The changing wage distribution and the decline of marriage*

Edoardo Ciscato†

January 21, 2021

Abstract

In this paper, we introduce a search-and-matching framework to study educational choices, and household formation and dissolution over the life-cycle. In this model, sorting patterns on the marriage market are driven by both economies of scale in household production and consumption, and search frictions. We discuss the identification of both objects with panel data on marriages and divorces. We estimate the model with PSID data and document how the gains from marriage, marriage market segmentation, and returns to college have changed in the U.S. over the last 50 years. Through our structural model, we show that changes in the wage distribution alone would lead to an increase of college education equally for both genders, could explain almost 20% of the total decline of marriage, but would also reduce women's marriage market returns to college. On the other hand, we show that changes in household production, gender norms, and search frictions have given female college graduates an edge on marriage markets. The positive interaction between the increasing college wage premium and the changes in marriage market fundamentals explains the fast rise of female college education.

Keywords: marriage markets, divorce, remarriage, human capital investment, search and matching.

JEL Classification: D13, J11, J12.

---

*I thank Pierre-André Chiappori, Alfred Galichon and Jean-Marc Robin for their guidance. I am grateful to Laurens Cherchye, Arnaud Dupuy, Chris Flinn, Marion Goussé, Maurizio Mazzocco, Costas Meghir, Frederic Vermeulen for their help and advice. I also thank all participants at various seminars and conferences. I thank the Doctoral School of Sciences Po and the Alliance Program for funding provided. This draft is based on a working paper published in the HCEO Working Group series (Ciscato, 2018). All errors are my own.

†KU Leuven. Email: edoardo.ciscato@kuleuven.be. Website: https://sites.google.com/site/ciscatoedoardo.

1. Introduction

Over the last fifty years, the share of single households has increased greatly in the U.S. While individuals have been further delaying household formation, separation risk has increased for both married and cohabiting couples, with a particularly strong increase for individuals with a low human capital (Lundberg et al., 2016). These trends have a number of important implications: firstly, low-wage individuals are less likely to rely on a partner that can help them achieve higher consumption levels, cope with income risk, and raise children (Blundell et al., 2016a,b). Secondly, rising wage inequality has resulted in a sharper divide between unstable families with little human capital and stable families with high human capital (Fernández and Rogerson, 2001; Fernández et al., 2005). Thirdly, the incentives to invest in college education have changed, as the association between college education and better marriage prospects is stronger than in the past (Bronson, 2014; Chiappori et al., 2017).

In this paper, we introduce a unified framework to empirically study the determinants of education, marriage, and divorce decisions over the life-cycle. Given the increasing marital instability and heterogeneity in marital histories observed in the data, it is increasingly difficult to represent the marriage market outcome as a one-time assignment as originally suggested by Becker (1973). Hence, this paper presents a new framework to study both cross-sectional (*whether and whom people marry*) and longitudinal aspects of matching behavior over the life-cycle (*when people marry, whether they divorce and remarry*).

The first goal of the paper is to empirically assess the determinants of dynamic matching behavior, and document how they have changed over the last five decades in order to better understand household formation and dissolution trends. As per Becker (1973) and Choo and Siow (2006), we interpret the marriage patterns observed in the data as the outcome of a competitive matching game. However, we allow for two different, nonexclusive driving forces behind sorting patterns on marriage markets. The first is the presence of economies of scale in consumption and household production, while the second is the presence of frictions. By pooling their resources together, couples can expand their consumption frontier and insure each other against income risk; while we do not explicitly model fertility, joining efforts in raising children can be thought of as a key component of such gains. Hence, changes in household technology, as well as time-use preferences and norms about family roles alter the incentives behind the choice of whether and whom to marry.[1] However, positive, or negative, assortative mating can

[1] Greenwood et al. (2005, 2016) argue that the increased availability of cheap and productive home appliances (e.g., dishwashers, washing machines, microwave ovens) induced a major change in both the gains from marriage and female labor supply. Chiappori et al. (2017) point at the increasing importance of the complementarity between parents' human capital in childrearing in order to explain increased assortative mating.

also be explained by the presence of frictions, i.e., individuals with similar characteristics might have higher chances to meet and a lower cost of learning about each other. The economic literature has often overlooked search frictions as a source of assortativeness,[2] although it is natural to think that marriage markets are, to some extent, segmented along demographic and socioeconomic observables.

While we think of the size and nature of the economies of scale and frictions as the primitive parameters of our model, matching decisions are made in a dynamic competitive environment and also depend on forward-looking agents' expectations about both their current match and their (re)marriage prospects. In a framework where partners cannot commit to an unproductive relationship, they need to consider if present-day economies of scale are large enough to compensate possibly meager expectations about future household income, as well as if they should instead search for a better march (Becker et al., 1977; Burdett and Coles, 1997; Shimer and Smith, 2000). Understandably, age and wage expectations can play an important role in shaping these trade-offs. Hence, we disentangle between static and dynamic components of the gains from marriage and discuss how they differ across population groups.

We estimate our model with data from the Panel Survey of Income Dynamics and show that search frictions are weaker among likes; single agents with similar education, ages, family backgrounds, and wages are more likely to meet on marriage markets. Hence, search frictions can partly explain positive assortative mating. At the same time, we find wages to be substitutes in household production, so that couples have incentives to specialize, as suggested by Becker (1973). However, households also benefit from insuring themselves against future income shocks; couples formed by two high-earners are less exposed to labor income risk as wage shocks are persistent. This said, we also find evidence of gender asymmetries in household production: *who* the primary earner is matters, and the husband's wage matters more than the wife's in determining the household economies of scale. We find gender roles in household production to be strong in the 1970s and weaker, but still present, in the 2010s.

The second goal of the paper is to leverage on our dynamic matching framework to empirically study both labor and marriage market returns to college, as well as track their evolution over time. We model the decision of going to college so that young agents take into account

[2]Economists did study the role of frictions as a determinant of marital turnover both theoretically (Burdett and Coles, 1997; Shimer and Smith, 2000) and empirically (Goussé et al., 2017). Few papers try to address whether meetings are assortative: Hitsch et al. (2010) and Belot and Francesconi (2013) discuss the empirical patterns of dating behavior using data from an online dating site and a speed dating agency, respectively. Sociologists and demographers put greater emphasis on market segmentation as a determinant of homogamy (Schwartz, 2013; Gonalons-Pons and Schwartz, 2017). Rosenfeld and Thomas (2012) analyze where and how couples meet in the U.S. with data from a nationally representative survey.

both the associated wage premium and the potential benefits of higher education on their future marriage prospects. Acquiring college education has far-reaching implications on future marriage market outcomes. Education has a direct impact on the economies of scale and search frictions, thus playing an important role in partner selection. However, it also has an indirect impact on marriage market outcomes as it influences all future wage realizations.

In our empirical analysis, we document ex-ante welfare changes for young agents between the 1970s and the 2010s, and define the returns to college for 17-year-old agents as the excess value from expected life-cycle consumption associated with going to college. We decompose them into a labor market and a marriage market component; the latter accounts for 19pp out of a total 82pp increase for men, and for 24pp out of a total 138pp increase for women.

In our model, lower match rates and higher separation rates result in welfare losses due to foregone economies of scale only achievable through marriage. Men without a college degree have experienced the largest losses as their worse labor market outlook compounds with poorer marriage prospects and increased marital instability. Women without a college degree have benefited from modest wage improvements and are thus better equipped to deal with singlehood. However, these gains are not large enough to compensate for the welfare losses they experience on marriage markets. Such women tend to match with low-wage men due to marriage markets being segmented along socioeconomic observables; their low-wage husbands are no longer the undisputed primary earners, so gains from specialization are weak. However, the gains from mutual insurance are also weak, because of the sizable threat of uninsurable negative shocks that would completely erase the match surplus. In contrast, college graduates start a family later, but most of them eventually engage in long-lasting relationships even in more recent periods. We show that, contrary to the 1970s, women's investments in human capital have started to pay off on marriage markets in the form of more stable matches and increased consumption opportunities for their families. As gender roles within the household have become more symmetric, high-wage couples are less specialized than in the past, but they benefit from the presence of two high earners able to sustain household consumption in case of a shock.

Finally, we ask the following counterfactual question: How would people marry, divorce and invest in college education if the only thing that changed since the 1970s were their full-time wages? We compute the counterfactual steady-state equilibrium of our model corresponding to this scenario and show that the increase in college wage premia between the 1970s and the 2010s results in a rise in the share of college graduates of about 20pp for both genders. When we look at marriage market outcomes, we find that changes in wages lead to a 2.4pp increase in the number of single adults, corresponding to about 17% of the total increase observed between the 1970s and the 2010s. In a counterfactual world where marriage is still based on "traditional" household specialization, the high wages of female college graduates constitute a hindrance to marriage, while non-graduate men have lost an edge and are less fit to be the primary earner in the household as their wages have decreased; as a result, the overall number of matches declines. However, changes in the wage distribution alone cannot explain either the faster rise of college education among women, or the current better marriage market prospects of higher-educated women. In fact, when considered in isolation, changes in wages cause a decline in the marriage market returns to college. On the other hand, the factual changes in household production, gender norms, and search frictions have given female college graduates an edge on marriage markets. When these changes are taken into account, a higher human capital is no longer a hindrance to women's family lives; on the contrary, it pays off on both labor and marriage markets.

