher-ability

TABLE 1
Utility Shifters: Preferences for Remaining Single and for Marrying Similarly Educated Spouses

	Men by Ability		Women by Ability	
	Low	High	Low	High
	Preference for Remaining Single, by Education $(\varphi_F(H_F(\theta_f, s_f)))$			
Secondary	1.468	−1.267	−.238	1.700
	(.065)	(.100)	(.096)	(.063)
High school	2.089	−.849	1.048	3.748
	(.067)	(.111)	(.120)	(.092)
University	3.288	−.290	3.573	5.608
	(.082)	(.079)	(.201)	(.135)
	Preference for Differently Educated Spouses $(\tilde{\varphi}_F(\theta_f, \lvert s - s_f \rvert))$			
One educational level difference	.417	−.775	.032	−.143
	(.040)	(.052)	(.057)	(.060)
Two educational levels difference	−2.719	−1.323	−4.318	−.851
	(.052)	(.079)	(.041)	(.108)

Note.—Asymptotic standard errors in parentheses are computed using the bootstrap.

men have a lower preference for being single, while the contrary is true for higher-ability women. The parameters on the lower panel reflect the utility cost of educational disparity within a couple as perceived by each partner: the more disparate the educational levels, the higher is the utility cost, but this differs substantially by gender and ability. One exception to the preference for similarity is lower-ability men who prefer a spouse who is one education level above or below them.

Table 2 shows the costs of education implied by the estimates. A reduction in parental income at age 16 reduces both high school participation and college attendance (increases the cost of education). Although the coefficients for the effect on college are lower, this is no surprise since college attendance takes place 2–3 years later and hence the effect of the shock may have been attenuated by that time.[25]

B. *The Marital Surplus*

We start by ranking individuals by their human capital as measured by their life cycle earnings capacity, which are a function of ability and education. For women, earnings capacity increases with education and ability. However, being a university graduate implies higher earnings, whatever the level of ability. For men, high-ability high school graduates have a higher earnings capacity than lower-ability university graduates. Table 3 reports the ranking of human capital by education and ability and the value of being single. The value of being single increases monotonically with individual human capital for both men and women. However, the increase is much steeper for women, which is part of the reason why single women tend to be drawn from the higher part of the human capital distribution.

Table 4 presents the economic surplus for all 36 possible combinations of human capital for couples. This is the economic value of marriage, over and above the value of remaining single. There are two important conclusions from this. First, the gradient of the surplus is much steeper with respect to female human capital than it is with respect to male. The reason is that the impact of education on female earnings (conditional on employment) and on employment itself is much higher for women than it is for men. We show this in figures 1 and 2.[26] Hence a large part of the variation

[25] Remember that the parental income is a residual, where the effects of family background have been removed. Moreover, we observe family income at the time the child was 16 only for a (younger) subsample of the data. In a richer model it would be desirable to also control for family background, which would affect wages and preferences potentially. However, this would increase the state space and the resulting possible matches beyond the capabilities of our data.

[26] These being earnings, they include hours dimensions as well, which are not modeled here. In interpreting male and female differences it is important to note that many women work part-time, at varying degrees over the life cycle, while men nearly always work full-time.

TABLE 2
UTILITY COST OF EDUCATION

	Men		Women	
	High School	University	High School	University
Constant	1.052	3.916	1.812	6.140
	(.017)	(.025)	(.024)	(.036)
Parental income at 16 (residual)	−.321	−.158	−.238	−.076
	(.115)	(.006)	(.013)	(.004)

NOTE.—Asymptotic standard errors in parentheses are computed using the bootstrap.

in the surplus is explained by the human capital of the woman. Second, the surplus is generally supermodular, particularly for higher levels of human capital. This will push toward positive assortative matching if it were not for preferences for marriage as implied by the random preferences β_i^H. This can be seen by noticing that for most 2×2 submatrices, the sum of diagonal terms exceeds the sum of off-diagonal ones. In particular, all 2×2 matrices at the top of the human capital distribution (i.e., those including the top three levels for each gender) are supermodular. Similarly, all 2×2 submatrices involving human capital levels not immediately adjacent are positive, suggesting that violations of supermodularity, although possible, are mostly "local."

C. *Marital Patterns*

The share of the surplus and the marital patterns drive the choice of partner. The Pareto weights implied by the choices of males are not restricted to be the same as those implied by the observed choices of the females. In equilibrium they should be, but since the model is heavily overidentified, this will in general not be the case in a finite sample even if the restrictions

TABLE 3
HUMAN CAPITAL AND THE VALUE OF BEING SINGLE

	Human Capital Rank					
	1	2	3	4	5	6
	Women					
Education (ability)	Secondary (L)	High school (L)	Secondary (H)	High school (H)	University (L)	University (H)
Value of single	33.4	61.9	88.7	88.7	191.7	293.3
	Men					
Education (ability)	Secondary (L)	High school (L)	Secondary (H)	University (L)	High school (H)	University (H)
Value of single	115.3	150.6	221.2	277.9	301.4	443.2

NOTE.—H = high ability; L = low ability; higher rank corresponds to higher human capital.