1.1. Relationship to the existing literature

This paper contributes to three different strands of literature. First, we build a novel search-and-matching model of the marriage market equilibrium based on Shimer and Smith (2000) and Chade and Ventura (2002). Similarly to Goussé et al. (2017), henceforth GJR, we introduce a random match quality component in the gains from marriage. In the empirical analysis, this helps us rationalize why observationally identical pairs do not always match. As in GJR, we let couples bargain over consumption and time use. The intrahousehold allocation crucially depends on both partners' (re)marriage prospects, which in turn depend on the number and characteristics of the available singles in equilibrium.[3] Changes in the quality of the match can also explain why some couples break up. Since partners can transfer utility to each other at no cost, separations occur when the match surplus turns negative and both partners are better-off living alone (Becker et al., 1977).

This paper expands the previous work of GJR in several directions. Firstly, we introduce aging in order to analyze the timing of household formation and dissolution over the life-cycle, the degree of marriage market segmentation into different cohorts, and the patterns of age hypergamy. Coles and Francesconi (2011), Díaz-Giménez and Giolito (2013), and Shephard (2018) are related papers that study gender-asymmetric incentives to wait before marriage; while the first two papers do not discuss the issues of divorce and remarriage, the work of

[3]The characterization of the distribution of singles as an equilibrium object differentiates our setting from dynamic discrete choice models where the pool of singles is treated as exogenous. Several papers in this literature explicitly model marriage and divorce decisions while studying savings and dynamic labor supply decisions. A precursor is Van der Klaauw (1996), while more recent examples are Sheran (2007), Keane and Wolpin (2010), Bronson (2014) and Fernández and Wong (2017). Beauchamp et al. (2018) take an intermediary approach, where the distribution of singles is an endogenous equilibrium object, but marriage across different generational cohorts is ruled out.

Shephard (2018) is the most closely related to ours, and differences are discussed in this section and throughout the paper. Secondly, we introduce wage shocks, so that mutual insurance becomes an important motive for marriage; yet, shocks can cause couples to break up if insuring the partner becomes too costly. More specifically, we assume there is no commitment to the initial intrahousehold allocation and couples are free to adjust it every time a wage shock hits. Hence, the mutual insurance coverage is below the ex-ante efficient level characterized by Ligon et al. (2002) and Mazzocco (2004).[4] Thirdly, we let education be the result of a human capital investment decision based on both labor and marriage market gains. In this regard, this paper differs from previous works on the interplay between marriage and education decisions by Chiappori et al. (2017, 2018a), as we allow agents to get (re)married and divorced at any stage of their life-cycle. In contrast to Guvenen and Rendall (2015) and Greenwood et al. (2016), we explicitly model aging and let age be an important determinant of both the wage process, the gains from marriage and search frictions. Our setting also differs from that of Shephard (2018), who embeds a dynamic labor supply framework in his search-and-matching model of the marriage market in order to study on-the-job human capital accumulation, but does not model educational choices.

The second contribution of the paper is a formal discussion on how to jointly identify and estimate match gains and market frictions using panel data on new matches and separations. The main intuition extends and generalizes the empirical strategy of GJR in order to study marriage market segregation along observed characteristics. In our data, we observe that, for a large share of new couples, the educational levels of the partner are the same. Two alternative explanations can rationalize this pattern; either the gains from marriage are comparatively stronger for homogamous couples, or similar people face low market frictions. Yet, if we also observe the stability of realized matches, we can disentangle between the two explanations; if homogamous couples show low separation rates, then the high match rates are explained by a high match surplus. Alternatively, if homogamous couples display high separation rates, then the prevalence of these matches in the data is due to weak frictions among likes. Importantly, this approach to the problem can be seen as an extension of the seminal work on the econometrics of matching models by Choo and Siow (2006), Choo (2015), and Galichon and Salanié (2015). The applied literature grounded on these papers uses cross-sectional matched data to identify the match surplus under the assumption that the matching process is frictionless, as

[4] When the parties cannot commit to insure each other in all future states of the world, the ex-ante efficient limited commitment contract dictates that the parties will re-bargain over the within-household allocation only if one of them prefers to quit under the current terms. Voena (2015), Reynoso (2017) and Shephard (2018) model their marriage contracts in this fashion. In this paper, even though the insurance motive behind marriage is present, the repeated Nash bargaining contract is not ex-ante efficient and the volatility of private consumption implied by the marriage contract is higher than in the limited commitment case.

explained by Jaffe and Weber (2019). In our paper, we relax this assumption and use data on household formation and dissolution to jointly identify the match surplus and market frictions.

Our approach is related to other papers that use panel data on match formation and dissolution to identify parameters associated with a frictional search process (Wong, 2003) or with the presence of switching costs (Bruze et al., 2015). Our paper provides a general discussion on the identification of search frictions and differs from others in two further directions. Firstly, to estimate our free parameters, we develop a moment matching estimator of both the match surplus and meeting function that only relies on the aggregate stock-flow equations (i.e., the equations defining flows in and out of marriage) implied by the aggregate demographics of the model. The identifying restrictions used in the estimation hold regardless of whether the market is at its steady-state equilibrium. This results in some tangible advantages, such as that our estimator is fast to compute and allows for rich parametric specifications of both the meeting and surplus function. Secondly, using the estimated model, we can derive predictions about untargeted moments of the steady-state distribution of couples and singles, as well as about the intrahousehold allocation of resources. Comparing these predictions with the corresponding data moments, we can gain insight into whether the model steady-state equilibrium is an accurate description of the marital patterns observed in the data. Our estimation results imply a household sharing rule that closely resembles the intrahousehold distribution of private leisure time observed in the data in different periods. This differentiates our approach from that of GJR and Shephard (2018), who use time-use data to estimate structural parameters of household production and time-use preferences, as well as search frictions.

The third contribution of the paper is empirical. We estimate the model with multiple data cuts - one for each decade starting from the 1970s - from the U.S. Panel Survey of Income Dynamics (PSID), and provide a detailed account of the structure of the gains from marriage and search frictions in the U.S. across five decades. Previous empirical works have not considered search frictions as a determinant of assortative mating (Chiappori et al., 2017; Eika et al., 2019; Ciscato and Weber, 2020). At least in the applied economic literature, our work represents a first attempt to empirically study marriage market segmentation along socioeconomic observables, such as educational levels, family backgrounds, and wages. Moreover, even without using consumption or time-use data, we are able to delve into the nature of the marriage surplus by disentangling between its static and dynamic components. Hence, our structural findings can be related to the household economics literature, which suggests that wages are substitutes in household production (Becker, 1973) and that gender roles are asymmetric (Bertrand et al., 2015). On the other hand, we also show that the joint presence of two high earners in the household generates gains through insurance against income shocks (Blundell et al., 2016b).

Our empirical analysis is closely related to a strand of the literature that has studied the fundamental changes leading to the decline of marriage in the U.S. over the last 50 years (Aiyagari et al., 2000; Regalia et al., 2001; Greenwood et al., 2003; Heathcote et al., 2010; Guvenen and Rendall, 2015; Greenwood et al., 2016). In particular, the mechanism discussed in our counterfactual analysis is qualitatively similar to that of Greenwood et al. (2016); in their model, the shrinking gender wage gap and increased college wage premium can partly explain the decline of marriage and the increase in cross-sectional income inequality. At the same time, they argue that only technological process in the home can trigger such a fast rise in women's college education and labor market participation rates. Finally, while most papers have been interested in cross-sectional income inequality between households, we discuss differences in ex-ante welfare among young individuals. This approach is similar to that of Fernández and Wong (2017), who study the distributional impact of divorce law reforms. The ex-ante welfare takes into account the expectations about the timing, the duration, and the number of matches an agent will have during the entire life-cycle. Given the differences in family outcomes across human capital levels outlined at the beginning of this paper, we believe that looking at the distribution of ex-ante welfare is particularly insightful to understand the diverging destinies of different population groups.

2. The Model

2.1. Outline

Every agent is born with a given family background and chooses whether to go to college at the age of 17. Her decision depends on both labor and marriage market gains from college education. The cost of her graduate studies includes monetary tuition fees, which depend on the state where she was raised, and some non-monetary component, which partly depends on her family background and is partly random.

After completing her studies, the agent starts her adult life as a young bachelor(ette). The breakdown of a working-life period, corresponding to two years in our application, is presented in Figure 1. In every period, the agent makes consumption and time-use decisions based on her full-time wage. The latter depend on her gender, education, age, and family background, and evolve according to a Markov process. At the same time, the agent can now look for a partner on the marriage market. Search is random and time-intensive; in every period, the agent meets at most one person of the opposite gender drawn from the pool of single agents available on the marriage market. On the other hand, search frictions are such that meetings can be positive, or negative, assortative; for instance, a 50-year-old college graduate is more likely to have a "date" with someone around her age and with the same education. However, meeting prospects depend on the supply of singles in the market; for instance, relatively old agents might face a tighter market, as many of their peers are already married. We do not impose any restriction on the strength and direction of these search frictions. Whether agents are more likely to meet their likes is a question that will be answered in the empirical analysis.

On dates, two agents learn about the quality of their relationship and decide whether to form a new couple. Couples pool resources and benefit from economies of scale in consumption and household production, as well as non-monetary gains from being together. The match surplus is split endogenously, with the equilibrium consumption share of each partner being a function of her own reservation value, i.e., the value of being single again. Likewise, we do not impose any ex-ante restriction on the nature of the match gains; how these depend on observed characteristics, such as the partners' education, age, wage and family background, is addressed in the empirical analysis.

Couples grow older together and face both wage and match-quality shocks. Hence, in every period, they need to decide whether to continue their relationship or to break up. In case of a separation, both partners are single again and can look for a new partner, although they are now older. At the age of 63, the agent retires: from then onwards, her earnings correspond to the full-time wage she earned in her last working-age period. If she is in a relationship, she is free to continue it. However, if she is single, divorces or becomes widowed, she can no longer look for a new partner.

Figure 1: Breakdown of a working-life period

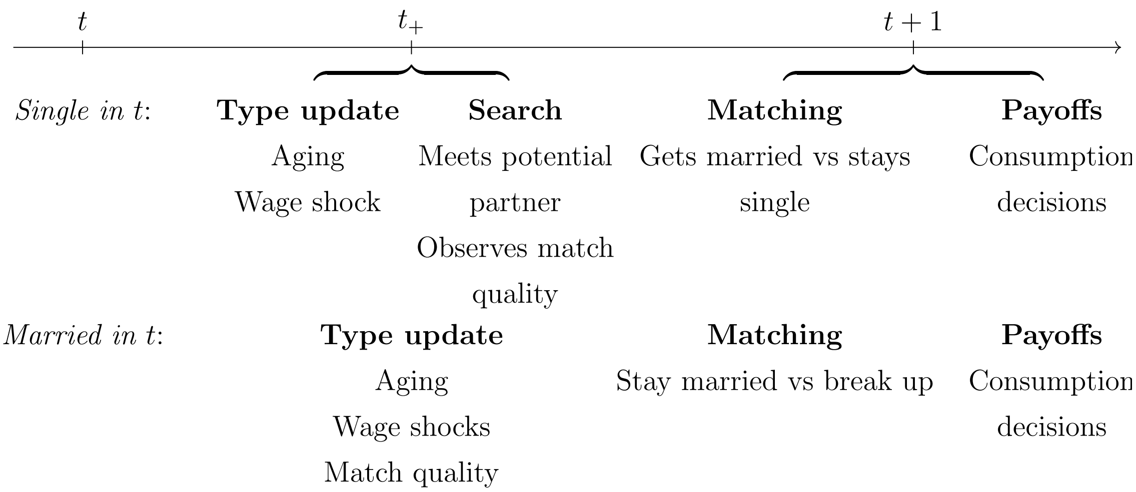

2.2. Working life

We assume working life and retirement stretch across multiple time periods and adopt the following notation: when a variable v refers to period t, v' refers to $t+1$. Young individuals enter the labor and marriage market at the age of 19 if they did not go to college, and at the age of 21 if they did. They also draw an initial wage w from a distribution that depends on both their education h and family background x. We denote $i = \{w, h, x, a\} \in \mathcal{I}^g$ the *type* of an agent, with $\mathcal{I}^g$ being i's discrete support for gender $g \in \{m, f\}$. To avoid confusion, i denotes the type of a man and j the type of a woman when the distinction is needed.

At the end of each period, an agent dies with probability $1 - \psi_i^g$, which depends on the agent's gender g and type.[5] We impose a terminal age $\bar{a} = 83$, so that $\psi_i^g = 0$ if $a = \bar{a}$. In every new period, surviving agents get older, so that $a' = a + 1$, and draw a new wage w'. Wages follow a first-order Markov process; the distribution of w' depends on the last realization w, as well as on the agent's gender, age, family background, and education. Since types are discrete, we introduce transition matrices $(\pi_{i,i'}^g)$ for each gender g that describe the evolution of types over time as a discrete-time Markov chain over the discrete space $\mathcal{I}^g$.[6]

2.3. Intra-household allocation

Agents that enter the market in period t are single, but older agents may be either in a relationship or single, as a result of their personal marital history. Before describing matching decisions, we first characterize the payoffs from living alone and from being in a relationship. Single men (women) of type i (j) enjoy per-period utility $v_{i,0}^m \in \mathbb{R}$ ($v_{0,j}^f \in \mathbb{R}$). A woman of type j married with a man of type i enjoys per-period utility $v_{i,j}^f(z) \in \mathbb{R}$, while her husband enjoys $v_{i,j}^m(z) \in \mathbb{R}$, where $z \in \mathbb{R}$ is a scalar denoting the current quality of the match.

Contrarily to the frictionless case of Shapley and Shubik (1971) and Becker (1973), where stability conditions pin down a rule to divide match surplus, the presence of search frictions requires us to introduce an additional assumption about the bargaining process (Shimer and Smith, 2000). We start by introducing the following notation: $V_{i,j}^g(z)$ is the present discounted value of being in a match $\{i, j, z\}$ for a partner of gender g, $V_{i,0}^m$ and $V_{0,j}^f$ a man i's and a woman j's present discounted value of being single. We define the *total surplus* for a match $\{i, j, z\}$ as

$$S_{i,j}(z) \equiv V_{i,j}^m(z) + V_{i,j}^f(z) - V_{i,0}^m - V_{0,j}^f. \tag{2.1}$$

We make the following assumption about the way partners choose the household allocation.

Assumption 1 (Bargaining)**.** *The preferred allocation, associated with payoffs* $\{v_{i,0}^m, v_{0,j}^f\}$*,*

[5] In the empirical application, we let $1 - \psi_i^g$ be the mortality rate of gender g at age a, and do not let ψ^g depend on other observed characteristics. This model can easily be extended to allow for both (exogenous) outflows and inflows of agents at any age. This can help to account for migration, active duty in the army, incarceration, long-term stay in health-care institutions, etc.

[6] Since h and x are time-invariant, we have $\pi_{i,i'}^m = 0$ if $h' \neq h'$ or $x' \neq x$. Moreover, since aging means that agents either get older or die in every new period, we have $\pi_{i,i'}^m = 0$ if $a' \neq a + 1$.

1. *lies on the household Pareto frontier;*
2. *is s.t. Individual Rationality is verified:*

$$V_{i,j}^m(z) \geq V_{i,0}^m \tag{2.2}$$

$$V_{i,j}^f(z) \geq V_{0,j}^f; \tag{2.3}$$

3. *is s.t. the total match surplus is split as follows:*

$$V_{i,j}^m(z) - V_{i,0}^m = \theta S_{i,j}(z) \tag{2.4}$$

with $\theta \in [0, 1]$.

The value function V^g is an equilibrium quantity; it does not only depend on the (exogenous) bargaining parameter θ and household technology, but also on agents' expected (re)marriage prospects over the life-cycle. The exact relationship between v^g and V^g will be determined later using the Bellman equations.

Finally, we require partners' preferences to verify the Transferable Utility property. In other words, we allow agents to make costless monetary transfers to each other in order to reach the allocation they agreed upon. The following restriction on the shape of the household Pareto set is a necessary and sufficient condition for TU (Demuynck and Potoms, 2020).

Assumption 2 (Transferable Utility)**.** *For each* $\{i, j, z\}$*, there exist two continuous and increasing functions* $\phi^m : \mathbb{R}_+ \to \mathbb{R}_+$ *and* $\phi^f : \mathbb{R}_+ \to \mathbb{R}_+$ *and a constant* $H_{i,j}(z)$ *s.t.*

$$\phi^m(v_{i,j}^m(z)) + \phi^f(v_{i,j}^f(z)) = H_{i,j}(z). \tag{2.5}$$

Importantly, the constant $H_{i,j}(z)$ plays a key role in determining the gains for a match $\{i, j, z\}$; the larger $H_{i,j}(z)$, the wider the couple's Pareto set. In other words, the larger $H_{i,j}(z)$, the greater the gains associated with current joint household production and being in a relationship. In fact, $H_{i,j}(z)$ can be interpreted as the household budget after taking into account the monetary value of the match. Note that ϕ^g depends on the curvature of the utility function, so that $\phi^g(v_{i,j}^g(z))$ represents the private budget of partner g. For this reason, we can define the sharing rule as follows:

$$\gamma_{i,j}(z) \equiv \frac{\phi^m(v_{i,j}^m(z))}{H_{i,j}(z)}. \tag{2.6}$$

2.4. Household formation and dissolution

We denote t_+ the instant that follows the entry and exit of individuals and the update process of types; we write $\tilde{v} \equiv v_{t_+}$ when a variable v refers to period t and $\tilde{v}$ to t_+. In

t_+, individuals are free to make matching decisions following changes in their types, but the market has not cleared yet (see Figure 1). A single man of type i meets a single woman of type j with probability $M_{i,j}(\tilde{n})/\tilde{n}_i^m$, while a single woman of type j meets a single man of type i with probability $M_{i,j}(\tilde{n})/\tilde{n}_j^f$. The function $M : \mathcal{I}^m \times \mathcal{I}^f \to \mathbb{R}_+$ yields the number of meetings between men i and women j in each period; $M_{i,j}$ takes $\tilde{n}$ as an argument, where $\tilde{n}$ denotes the distribution of singles on the market in t_+. The *meeting function* M needs to be such that the following feasibility constraints are respected:

$$\tilde{n}_i^m \geq \sum_j M_{i,j}(\tilde{n}) \tag{2.7}$$

$$\tilde{n}_j^f \geq \sum_i M_{i,j}(\tilde{n}). \tag{2.8}$$

Upon a meeting, two singles draw an initial match quality z from a distribution G^{new}. Hence, they decide whether to match or not based on their initial state $\{i, j, z\}$. Similarly, in t_+, matched agents decide whether to continue their relationship after drawing new types of both partners (i and j) and a new match quality (z). In particular, they draw z from a distribution G, possibly different from G^{new}. At the same time, we rule out serial correlation between match quality draws.[7]