TABLE 4
Economic Surplus from Marriage

Men's Education and Ability	Women's Ability and Education: Secondary (L)	High School (L)	Secondary (H)	High School (H)	University (L)	University (H)
Secondary (L)	85.24	149.18	189.63	189.48	197.67	246.00
High school (L)	82.84	144.69	190.01	186.44	200.48	249.98
Secondary (H)	129.88	210.87	267.49	265.53	300.69	371.89
University (L)	101.79	177.37	241.87	233.00	269.38	340.10
High school (H)	139.45	221.56	289.03	281.81	327.76	406.71
University (H)	143.46	235.52	318.10	306.32	367.30	462.56

Note.—Rows and columns are ordered by male and female human capital, respectively. L and H signify low and high ability, respectively.

we impose are true in the population. In a final step of estimation we choose the Pareto weights that minimize the difference in implied marital patterns when comparing male and female choices. The resulting marital patterns implied by the choices of men and women exactly coincide; they are shown in table 5. For nearly 50 percent of married couples both partners have the same level of education; however, there is a substantial number of marriages that do not follow this rule. Hence along the educational

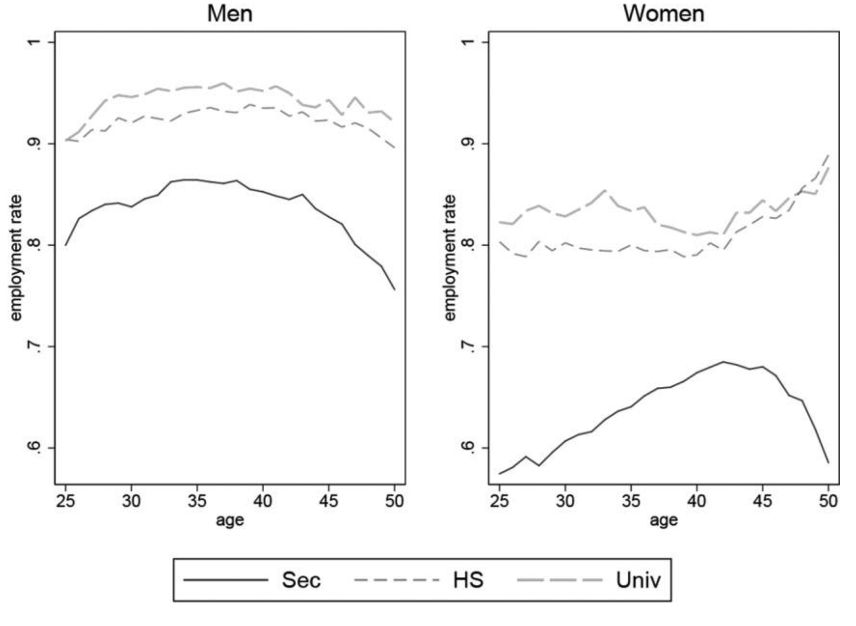

Fig. 1.—Employment of men and women over the life cycle. Color version available as an online enhancement.

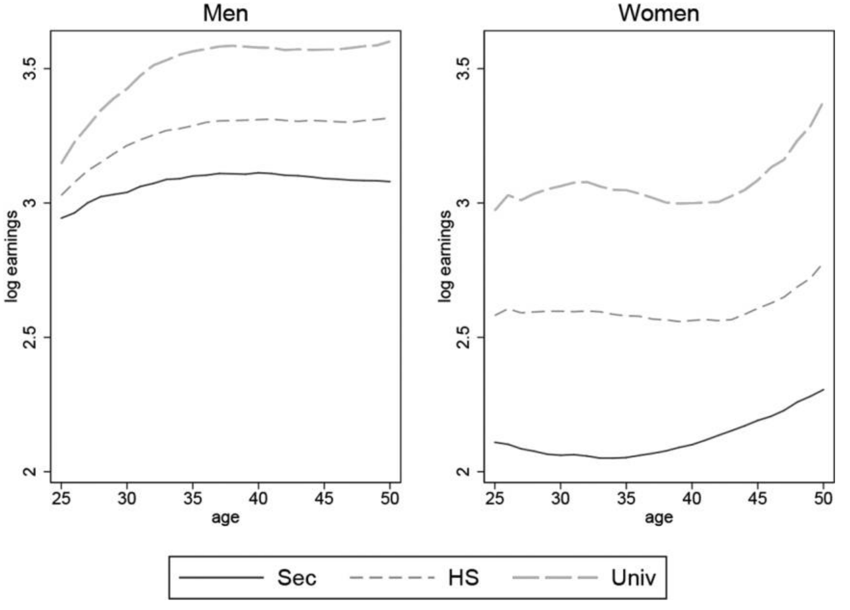

FIG. 2.—Log annual potential earnings for men and women. Color version available as an online enhancement.

dimension the sorting patterns are not perfectly assortative and the model is able to fit this.

Table 6 shows the composition of the singles sorted by their level of human capital:[27] 68 percent of single men are low-ability (human capital levels 1, 2, and 4) compared to about 30 percent in the population. By contrast among single women, only 18 percent are low-ability (human capital levels 1, 2, and 5) compared to about 40 percent in the population. Once ranked by the value of human capital, as measured by potential earnings over the life cycle, we still see that the majority of single women have high human capital and the majority of single men are at the lower end. Finally, the complete set of marital patterns conditional on being married, as implied by our model, are presented in table 7. The matches that actually form will depend on the Pareto weights, and we turn to these now.

D. *Sorting and the Sharing Rule*

The estimated Pareto weights reveal the allocation of welfare within the household in the context of the equilibrium observed in the data. This

[27] Here and in what follows we rank individuals by the level of their potential life cycle earnings, which depends on education and ability. The lowest level is denoted by 1 and increases up to 6.

TABLE 5
Marital Matching Patterns

Men's Education	Women's Education: Secondary	High School	University	Secondary	High School	University
	Simulated Proportions			Data Proportions		
Secondary	.315	.070	.002	.291	.094	.014
High school	.158	.125	.029	.156	.126	.032
University	.009	.048	.052	.019	.044	.053

Note.—The numbers represent cell proportions.

takes into account the public good and the labor supply/leisure decision. Nonmarket time and private consumption are perfect substitutes, while both are complements of public consumption. Table 8 shows the men's share in the gains from marriage that clear the market.[28]

In principle, the relationship between a person's human capital and share of welfare need not be strictly monotonic; the share also reflects relative scarcity of spouses at each level of human capital and therefore depends on the entire distribution. Still, we see that in most (but not all) cases the male share declines in his wife's human capital. Among couples of college graduates with higher ability, the share favors women. Low-skill men marrying the lowest-skill women (col. 1) benefit from very high shares. However, if a low-skill man marries a highest-skill woman (a very rare combination: 0.01 percent of the population), she gets 82 percent of welfare, reflecting a very high Pareto weight for her. Among couples in which the husband is much more skilled than the wife, most of her utility comes from time off work (she is indeed less likely to work), public consumption, and her marital preference.

Risk.—The sorting we observe and the resulting Pareto weights are driven by the structure of the surplus. We have already seen how this varies as a function of human capital. The way it changes across groups is driven both by human capital at the time of matching and by its stochastic properties, since marriage allows risk sharing. In figure 3 we show how the aggregate surplus varies when we change the variance of earnings of men and women by the same factor.[29] The figure shows that as uncertainty rises, the economic value of marriage relative to being single increases because of risk sharing. Halving the variance reduces the aggregate surplus from marriage by 7 percent, and doubling it increases it by 13 percent.

[28] Note that, given supermodularity of the economic component (when this is the case), a marriage between spouses of very different skills signals large values of the corresponding marital preference.

[29] Aggregate surplus is a weighted average across all matches, with weights the probability of a match with baseline risk and conditional on marriage.

TABLE 6
Distribution of Human Capital among Singles

	Level of Human Capital					
	1	2	3	4	5	6
Women	.08	.06	.13	.44	.04	.25
Men	.22	.27	.09	.18	.19	.05

Note.—Levels of human capital in increasing order: (A) men 1: secondary low ability; 2: high school low ability; 3: secondary high ability; 4: university low ability; 5: high school high ability; 6: university high ability; (B) women 1: secondary low ability; 2: high school low ability; 3: secondary high ability; 4: high school high ability; 5: university low ability; 6: university high ability.

TABLE 7
Sorting Patterns Conditional on Marriage by Education and Ability

	Women's Education and Ability					
Men's Education and Ability	Secondary (L)	High School (L)	Secondary (H)	High School (H)	University (L)	University (H)
Secondary (L)	.033	.026	.043	.015	.000	.000
High school (L)	.031	.009	.035	.006	.004	.004
Secondary (H)	.134	.030	.181	.015	.001	.001
University (L)	.000	.011	.000	.007	.003	.004
High school (H)	.054	.080	.075	.061	.013	.014
University (H)	.000	.025	.011	.016	.025	.032

Note.—Rows and columns are ordered by male and female human capital, respectively. L and H signify low and high ability, respectively. Cell proportions are reported.