2.5. Value functions

As a consequence of Assumption 1, matching decisions only depend on the sign of the total match surplus: if $S_{i,j}(z) > 0$, a couple $\{i, j, z\}$ is willing to get married (or to continue the relationship). Averaging over match quality z, we can express the odds of getting matched and continuing a relationship as follows:

$$\alpha_{i,j}^{new} \equiv \int \mathbb{1}\{S_{i,j}(z) > 0\} dG^{new}(z) \tag{2.9}$$

$$\alpha_{i,j} \equiv \int \mathbb{1}\{S_{i,j}(z) > 0\} dG(z), \tag{2.10}$$

[7]The model is still fully tractable after introducing serial correlation, as showed by Shephard (2018). The main identification findings of this paper can be extended to the case with serial correlation. However, since z is treated as a term unobserved by the econometrician in the empirical analysis, this poses additional challenges in the estimation. In contrast, most findings can be easily extended to the case where the distribution G depends on match duration, provided that the latter is observed by the econometrician. The work of Bruze et al. (2015) can be regarded as an instance of this approach, although search frictions are ruled out.

while the expected surplus from a new match and from continuing a relationship are respectively denoted as:

$$\bar{S}^{new}_{i,j} \equiv \int \max\{S_{i,j}(z), 0\} dG^{new}(z) \tag{2.11}$$

$$\bar{S}_{i,j} \equiv \int \max\{S_{i,j}(z), 0\} dG(z). \tag{2.12}$$

The value functions V^m and V^f solve the Bellman equations:

$$V^m_{i,j}(z) = v^m_{i,j}(z) + \beta \mathbb{E}_{i'|i} V^m_{i',0} + \beta\theta \mathbb{E}_{i',j'|i,j} \bar{S}_{i',j'} \tag{2.13}$$

$$V^f_{i,j}(z) = v^f_{i,j}(z) + \beta \mathbb{E}_{j'|j} V^f_{0,j'} + \beta(1-\theta) \mathbb{E}_{i',j'|i,j} \bar{S}_{i',j'}, \tag{2.14}$$

where the last term on the rhs of both equations corresponds to the partner's continuation value from the current match.[8]

Match prospects for singles depend on the supply of unmatched agents on the market, $\tilde{n}^g$. In what follows, we focus on the steady-state search equilibrium and conjecture that single agents looking for a partner face the same distribution of available singles $\tilde{n}^g$ in every future period. The value functions $V^m_{i,0}$ and $V^f_{0,j}$ for single agents solve the Bellman equations:

$$V^m_{i,0} = v^m_{i,0} + \beta \mathbb{E}_{i'|i} V^m_{i',0} + \beta\theta \mathbb{E}_{i'|i} \sum_j \frac{M_{i',j}(\tilde{n})}{\tilde{n}^m_{i'}} \bar{S}^{new}_{i',j} \tag{2.15}$$

$$V^f_{0,j} = v^f_{0,j} + \beta \mathbb{E}_{j'|j} V^f_{0,j'} + \beta(1-\theta) \mathbb{E}_{j'|j} \sum_i \frac{M_{i,j'}(\tilde{n})}{\tilde{n}^f_{j'}} \bar{S}^{new}_{i,j'} \tag{2.16}$$

where the last term represents the expected surplus from search on the marriage market in the next period.

2.6. Educational choices

At the age of 17, a single agent with family background x needs to decide whether to go to college. The gains from college education correspond to the difference between the (expected) present discounted value of starting working life with or without a degree. Since all agents start their working lives as singles, such gains correspond to:

$$\Delta^g(x) \equiv \mathbb{E}_{i'} \left[V^g_{i',0} \middle| h = CG, a = 17, x \right] - \mathbb{E}_{i'} \left[V^g_{i',0} \middle| h = NC, a = 17, x \right]. \tag{2.17}$$

[8] As agents face both mortality and wage risk, their expected surplus is computed as

$$\mathbb{E}_{i',j'|i,j} \bar{S}_{i',j'} = \psi^m_i \psi^f_j \sum_{i',j'} \pi^m_{i,i'} \pi^f_{j,j'} \bar{S}_{i',j'}$$

where, we remind the reader, $1 - \psi^g_i$ is the death probability for an agent of gender g with type i and $\pi^g_{i,i'}$ is the probability of transitioning from type i to type i' for an agent of gender g. Similar formulas hold for the expected reservation values.

Going to college is costly, and the gender-specific cost function C^g depends on both the family background x, the state s where the agent was raised, and an idiosyncratic term $\eta \in \mathbb{R}$. Foregone earnings and match opportunities constitute additional implicit costs of college education, as college students need to wait until they are 21 to get their degree and start working and looking for a partner. In summary, the agent goes to college if:

$$\Delta^g(x) - C^g(x, s, \eta) > 0. \tag{2.18}$$

2.7. Population dynamics

We denote $\ell^m : \mathcal{I}^m \to \mathbb{R}_+$ and $\ell^f : \mathcal{I}^f \to \mathbb{R}_+$ the marginal distributions of men and women on the marriage market. These distributions result from both the endogenous educational choices and the exogenous initial composition of new cohorts, as well as cohort dynamics as described by π^g. We assume that such marginals are stationary.

Assumption 3 (Populations). *For each Δ^g, the marginal distribution on each side of the market g is described by a stationary PDF ℓ^g. In particular, we assume that:*

a. size and composition of a new cohort of individuals is constant over time;
b. every cohort faces identical mortality and wage risk.

The first part of the assumption ensures that a new cohort of young individuals has the same size and gender ratio of the previous. Moreover, every new cohort must face the same distribution of college costs. The second part of the assumption demands that, given its initial distribution and educational choices, every cohort follows the same aggregate trajectory. In other words, all cohorts will have the same life expectancy and wage risk. On the other hand, assumption 3 does not exclude changes in composition within the same cohort. Since men typically have a lower life expectancy, the cohort's gender ratio can be tilted in favor of women as the cohort gets older. Moreover, the wage process does not need to be stationary, and the wage mean and variance can both increase as the cohort gets older.

We denote $m : \mathcal{I}^m \times \mathcal{I}^f \to \mathbb{R}_+$ the number of couples of type $\{i, j\}$ at the end of the period. The aggregate count $m_{i,j}$ result from the following aggregate law of motion:

$$m_{i,j} = \tilde{m}_{i,j} + \underbrace{\alpha_{i,j}^{new} M_{i,j}(\tilde{n})}_{MF_{i,j}} + \underbrace{(1 - \alpha_{i,j})\tilde{m}_{i,j}}_{DF_{i,j}} \tag{2.19}$$

where $\tilde{m}$ represents the number of couples of a given type at risk of separation at the beginning of the period, whereas $MF_{i,j}$ and $DF_{i,j}$ respectively denote the number of new matches (the "marriage flow") and separations (the "divorce flow").

The functions $n^m : \mathcal{I}^m \to \mathbb{R}_+$ and $n^f : \mathcal{I}^f \to \mathbb{R}_+$, respectively, represent the number of male singles and female singles at the end of each period. These functions result from the following accounting restrictions, which must hold in any period:

$$\ell_i^m = n_i^m + \sum_j m_{i,j} \tag{2.20}$$

$$\ell_j^f = n_j^f + \sum_i m_{i,j}. \tag{2.21}$$

To complete the description of population dynamics, we express $\tilde{n}^g$ as a function of the end-of-the-period distribution of singles n^g. Similarly, we write $\tilde{m}$ as a function of m:

$$\tilde{n}_{i'}^g = \begin{cases} \sum_i \psi_i^g \pi_{i,i'}^g n_i^g & \text{if } a' > 17 \\ \ell_i^g & \text{if } a' = 17 \end{cases} \tag{2.22}$$

$$\tilde{m}_{i',j'} = \sum_{i,j} \psi_i^m \psi_j^f \pi_{i,i'}^m \pi_{j,j'}^f m_{i,j}. \tag{2.23}$$

2.8. Equilibrium

We define the steady-state search equilibrium by combining the key equations outlined throughout Section 2.

Definition 2.1 (Search Equilibrium). *A steady-state search equilibrium is given by time-invariant measures of singles n^m and n^f over $\mathcal{I}^m$ and $\mathcal{I}^f$ respectively, so that all the following conditions hold.*

- *Matching decisions only depend on the sign of the match surplus S.*
- *For given values of being single $V_{i,0}^m$ and $V_{0,j}^f$, the surplus function S solves a system of $|\mathcal{I}^m| \times |\mathcal{I}^f|$ equations implied by the Pareto frontier (2.5) and the Bellman equations for married agents (2.13) and (2.14).*
- *For a given distributions of singles n^m and n^f, $V_{i,0}^m$ and $V_{0,j}^f$ solve a system of $|\mathcal{I}^m| + |\mathcal{I}^f|$ equations given by the Bellman equations for single agents (2.15) and (2.16).*
- *For given values of being single $V_{i,0}^m$ and $V_{0,j}^f$, the marginals of the matching game ℓ^m and ℓ^f are built to reflect agents' educational choices as described by (2.18).*
- *Finally, n^m and n^f solve a system of $|\mathcal{I}^m| + |\mathcal{I}^f|$ equations implied by the accounting restrictions (2.20) and (2.21), where m is a function of n given by the law of motion (2.19).*

This definition suggests that the equilibrium is a fixed-point (n^m, n^f) of an operator T over the support $\mathbb{R}_+^{|\mathcal{I}^f|+|\mathcal{I}^f|}$. Recently, Manea (2017) and Shephard (2018) generalized the original proof of existence by Shimer and Smith (2000) for this class of models, but uniqueness is not guaranteed. In practice, iteration of the fixed-point operator always results in convergence to the same distribution n for a given set of primitive parameters, regardless of the initial values chosen.

3. Identification

3.1. Match surplus

The identification of the surplus function in matching models starting from cross-sectional matched data relies on basic demographic restrictions. In our specific case, if we observe the matching outcome $(\hat{m}_{i,j})$, $(\hat{n}_i)$ and $(\hat{n}_j)$ in the data, we can recover the matching rules α and α^{new} from the law of motion.[9] However, this is possible only if the meeting function M and the distribution of the initial match quality draws G^{new} were known to the econometrician.