TABLE 8
Sharing Rule

	Women's Ability and Education					
Men's Education and Ability	Secondary (L)	High School (L)	Secondary (H)	High School (H)	University (L)	University (H)
Secondary (L)	.796	.374	.539	.184	.273	.184
	(.030)	(.020)	(.017)	(.018)	(.014)	(.014)
High school (L)	.905	.592	.596	.406	.069	.050
	(.024)	(.021)	(.015)	(.017)	(.011)	(.013)
Secondary (H)	.580	.461	.448	.302	.141	.128
	(.011)	(.009)	(.008)	(.007)	(.049)	(.027)
University (L)	.933	.855	.943	.668	.452	.372
	(.013)	(.011)	(.017)	(.012)	(.010)	(.010)
High school (H)	.761	.481	.585	.368	.244	.213
	(.008)	(.008)	(.007)	(.006)	(.004)	(.005)
University (H)	.703	.762	.743	.624	.416	.372
	(.014)	(.005)	(.009)	(.005)	(.004)	(.004)

Note.—Male share of surplus. Asymptotic standard errors in parentheses are computed using the bootstrap. Cells are ordered by male and female human capital, respectively. L and H signify low and high ability, respectively.

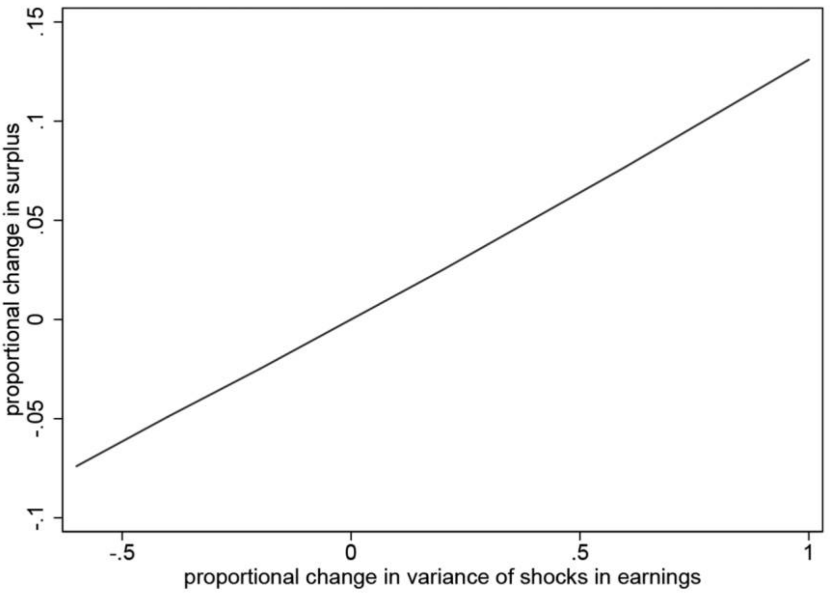

FIG. 3.—The impact of risk on the marital surplus. Color version available as an online enhancement.

In figure 4 we then show the resulting changes in education and marriage rates as risk changes. Increases in risk increase the proportion of college graduates for both men and women and of high school for the latter. Marriage rates also go up in response to the increased gains from marriage. However, the marriage rates increase most for higher education groups for men, while for women the marriage rates increase most among high school graduates.

Marital returns to education.—Finally, it is crucial to keep in mind that the Pareto weights, and more generally the patterns of intrahousehold distribution of resources and welfare, are not structural parameters but endogenous entities reflecting the conditions in the marriage market. The present estimations reflect the patterns we see in the data. In what follows we carry out a counterfactual simulation that will yield new Pareto weights and marital patterns.

One of the key points of our approach is that part of the returns to education can be accounted for by marriage and in particular by the sharing of the marital gains. Thus, ignoring the preferences for marital status, marital returns to high school account for 46 percent of the entire return to high school for women, assuming optimal choice of partner. Marital returns account for 58 percent of the female college premium. Both these

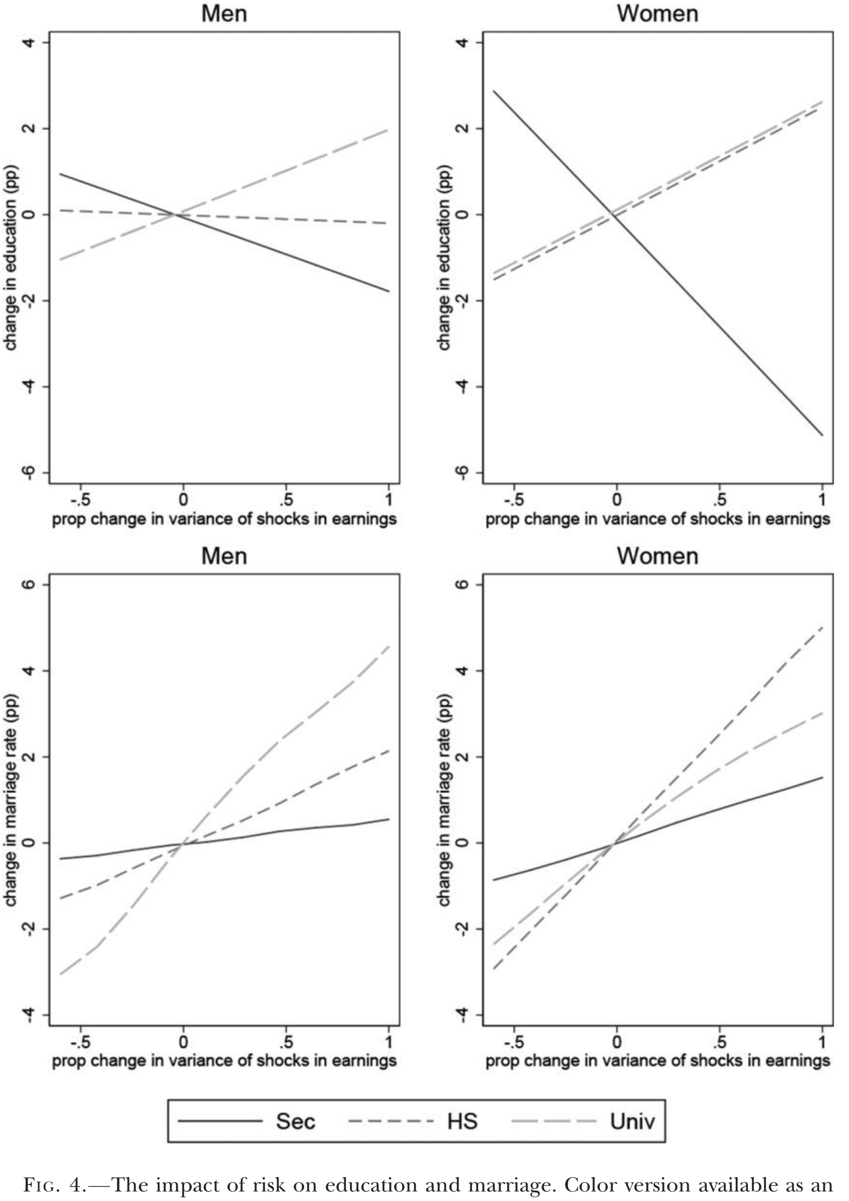

FIG. 4.—The impact of risk on education and marriage. Color version available as an online enhancement.

numbers demonstrate the importance of marriage in determining the returns to education for women. For men there is a similar impact, but it is smaller: the respective effects are 23 percent for high school and 44 percent for college.

E. *Uniqueness: A Brief Discussion*

As indicated above, uniqueness of the Nash equilibrium is not formally guaranteed. Besides the theoretical aspects, this fact might affect our estimation strategy, since the latter assumes efficiency. If, in the real world, individuals happen to miscoordinate on some inefficient equilibrium, our estimates may be biased, since we will recover the parameters for which the observed decisions are efficient. In practice, this would mean that (i) the real benefits of education are (much) larger than what we find, but (ii) agents are stuck in an inefficient equilibrium in which they fail to fully exploit these benefits. This possibility is probably mitigated by two remarks. First, there is little evidence of massive underinvestment in education. Second, from a more theoretical perspective, the Nöldeke-Samuelson approach clearly suggests that coordination failures are particularly likely in situations in which the only benefit of the investment is perceived on the matching market; however, our estimates clearly indicate that even singles gain from education—these gains being reaped in the labor market. In particular, a no-investment outcome is a theoretical possibility in the general model but cannot happen in our context because of the size of the labor market returns. Another theoretical possibility would be that all agents choose their autarchy investments and remain unmatched; but this is precluded in our context since marriage is always beneficial in economic terms, because of both risk sharing and the presence of a public good.

Still, the possibility of multiple equilibria cannot be formally excluded. In order to further investigate this issue, we investigated whether our model evaluated at the estimated parameters implied the existence of other equilibria. We start from a situation in which all agents are at the minimum education level, derive the resulting labor supply and matching patterns, compute the corresponding returns on education (including the marital returns), solve for optimal education choices given these returns, and iterate. We failed to find an alternative equilibrium; in particular, the process always converges to a distribution of education similar to the one observed. We then address a further possible concern raised by this approach: if the observed distribution of education was in fact generated by an inefficient equilibrium, our assumption of efficiency would bias the costs of education upward in order to be able to "justify" the observed education patterns as the outcome of a social optimum. We repeated the exercise for lower education costs, similar to those we use in the counterfactual simulations described below. Again, we find a unique equilibrium, described in the next section. Our conclusion, thus, is that while the possibility of multiple equilibria cannot be formally eliminated, its relevance for our results is unlikely to be high.[30]

[30] An alternative approach would rely on a matching maximum score estimator approach à la Fox (2016), although it would require a different stochastic structure.