Theorem 3.1. *Assume G is known and $S_{i,j}(z)$ is nondecreasing in z. If M and G^{new} are also known and the inequality condition*

$$M_{i,j}(\tilde{n}) + \tilde{m}_{i,j} > m_{i,j} \tag{3.1}$$

holds for any $\{i,j\}$, then the functions α, α^{new} are identified over $\mathcal{I}^m \times \mathcal{I}^f$ with data on the stock of matches $(\hat{m}_{i,j})$ and unmatched agents $(\hat{n}_i)$ and $(\hat{n}_j)$.

Proof. Start from noting that (2.9) and (2.10) share the same integrand, the indicator function $\mathbb{1}\{S_{i,j}(z) > 0\}$. Since $S_{i,j}$ is nondecreasing in z, there exists a threshold $\bar{z}_{i,j}$ s.t. $S_{i,j}(z) > 0$ for any $z > \bar{z}_{i,j}$. Hence, we input (2.9) and (2.10) into the law of motion (2.19) and obtain

$$m_{i,j} = \tilde{m}_{i,j}(1 - G(\bar{z}_{i,j})) + M_{i,j}(\tilde{n})(1 - G^{new}(\bar{z}_{i,j}))$$

that can be rearranged into

$$G(\bar{z}_{i,j}) = 1 - \frac{m_{i,j}}{\tilde{m}_{i,j}} + \frac{M_{i,j}}{\tilde{m}_{i,j}}(1 - G^{new}(\bar{z}_{i,j})).$$

Note that the lhs is an increasing function of $\bar{z}_{i,j}$ with image $[0,1]$, while the rhs is a decreasing function of $\bar{z}_{i,j}$ with image

$$\left[\frac{\tilde{m}_{i,j} - m_{i,j}}{\tilde{m}_{i,j}}, \frac{M_{i,j} + \tilde{m}_{i,j} - m_{i,j}}{\tilde{m}_{i,j}}\right] \equiv [\underline{b}, \bar{b}].$$

[9] As already seen in Choo (2015), there are two possible approaches. The econometrician can work with repeated cross-sections that track the evolution of $\hat{m}_{i,j,t}$ over different data years t. Alternatively, the econometrician can assume the data patterns correspond to the steady-state equilibrium matching and only use variation in $\hat{m}_{i,j}$ across ages within one data year.

Hence, $\bar{z}_{i,j}$ exists and is unique if the two sets are not disjoint, which can easily be checked. The inequality $\bar{b} < 1$ always holds and the inequality $\bar{b} > 0$ always holds by assumption (3.1). Given $\bar{z}_{i,j}$, we can use (2.10) and (2.9) to build α and α^{new}, respectively. □

Condition (3.1) introduces inequality restrictions for the meeting function M; if the researcher's prior over $M_{i,j}$ is incorrect, then $\alpha_{i,j}$ is not identified. The economic intuition is that, if the function M cannot generate a sufficient number of meetings between types $\{i, j\}$ to explain the difference $m_{i,j} - \tilde{m}_{i,j}$ between new couples of type $\{i, j\}$ and couples of the same type from the last period, then M cannot rationalize the patterns observed in the data.

Our first result is a variation on the standard results obtained by Choo and Siow (2006) and Galichon and Salanié (2015) on the identification of the static Becker-Shapley-Shubik model with cross-sectional matched data. Choo (2015) and Bruze et al. (2015) extend the identification results to its dynamic version. The identification of the surplus function in this class of models relies on a distributional assumption on the random components of matching preferences. In this paper, we maintain this assumption and assume the distribution G is known. However, as discussed by Jaffe and Weber (2019), another key identifying assumption is that the marriage market is frictionless and meetings are random. In the next section, we relax this second assumption, allowing for frictions.

3.2. Search frictions

In a vast majority of cases, the econometrician does not observe the number of meetings $M_{i,j}$ between types and is interested in learning about the importance of search frictions. In fact, the following question is often left unanswered by the economic literature on marriage markets: Are matches between two specific types i and j frequent (rare) because of a high (low) match surplus or because these types face low (high) frictions?

In what follows, we develop and formalize the key intuition of GJR that data on new matches and separations can help solve this identification puzzle. Consider an observational setting where the econometrician observes data on gross flows in and out of marriage, i.e., the number of new matches and separations by type, $(\widehat{MF}_{i,j})$ and $(\widehat{DF}_{i,j})$. This typically requires panel data with at least two periods ($T \geq 2$) to observe changes in household composition between consecutive waves. We use the additional information contained in gross flows to disentangle the role of search frictions in explaining the observed matching patterns.

Theorem 3.2. *Assume G is known and $S_{i,j}(z)$ is nondecreasing in z. If G^{new} is also known, M, α and α^{new} are identified over $\mathcal{I}^m \times \mathcal{I}^f$ with matched data on both stocks $(\hat{m}_{i,j})$, $(\hat{n}_i)$, $(\hat{n}_j)$ and flows $(\widehat{MF}_{i,j})$, $(\widehat{DF}_{i,j})$.*

On the other hand, we also warn against the possibility of disentangling between different sources of frictions; if G^{new} differs from G, the identification of M will only be feasible with additional parametric restrictions on either G^{new}, M, or both. Yet, learning whether G^{new} differs from G might help us better understand the nature of marriage market frictions. For instance, new marriages can be associated with a sunk cost due to learning or to the necessary investment to start off the household. Additionally, the distribution of new match quality draws, G^{new}, can be associated with a higher variance.

Corollary 3.2.1. *Assume G is known and $S_{i,j}(z)$ is nondecreasing in z. Up to $\mathcal{I}^m \times \mathcal{I}^f$ parameters of M and G^{new} are identified with matched data on both stocks $(\hat{m}_{i,j})$, $(\hat{n}_i)$, $(\hat{n}_j)$ and flows $(\widehat{MF}_{i,j})$, $(\widehat{DF}_{i,j})$.*

Proof. Recall from the law of motion (2.19) the analytical expressions of marriage and divorce flows.

$$MF_{i,j} = \alpha_{i,j}^{new} M_{i,j}(\tilde{n}), \quad DF_{i,j} = (1 - \alpha_{i,j})\tilde{m}_{i,j}.$$

First, as G is known, we can recover $\bar{z}_{i,j}$ starting from the analytical expression of DF:

$$\bar{z}_{i,j} = G^{-1}\left(\frac{DF_{i,j}}{m_{i,j}}\right). \tag{3.2}$$

This leads us to recover $\alpha_{i,j}$. Then we proceed in two alternative ways depending on the target of the identification.

1. When G^{new} is known, then $\alpha_{i,j}^{new}$ can be computed as $1 - G^{new}(\bar{z}_{i,j})$ and $M_{i,j}$ follows as $MF_{i,j}/\alpha_{i,j}^{new}$, which proves theorem 3.2.
2. When G^{new} is unknown, we end up with a system of $\mathcal{I}^m \times \mathcal{I}^f$ equations as we replace $\alpha_{i,j}^{new}$ with $1 - G^{new}(\bar{z}_{i,j})$ in the stock-flow equation implied by the definition of $MF_{i,j}$. This puts an upper bound on the number of free parameters of the unknown functions G^{new} and M, and proves corollary 3.2.1.

□

The intuition behind theorem 3.2 is the following; both marriage and divorce flows contain information about the underlying marital surplus. However, the number of new matches also informs us about the amount of frictions that single individuals need to overcome when search is costly. Information about flows in and out of marriage can thus be combined to determine the role of frictions and of match surplus. In applied work, corollary 3.2.1 offers guidance in the choice of the number of free parameters when choosing a specification for G^{new} and M; the total number of free parameters in G^{new} and M cannot exceed $|\mathcal{I}^m| \times |\mathcal{I}^f|$.

Similar empirical strategies have already been used in the matching literature. In order to identify the arrival rate of meetings in their random search models, Wong (2003) uses data on the duration of marriages, while GJR uses marriage and divorce rates. Bruze et al. (2015) assume away search frictions and use data on the duration of marriage to identify switching costs as a function of duration. Finally, Greenwood et al. (2016) and Shephard (2018) assume that the unobserved match-quality is autocorrelated over time. They use information on marriage duration and divorce risk to identify both the parameters of the match-quality process and the arrival rate of meetings.

3.3. Household production

So far, we did not have to impose that the marriage market is at the steady-state to identify the objects of interest. With data spanning across different years, both $\alpha_{i,j,t}$ and $M_{i,j,t}$ are identified with matched data for any data year t using the stock-flow equations defining the flows in and out of marriage, named MF and DF, respectively. These equations hold regardless of whether the market is at the steady-state equilibrium. However, if we want to decompose the surplus S and distinguish its dynamic components from the static gains from marriage for a given data year t, an assumption is needed to model the agent's expectations about the distribution of singles on the market and meeting probabilities in future periods $s > t$. Imposing stationarity is one solution to the problem; in this way, the estimated meetings M from the current period t are enough to proceed.[10]

In this paper, we follow the recent marriage market literature (Galichon and Salanié, 2015; Chiappori et al., 2018b) and additionally assume that the random match quality component is additively separable:

Assumption 4 (Separability)**.** *Match quality z enters the match surplus function additively:*

$$S_{i,j}(z) = S_{i,j} + z. \tag{3.3}$$

Assumption 4 makes it straightforward to obtain $S_{i,j}$, the deterministic component of the match surplus, by inverting (2.10).