VIII. Counterfactual Simulations

The model offers us a way of interpreting the data as well as the possibility of counterfactual analysis with an emphasis on longer-run outcomes. Here we examine the impact of reducing the costs of university education by 10 percent. The key mechanism that can cause the realignments in the marriage market and indeed change the welfare of men and women is the increased supply of college graduates of both genders, which will affect the types of individuals that enter the marriage market. The changes in the implied Pareto weights will then feed back into the education choice.

The final equilibrium distribution of education shows an increase in the supply of both male and female college graduates (see table 9). These supply changes are associated with both changes in the matching patterns and the welfare share for each type of match. The former are shown in tables 10 and 11 and the latter in table 12. Generally there is an increase in the proportion of college-educated women (men) marrying non-college-educated men (women), as well as an increase in matches among college-educated men and women. As shown in table 11, the proportion of singles increases.

Table 12 reveals very interesting changes in the way welfare is distributed within the household. A negative value means a decrease in the male share in favor of the female one. These results imply that subsidizing women to increase college attendance can increase their share of welfare, particularly for the lower-skill ones marrying high-skill men. For example, the welfare share of low-ability male college graduates marrying low-ability high school graduates declines by 5 percentage points. Female college graduates, however, do not necessarily gain themselves: the share of college graduates declines. These patterns are driven by the change in the relative scarcity of partners at each skill level.

Underlying these results is the convergence to a new long-run equilibrium, with changes in educational attainment relative to the immediate effect induced by the subsidy, as the marital patterns change and the Pareto weights adjust. In future research it will be important to examine how

TABLE 9
EDUCATION DISTRIBUTION

	Men		Women	
	Baseline	Low-Cost University	Baseline	Low-Cost University
Secondary	.448	.403	.522	.472
High school	.400	.371	.341	.314
University	.152	.227	.137	.214

TABLE 10
CHANGES IN THE MATCHING PATTERNS

MEN'S EDUCATION AND ABILITY	WOMEN'S EDUCATION AND ABILITY					
	Secondary (L)	High School (L)	Secondary (H)	High School (H)	University (L)	University (H)
Secondary (L)	−.18	−.21	−.32	−.12	.00	.00
High school (L)	−.12	−.07	−.18	−.06	.11	.05
Secondary (H)	−.66	−.22	−1.30	−.14	.04	.01
University (L)	.00	.18	−.01	.14	.16	.20
High school (H)	−.28	−.44	−.52	−.42	.26	.27
University (H)	.00	.53	.21	.34	1.26	1.51

NOTE.—Numbers correspond to changes in the proportion of each cell. Cells are ordered by male and female human capital, respectively. L and H signify low and high ability, respectively

TABLE 11
MARITAL PATTERNS

	MEN		WOMEN	
	Baseline	Low-Cost University	Baseline	Low-Cost University
Proportion remaining single	.185	.194	.202	.221
Proportion marrying equally educated spouse	.495	.486	.489	.477

TABLE 12
CHANGES IN THE SHARING RULE: PERCENTAGE POINTS

MEN'S EDUCATION AND ABILITY	WOMEN'S EDUCATION AND ABILITY					
	Secondary (L)	High School (L)	Secondary (H)	High School (H)	University (L)	University (H)
Secondary (L)	1.2	1.3	.4	2.3	−.6	2.1
High school (L)	−.4	−.1	.0	1.5	3.4	4.5
Secondary (H)	.4	1.0	.1	1.6	6.1	3.3
University (L)	.2	−5.4	.3	−2.4	−1.3	1.0
High school (H)	.1	.2	−.1	.3	4.1	2.9
University (H)	−2.0	−3.7	−3.6	−3.1	.4	−.1

NOTE.—Cells are ordered by male and female human capital, respectively. L and H signify low and high ability, respectively. A negative number corresponds to a decline in the male share.

changes in welfare benefits and their targeting affect marital patterns and life cycle work and consumption decisions.

IX. Concluding Remarks

In this paper we have presented an equilibrium model of education choice, marriage, life cycle labor supply, savings, and public goods in a world with uncertainty in the labor market. Our framework relies on a transferable utility setting, which allows us to simulate policies that change the economic environment at any stage of the life cycle. Matching in the marriage market is stochastic but frictionless and trades off the economic value of marriage with random preferences for type of mate (defined by their human capital). On the economic side, the final structure of matching is driven both from the demand for public goods and from a risk-sharing motive.