The calibration of the remaining parameters, notably the functions ϕ^m, ϕ^f, and the bargaining parameter θ, allows us to determine the location of the Pareto frontier $H_{i,j}(z)$ for a household $\{i, j, z\}$, and the spouses' payoffs $v^m_{i,j}(z), v^f_{i,j}(z)$, starting from $S_{i,j}(z)$.[11] Importantly,

[10] Alternatively, if agents are rational and have perfect foresight, the time series of M for the period ahead of t is needed in order to build agents' expected surplus in a non-stationary environment.

[11] Additional data on household behavior, such as time-use and consumption data, could help us recover ϕ^m,

if these additional parameters are known, we can disentangle the component of the match surplus that comes from household production from both the partners' reservation values and the match continuation value. More explicitly, if these parameters are known, we can establish a one-to-one map between S and H over the space $\mathcal{I} \times \mathcal{J} \times \mathbb{R}$. Details on how to proceed can be found in Appendix A.1.

4. Empirical specification and estimation

4.1. Outline

We estimate the model with five separate sample cuts from the main PSID database, each covering one of the following time periods: 1971-1975, 1981-1985, 1991-1995, 2001-2005 and 2011-2015. For every sample cut, the estimation is carried out in separate steps. Firstly, we estimate the parameters of the wage process using panel data on wages. Secondly, we jointly estimate α, the match probabilities, and M, the meeting function, with panel data on household formation and dissolution. Thirdly, we estimate the cost of going to college with data on educational achievement and using estimates for labor and marriage market returns from the two previous steps.

4.2. Data

From the main PSID sample, we only keep observations from the SRC (Survey Research Center), whereas we exclude observations from the SEO (Survey of Economic Opportunity), the immigrant and Latino samples. Recent waves of the original SRC sample are known to not be representative of the American population as they do not account for recent waves of immigration (Brown et al., 1996). On the other hand, limiting the analysis to the SRC sample allows for easier comparison across different time periods and greatly limits attrition issues. We keep all individuals aged between 19 and 84 that live at home or in an institution at the moment of the interview.

Individuals are classified as either "matched" or "unmatched", depending on the "relationship-to-head" variable. Couples of matched individuals are identified through the presence of the "head's partner". Individuals that are not the head nor his/her partner are classified as "unmatched".[12] The partner can be either a legal spouse or an unmarried cohabitor. It is only

ϕ^f and θ in the same way discussed by GJR and Shephard (2018). In this paper, we follow a different approach and leave time-use data out of the estimation; however, we use them to check whether the predictions of the estimated model about the sharing rule are correct.

[12] In particular, members of secondary families are counted as "unmatched", although some of them may be married: e.g., a young spouse may live with his/her in-laws. In this case, we only count them as "matched"

possible to distinguish between the two types of partner starting from 1983, and we choose not to distinguish between legal marriage and cohabitation in our benchmark sample. Starting from the 1990s, we can build alternative samples where a match corresponds to a legal marriage; in these samples, cohabiting individuals are treated as singles, similarly to Chiappori et al. (2017).

We study couple formation and dissolution by observing changes in matches over two consecutive waves. Because interviews have been conducted every two years since 1997, we focus on changes that occur over a two-year time span. "New matches" occur if: i) an individual lives alone in year y and with a partner in $y+2$; ii) an individual lives with a partner in y and with a different partner in $y+2$. Conversely, "separations" occur if: i) an individual lives with a partner in y and alone in $y+2$; ii) an individual lives with a partner in y and with a different partner in $y+2$. By observing changes in household composition, we can distinguish and thus count out relationships ending due to the death of a partner.

4.3. Marginals

The matching variables constituting the type i of an agent include the age group, family background, education, and hourly wage. As for age groups, we compute age as the interview year minus the year of birth, and then assign individuals to age groups spanning two years (e.g. a 29- and a 30-year-old belong to the same group) as a time period corresponds to two years in our application. We calibrate mortality rates using vital statistics from the Center for Disease Control and Prevention, and set the gender ratio at the age of 19 in order to match the overall gender ratio of the working-age population in the PSID sample.

We build a family background variable using the first principal component factor from a set of family background characteristics, namely parents' education and socioeconomic status, as suggested by Blundell et al. (2016a). We then assign people to a high or a low type depending on whether their score is above or below the mean of their age cohort. As for education, we use the panel dimension of the database and we define education as the highest educational level achieved in lifetime; we then divide respondents into two categories, those with a college degree and those without. We label people that worked an average of more than 15 weekly hours per week as "employed" and compute their hourly wages as the ratio between the (deflated) total yearly labor income and the total yearly number of hours worked.[13] Finally, we input the wages of non-employed individuals after imposing more structure on the wage process and estimating

once they establish their own household.

[13] In the PSID, information on earnings and labor supply is only available for surveyed heads and partners, but not for other members of the household. In particular, this matters for young individuals that are living with their parents. We assume that these people are not employed.

its parameters, as detailed below.

We assume the marginal distribution of wages is log-normal, with its moments depending on gender g, education h, family background x, and age a:

$$\ln w_{i,t} = \vartheta^g_{0,h}(x) + \vartheta^g_{1,h} a_{i,t} + \vartheta^g_{2,h} a^2_{i,t} + \sigma_{i,t}\xi_{i,t} \tag{4.1}$$

$$\ln \sigma_{i,t} = \varpi^g_{0,h} + \varpi^g_{1,h} a_{i,t} \tag{4.2}$$

$$\xi_{i,t} \sim \mathcal{N}(0,1) \tag{4.3}$$

where all coefficient are gender- and education-specific. We call $r_{i,t} \equiv \Phi^{-1}(\xi_{i,t})$ the wage rank of agent i in period t among those that share the same gender g, education h, family background x and age a, with Φ^{-1} denoting the inverse of the standard normal CDF. We model wage mobility by using a copula representation, as in Bonhomme and Robin (2009). In other words, the joint CDF of $r_{i,t}$ and $r_{i,t+1}$ is given by a copula function. In particular, we choose a Plackett's copula; technical details are provided in Appendix A.2. In this way, the degree of wage mobility only depends on one parameter, which is in its turn a strictly increasing function of the Spearman's rank correlation coefficient.[14] We allow the Plackett's copula parameter, and thus the degree of wage mobility, to differ by gender g, education h, and age a:

$$\ln \rho_{i,t} = \varrho^g_{0,h} + \varrho^g_{1,h} a_{i,t}. \tag{4.4}$$

We estimate equations (4.1), (4.2) and (4.4) by Maximum Likelihood Estimation, separately for each education-gender group and using working-age individuals (19- to 62-year-olds). In doing so, we address the issue of selection into employment on unobservables through a standard parametric control function approach (Heckman, 1974). In other words, we add a selection equation where we let the probability of being employed depend on a third-degree polynomial of age, family background, and a second-degree polynomial of residual household transfers. As in Chiappori et al. (2018a), this residual is obtained by regressing household transfers on household demographics, including the head's sex, whether (s)he has a partner, and the presence and age of children. The residual is meant to capture policy variation in the transfer system over time and across states.

After estimating the parameters of the wage process, we complete each observation's matching type i and assign an hourly wage to individuals who were not classified as employed. In order to do so, we regress the estimated residuals $\hat{\xi}_{i,t}$ of employed agents on a third-degree polynomial of their estimated propensity to work[15], in addition to time and state dummies.

[14] In their analysis of mobility in the US, Kopczuk et al. (2010) suggest that wage rank correlation between periods is indeed one of the most natural measures of wage mobility.

[15] This corresponds to the linear index of the previously estimated selection equation.

We use the predicted ξ for the non-employed to proxy their wages. Once all observations have been assigned an hourly wage, we discretize the support of the wage distribution using 11 bins and assign each observation to one of the groups.

4.4. Search frictions and match surplus

In our application, we assume $G^{new} = G$ and focus on search as the only source of frictions.[16] We specify the meeting function as follows:

$$\log M_{i,j}(\tilde{n}) = \lambda_{i,j} + c_i(\tilde{n}) + d_j(\tilde{n}) \tag{4.5}$$

$$\lambda_{i,j} \equiv \lambda' x^{\lambda}_{i,j}/2 \tag{4.6}$$

where $\lambda_{i,j}$ is a scalar that measures the severity of frictions for pairs $\{i, j\}$; the lower $\lambda_{i,j}$, the stronger search frictions for that type combination. The terms of $x^{\lambda}_{i,j}$ include dummies for the joint educational attainments and family background of (potential) spouses and first- and second-order polynomial terms of spouses' ages and wages, with λ being the vector of corresponding coefficients. The functions $c : \mathcal{I}^m \to \mathbb{R}_+$ and $d : \mathcal{I}^f \to \mathbb{R}_+$ ensure that the feasibility constraints (2.7) and (2.8) are always respected; c and d are computed numerically for any given $\tilde{n}$. In Appendix A.3, we provide an optimal-transport characterization of this meeting process, where $\lambda_{i,j}$ plays the role of the transport cost function: for a given population of singles described by $\tilde{n}$, meetings M minimize the total transportation costs (plus an entropy term). This characterization allows us to exploit well-known theoretical properties and fast computational methods from the optimal transport and matching literature (Galichon and Salanié, 2015).

We further impose that meetings do not occur at all ($\lambda_{i,j} \to -\infty$) if the age distance is too large; women do not date men who are more than 6 years younger or more than 14 years older than them. This is based on the empirical observation that marriages of this type are very rare in the data.[17] Finally, we also assume that retired agents are not active on marriage markets. In this case as well, marriages for people older than 63 are rare in the data.[18]

We parameterize the deterministic component of the surplus function $S_{i,j}$, the term that

[16] We tried several alternatives; in particular, we let the location and scale of G^{new} be different from those of G. In this case, $-\mathbb{E}_{G^{new}} z$ can be interpreted as the sunk cost to start a new household; alternatively, $-\mathbb{E}_{G^{new}} z$ can be thought of as the cost of a separation. However, since the parameters of the meeting function (4.5) are already numerous, adding more parameters through G^{new} resulted in a loss of precision, given that the identification of the two functions hinges on the same source of variation.