We find that the surplus from marriage is supermodular nearly everywhere, pushing toward positive assortative matching, with any departures from perfect sorting being driven by random preferences for mates and by some local departure from supermodularity. We also find that the human capital of women is a very strong determinant of marital surplus, more so than the human capital of men. The model is able to replicate matching patterns very well, despite the fact we do not allow for frictions.

Generally high–human capital women get more than half of the marital surplus, while men marrying low–human capital women get most of the surplus. These shares reflect the existing equilibrium in the data. However, the share of welfare is endogenous, and changes in the supply of men and women of different levels of human capital can change them. Thus in our counterfactual simulation, where we reduce the costs of education, inducing more to graduate from college, we find that the share of low–human capital women increases, while the share of low-ability college graduate women declines.

Finally, our model sheds light on the determinants of human capital investments. Two conclusions emerge. First, noneconomic factors play an important role in both the decision to marry and the marital patterns conditional on marriage and therefore indirectly affect the return to education. This is by no means surprising. However, our approach allows us to quantify the magnitude of these effects; we find them to be quite large. Second, and more importantly, we can explicitly decompose the returns to education into those perceived on the labor market and those reaped in addition on the marriage market (the "marital college premium"). Our results are unambiguous: the benefits perceived through marriage (through risk sharing and the joint consumption of public goods) are dominant. Our analysis, therefore, confirms the notion, put forth by Chiappori et al. (2009), that any empirical analysis that omits marital gains and concentrates exclusively on the labor market may be severely biased. An important implication is that a policy (e.g., a tax reform) that directly affects the returns on human capital investment will also alter the respective importance of economic and noneconomic factors for the determination of matching patterns, further influencing incentives to invest; these equilibrium effects will typically amplify the initial impact, resulting in potentially large long-term consequences that should not be ignored.

This paper is a first step toward a rich research agenda analyzing the interactions of marriage, labor markets, and educational choices. Generalizations will include allowing for imperfectly transferable utility, generalizing the model to allow for divorce, and finally allowing for limited commitment.[31] These are important issues that will lead to better understanding of marriage markets and intrahousehold inequality. However, they are also challenging. Our framework here shows, however, that such equilibrium models can be rich in implications and valuable for the understanding of the longer-term effects of policies.

Finally, the framework developed in this model, complicated as it may be, relies on two simple but extremely powerful insights. One is that marital sorting patterns—who marries whom—have an important, economic component, which can be analyzed in terms of "complementarity" or "substitutability" (in technical terms, super- or submodularity) of the surplus created within marriage; the second is that the intrahousehold allocation of resources (therefore of welfare) is related to the equilibrium conditions prevailing on the "marriage market" and should therefore be analyzed using the "theory of optimal assignments" (also know as matching models). Both insights are explicitly present in Becker's (1974) *JPE* masterpiece. That, more than 40 years later, we can still find much to learn in exploiting these insights is an obvious tribute to the importance of Becker's contribution.

References

Atakan, Alp. 2006. "Assortative Matching with Explicit Search Costs." *Econometrica* 74 (3): 667–81.

Attanasio, Orazio, Hamish Low, and Virginia Sánchez-Marcos. 2008. "Explaining Changes in Female Labor Supply in a Life-Cycle Model." *A.E.R.* 98, no. 4 (September): 1517–52.

Becker, Gary S. 1974. "A Theory of Social Interactions." *J.P.E.* 82, no. 6 (November/December): 1063–93.

[31] Regarding imperfectly transferable utility, see Galichon, Kominers, and Weber (2016) for a recent contribution along these lines. A recent contribution by Voena (2015) models divorce under limited commitment to study the impact of divorce laws on the choices of couples.

———. 1991. *A Treatise on the Family*. Cambridge, MA: Harvard Univ. Press.

———. 1994. *Human Capital: A Theoretical and Empirical Analysis, with Special Reference to Education*. 3rd ed. Chicago: Univ. Chicago Press.

Blundell, Richard W., Pierre-André Chiappori, and Costas Meghir. 2005. "Collective Labor Supply with Children." *J.P.E.* 113, no. 6 (December): 1277–1306.

Blundell, Richard W., Monica Costa Dias, Costas Meghir, and Jonathan Shaw. 2016. "Female Labor Supply, Human Capital, and Welfare Reform." *Econometrica* 84 (5): 1705–53.

Blundell, Richard, Alan Duncan, and Costas Meghir. 1998. "Estimating Labor Supply Responses Using Tax Reforms." *Econometrica* 66 (4): 827–61.

Chiappori, Pierre-André. 1988. "Rational Household Labor Supply." *Econometrica* 56, no. 1 (January): 63–90.

———. 1992. "Collective Labor Supply and Welfare." *J.P.E.* 100, no. 3 (June): 437–67.