[17] Only 3.2% (3.6%) of all couples in 1971-1975 (2011-2015) displayed an age gap outside of the considered interval.

[18] In the 2011-2015 sample, only 2.35% (1.51%) of all new marriages involved a man (woman) older than 63.

depends on the partners' types $\{i, j\}$ only, as follows:

$$S_{i,j} = \alpha' x_{i,j}^{\alpha} \tag{4.7}$$

where the terms of $x_{i,j}^{\alpha}$ include dummies for the joint educational attainments and family backgrounds of the spouses, as well as first- and second-order polynomial terms of spouses' ages and wages; α is the vector of corresponding coefficients. As stated in Assumption 4, the surplus function is additively separable in the match-quality parameter z, and we assume that it follows a logistic distribution with location and scale set to 0 and 1, respectively, so that the probability that the surplus is positive is given by the following equation:

$$\alpha_{i,j} = \frac{\exp(S_{i,j})}{1 + \exp(S_{i,j})}. \tag{4.8}$$

We estimate the parameters of the meeting and surplus function with a Matching Moment estimator.[19] For every choice of λ and α, we compute the number of meetings and the matching probabilities in order to evaluate the number of new matches MF and separations DF predicted by such parameters. Our estimates $\hat{\lambda}$ and $\hat{\alpha}$ are selected such that:

$$\mathbb{E}_{\widehat{MF}} x_{i,j,k}^{\lambda} = \mathbb{E}_{MF(\hat{\lambda},\hat{\alpha})} x_{i,j,k}^{\lambda} \tag{4.9}$$

$$\mathbb{E}_{\widehat{DF}} x_{i,j,k}^{\alpha} = \mathbb{E}_{DF(\hat{\lambda},\hat{\alpha})} x_{i,j,k}^{\alpha}. \tag{4.10}$$

where the lhs corresponds to the empirical moment and the rhs to the predicted moment evaluated at $(\hat{\lambda}, \hat{\alpha})$.

As implied by equations (4.9) and (4.10), the choice of the empirical moments to match is driven by the polynomial terms of the meeting and surplus function. Many of these moments have a natural interpretation. Among others, there are the average match and separation rate, the share of new matches and separations by spouses' educational levels, men's and women's average age at the moment of the match, and so on. This estimation procedure is advantageous for at least three reasons; inference is straightforward as the estimator belongs to the GMM class, the estimation does not require us to solve the model at every iteration, and it only uses the stock-flow equations that define MF and DF, without relying on any stationarity assumption. Details are provided in Appendix A.4.

4.5. Household production

We assume the per-period payoffs are logarithmic in the agent's private budget, thereby implying $\phi^g(.) = \exp(.)$. Hence, a single agent's per-period payoff simply corresponds to $\log w$.

[19] The identification strategy outlined in section 3 would suggest using all $|\mathcal{I}^m| \times |\mathcal{I}^f|$ identifying restrictions implied by the stock-flow relationships of equation (2.19). However, when the sample size is small relatively to the number of types, many of such restrictions are left undetermined due to a "curse of dimensionality".

On the other hand, since we are in a collective setting where the household allocation is efficient, the agent's private budget corresponds to the share of the household budget allocated to her private consumption in collective decision-making (Chiappori, 1992). The sharing rule formulated in equation (2.6) effectively describes the division of resources within the household.

In a context where agents only consume a private good q^g, $g \in \{m, f\}$, and a home-produced public good $F_{i,j}(Q)$, with Q being a monetary input and F the household technology, these payoffs can be derived from the following class of utility functions:

$$u^g_{i,j}(q^g, Q; z) = \log q^g + F_{i,j}(Q) + \frac{z}{2}$$

where the shape of the production function $F_{i,j}$ is left unrestricted and is the key source of the economic gains from marriage.[20] In this type of setting, Assumption 2 (Transferable Utility) requires agents to have identical preferences about the public good. As the shock enters the utility function additively, we ensure Assumption 4 is respected. As a result, the household Pareto frontier level can be decomposed as $H_{i,j}(z) = H_{i,j} \exp(z/2)$, where $H_{i,j}$ only depends on the couple's observed type $\{i, j\}$.

We calibrate the remaining parameters and set the annual discount factor β to 0.98 (Voena, 2015; Chiappori et al., 2018a) and the bargaining parameter θ to 0.5.[21] After calibrating these parameters, we are able to determine the deterministic component of the Pareto frontier level $H_{i,j}$ for any type $\{i, j\}$, starting from the estimates of the surplus and meeting function, as detailed in Appendix A.1.

4.6. Cost of college

After estimating the parameters of the wage process, the match surplus function, and the meeting function, we can compute the gains from marriage $\Delta^m(x)$ and $\Delta^f(x)$ defined by equation (2.17). The cost function C helps rationalize educational choices described by (2.18), given the estimated returns to college. In practice, we express the cost of going to college as the sum of monetary fees and a gender-specific nonmonetary component dependent on family

[20] Alternatively, in a model of labor supply where leisure l_g is a private good, the agent's private expenditure corresponds to $q^g + w^g l^g$. Moreover, note that the public good production of singles is normalized to zero. This normalization is necessary in this class of matching models (Galichon and Salanié, 2015). In fact, $F_{i,j}(Q)$ should be understood as the utility from the public good consumption surplus generated by the match with respect to what individuals can achieve when single.

[21] GJR use information on the distribution of partners' leisure time in order to infer θ. Their estimates for a sample of British individuals aged between 22 and 50 in the 1991-2008 period suggest that θ is not significantly different than 0.5.

background x[22] and an idiosyncratic term η distributed logistically with variance σ_η^2

$$C^g(x, s) = \log(CollegeFees_s) + \delta_0^g + \delta_1^g x + \delta_2^g x^2 + \delta_3^g x^3 + \eta. \tag{4.11}$$

where the monetary fees are calculated as the in-state public college tuition dependent on the state s where the agent was raised.[23] The parameters of C are estimated with a logit regression based on the binary choice (2.18).

5. Results

5.1. Marginals, college education, and wage process

We start by describing the parameters estimated at the first step, which include gender-specific mortality rates and gender- and education-specific wage process parameters. One important demographic fact to keep in mind is that higher mortality rates among men used to lead to an important gender ratio imbalance among elder individuals in the 1970s; while this gender imbalance has partly faded, it is still present in the 2010s, as shown in panel (b) of Figure 23.

Figure 2 shows that the share of female college graduates has doubled between the 1970s and the 2010s, while the share of male college graduates has only increased by about 50%; women have overtaken men starting from the 2000s. In Appendix, Table 6 shows that the total cost of college has gradually increased; the increase is stronger for men, particularly for those with a worse family background, and milder for women.

[22]We assume the cost function C^g depends on the continuous family score introduced in Section 4.3, whereas in later stages of the life-cycle, only a discretized version of the same variable matters. This is done essentially to reduce the state space of the agents and ease calculations when solving for the marriage market equilibrium.

[23]More precisely, they are calculated as the weighted average in-state first-year undergraduate tuition and mandatory fees across public universities in a given state from the Integrated Postsecondary Education Data System. The weights are computed as the shares of enrolled first-year undergraduate students in each state public university. For every sample cut, we use data from the year opening the corresponding decade (e.g., for the 1971-1975 sample cut, we use college tuition and fees data from 1970).

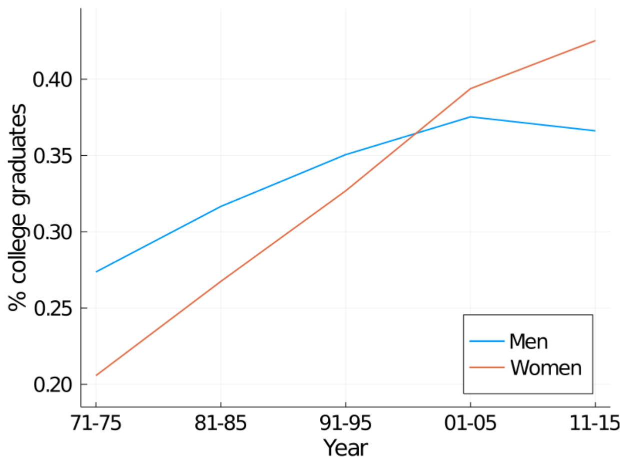

Figure 2: Share of college graduates

Changes in the wage distribution between the early 1970s and the early 2010s are well summarized by plot 3. The men's distribution has become flatter, while the women's distribution has both become flatter and shifted to the right. While men's wages around the median have stagnated, those below the median are now lower than previously. On the other hand, the 2010s women's wage distribution first-order stochastically dominates the 1970s distribution, and the gender wage gap has become narrower as measured at all percentiles.