Chiappori, Pierre-André, Murat Iyigun, and Yoram Weiss. 2009. "Investment in Schooling and the Marriage Market." *A.E.R.* 99 (5): 1689–1713.

Chiappori, Pierre-André, Robert McCann, and Lars Nesheim. 2010. "Hedonic Price Equilibria, Stable Matching, and Optimal Transport: Equivalence, Topology, and Uniqueness." *Econ. Theory* 42 (2): 317–54.

Chiappori, Pierre-André, and Bernard Salanié. 2016. "The Econometrics of Matching Models." *J. Econ. Literature* 54 (3): 832–61.

Chiappori, Pierre-André, Bernard Salanié, and Yoram Weiss. 2017. "Partner Choice, Investment in Children, and the Marital College Premium." *A.E.R.* 107 (8): 2109–67.

Choo, Eugene, and Aloysius Siow. 2006. "Who Marries Whom and Why." *J.P.E.* 114, no. 1 (February): 175–201.

Cole, H. L., G. J. Mailath, and A. Postlewaite. 2001. "Efficient Non-contractible Investments in Large Economies." *J. Econ. Theory* 101 (2): 333–73.

Doval, Laura. 2016. "A Theory of Stability in Dynamic Matching Markets." Working paper, Yale Univ.

Fox, J. T. 2016. "Estimating Matching Games with Transfers." Working paper, Rice Univ.

Galichon, Alfred, Scott Duke Kominers, and Simon Weber. 2016. "Costly Concessions: An Empirical Framework for Matching with Imperfectly Transferable Utility." Manuscript, Soc. Sci. Res. Network. http://dx.doi.org/10.2139/ssrn.2535394.

Goussé, Marion, Nicolas Jacquemet, and Jean-Marc Robin. 2016. "Marriage, Labor Supply, and Home Production: A Longitudinal Microeconomic Analysis of Marriage, Intra-household Bargaining and Time Use Using the BHPS, 1991–2008." Working Paper no. 16-01, CIRPEE, Quebec.

Graham, B. 2011. "Econometric Methods for the Analysis of Assignment Problems in the Presence of Complementarity and Social Spillovers." In *Handbook of Social Economics*, vol. 1B, edited by J. Benhabib, M. Jackson, and A. Bisin, 965–1052. Amsterdam: North-Holland.

Heckman, James J. 1979. "Sample Selection Bias as a Specification Error." *Econometrica* 47 (1): 153–61.

Hoppe, H., B. Moldovanu, and A. Sela. 2009. "The Theory of Assortative Matching Based on Costly Signals." *Rev. Econ. Studies* 76 (1): 253–81.

Lauermann, Stephan. 2013. "Dynamic Matching and Bargaining Games: A General Approach." *A.E.R.* 103 (2): 663–89.

Low, Hamish, Costas Meghir, and Luigi Pistaferri. 2010. "Wage Risk and Employment Risk over the Life Cycle." *A.E.R.* 100 (September): 1432–67

Low, Hamish, Costas Meghir, Luigi Pistaferri, and Alessandra Voena. 2018. “Marriage, Labor Supply and the Dynamics of the Social Safety Net.” Working Paper no. 24356, NBER, Cambridge, MA.

Mazzocco, M. 2007. “Household Inter-temporal Behaviour: A Collective Characterization and a Test of Commitment.” *Rev. Econ. Studies* 74, no. 3 (July): 857–95.

McFadden, D. 1989. “A Method of Simulated Moments for Estimation of Discrete Response Models without Numerical Integration.” *Econometrica* 57 (5): 995–1026.

Nöldeke, G., and L. Samuelson. 2015. “Investment and Competitive Matching.” *Econometrica* 83 (3): 835–96.

Pakes, A., and D. Pollard. 1989. “Simulation and the Asymptotics of Optimization Estimators.” *Econometrica* 57 (5): 1027–57.

Peters, Michael, and Aloysius Siow. 2002. “Competing Premarital Investments.” *J.P.E.* 110 (3): 592–608.

Schulhofer-Wohl, Sam. 2006. “Negative Assortative Matching of Risk-Averse Agents with Transferable Expected Utility.” *Econ. Letters* 92 (3): 383–88.

Shimer, Robert, and Lones Smith. 2000. “Assortative Matching and Search.” *Econometrica* 68, no. 2 (March): 343–69.

Townsend, Robert M. 1994. “Risk and Insurance in Village India.” *Econometrica* 62 (3): 539–91.

Villani, C. 2009. *Optimal Transport: Old and New*. Grundlehren der Mathematischen Wissenschaften, vol. 338. New York: Springer.

Voena, A. 2015. “Yours, Mine and Ours: Do Divorce Laws Affect the Intertemporal Behavior of Married Couples?” *A.E.R.* 105 (8): 2295–2332.