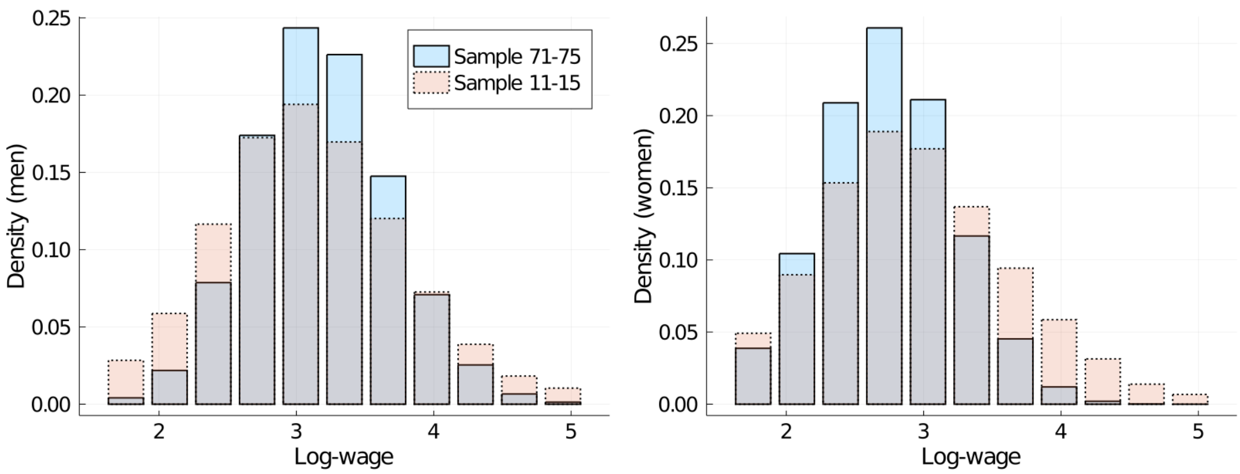

(a) Men

(b) Women

Figure 3: Full-time wage distribution

The overall increase in wage dispersion is due to both an increase in the college wage premium, an increase in college graduation rates and an increase in cross-sectional wage dispersion within each educational group. This is evident in panels (c) and (d) of Figure 25. At the same time, panels (e) and (d) of the same figure show that wage mobility has slightly decreased for both male and female college graduates, while it has increased for individuals without a college degree, particularly the youth. On the other hand, wage inequality between genders has decreased, partly due to women's college graduation rates increasing at a faster pace than men's.

5.2. Search frictions

Table 1 reports the estimated parameters of the meeting function, which correspond to the coefficients of the exogenous type-specific meeting rate $\lambda_{i,j}$; they account for both search frictions and assortativeness in the meeting function (4.5). First, we note that the constant is steady across different decades, suggesting that baseline search frictions for individuals without a college degree have not changed much over time. Having a college degree does not always help reduce frictions; female college graduates experience increasingly stronger frictions, while male college graduates, who previously suffered stronger frictions, now enjoy a slight advantage over non-graduates. Search frictions increase with age for almost all individuals, except for very young men, but they do so faster for women. While this age asymmetry is observed in all periods, it is weaker in more recent periods.

The interaction terms of $\lambda_{i,j}$ give information about sorting patterns. Our findings show that the meeting function is supermodular with respect to all observed dimensions, i.e., the levels of education, family backgrounds, age, and wages. Hence, people that are similar along these dimensions face lower search frictions and tend to have higher chances of meeting each other. The different columns of Table 1 also show that the degree of supermodularity with respect to both educational levels and wages is fairly steady over time. The degree of supermodularity with respect to age has slightly decreased, while with respect to family backgrounds, it has almost disappeared.

We can visualize how search frictions depend on agents' wages by plotting meeting rates $\lambda_{i,j}$ for different wage combinations. The plots in Figure 4 show that, in both periods, meeting rates are high along the diagonal, and thus search frictions are weaker among likes. In the 1970s, high-wage individuals, and particularly high-wage men, tend to face low frictions; this is not the case anymore in the 2010s, where meetings rates are fairly similar for individuals with different wages. In Appendix, Figure 31 contains similar plots describing the structure of search frictions with respect to age. This figure shows that search frictions are weaker along the diagonal in both periods. Interestingly, search frictions for older agents were stronger in the 1970s than in current days.

How should we interpret these findings? First of all, they reflect the way marriage markets are structured, and particularly how they are imperfectly segmented along demographic and socioeconomic dimensions. Search frictions being weaker among likes partly reflects the fact that, in everyday life, similar people have more chances to meet at school, in college, at work, in their neighborhood, or through the family and friends network. However, as suggested by Becker et al. (1977), search frictions can also capture differences in both learning and sunk costs necessary to start a new family. For instance, female college graduates facing stronger frictions in the 2010s might reflect a higher screening cost. In other words, female college graduates tend to screen fewer potential partners, which also makes them more committed to their relationships, leading to lower separation rates and long-lasting relationships.[24]

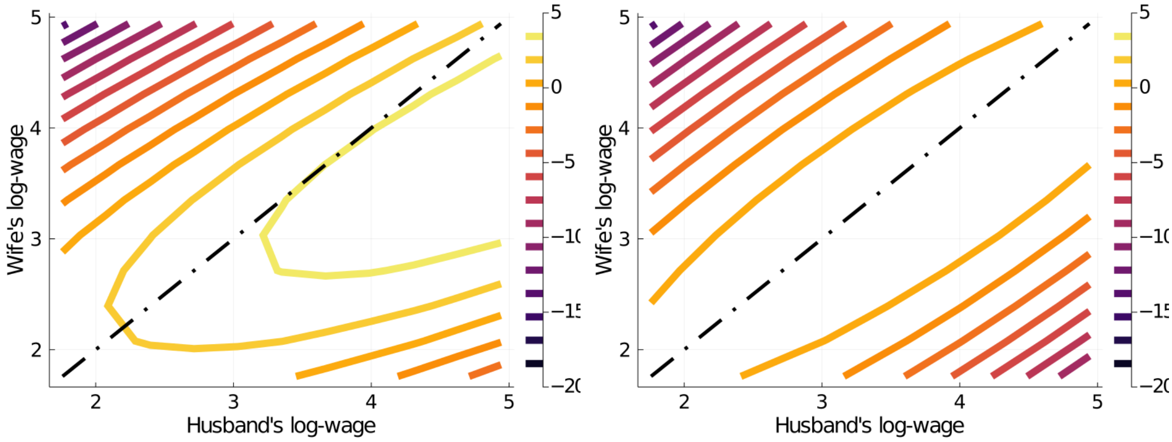

(a) 1971-1975

(b) 2011-2015

Figure 4: Estimated meeting rates $\hat{\lambda}_{i,j}$ as a function of partners' wages

Notes: the figure shows how the estimated $\hat{\lambda}_{i,j}$, the exogenous parameter of the meeting function $M_{i,j}$ that measures the strength of search frictions, changes with agents' wages. The higher $\hat{\lambda}_{i,j}$, the more likely a pair of single agents $\{i, j\}$ is to meet on the marriage market, all else constant. The average levels of $\hat{\lambda}_{i,j}$ are normalized to compare the two data periods.

5.3. Match surplus

Table 2 reports the estimated coefficients of the nonrandom component of the match surplus, $S_{i,j}$. The point estimates of the constant term suggest gains from marriage for non-graduates might be lower than in the past, but the decline is not significant. A college degree is associated with a higher match surplus in all periods, while having a better family background matters less. The importance of both male and female wages increases over time and is discussed later with the help of Figure 6.

In all periods, $S_{i,j}$ exhibits a supermodular structure with respect to the levels of education, family backgrounds, and age. The degree of complementarity is fairly constant over time for the first two variables, while it decreases for age. Figure 29 in Appendix shows that, as partners get

[24]Since we do not distinguish between legal marriage and cohabitation, differences in frictions can also reflect the differential taste for marriage over cohabitation. Categories that have a stronger preference for marriage will need to search longer, but, on the other hand, their relationships will also tend to last longer.

Table 1: Estimated coefficients of the meeting rate $\lambda_{i,j}$

Polynomial term $x^{\lambda}_{i,j}$	1971-1975 (1)	1981-1985 (2)	1991-1995 (3)	2001-2015 (4)	2011-2015 (5)
Constant	-12.14	-10.15	-11.11	-10.27	-11.71
	(0.90)	(0.64)	(0.58)	(0.72)	(0.63)
Wife has a college degree	0.06	-0.26	-0.53	-0.52	-0.39
	(0.22)	(0.18)	(0.18)	(0.20)	(0.22)
Husband has a college degree	-1.04	-0.73	-0.10	0.50	0.44
	(0.21)	(0.18)	(0.18)	(0.19)	(0.20)
Spouses have same education	2.40	1.81	2.13	1.97	2.43
	(0.16)	(0.14)	(0.14)	(0.15)	(0.17)
Wife has high family background	-0.08	-0.60	-0.59	-1.33	-1.14
	(0.18)	(0.16)	(0.16)	(0.17)	(0.17)
Husband has high family background	-0.21	-0.08	-0.33	-0.98	-0.45
	(0.17)	(0.15)	(0.16)	(0.17)	(0.17)
Spouses have same family background	1.04	0.70	0.11	-0.09	0.22
	(0.12)	(0.11)	(0.11)	(0.13)	(0.13)
Husband's age	0.45	0.30	0.45	0.36	0.34
	(0.08)	(0.07)	(0.08)	(0.08)	(0.08)
Wife's age	-0.85	-0.34	-0.54	-0.45	-0.37
	(0.08)	(0.07)	(0.08)	(0.09)	(0.09)
Husband's age (sq.)	-0.24	-0.15	-0.18	-0.17	-0.20
	(0.02)	(0.01)	(0.01)	(0.01)	(0.01)
Wife's age (sq.)	-0.26	-0.18	-0.18	-0.18	-0.22
	(0.02)	(0.01)	(0.01)	(0.01)	(0.02)
Husband's age × wife's age	0.51	0.32	0.35	0.34	0.41
	(0.04)	(0.02)	(0.02)	(0.03)	(0.03)
Husband's log-wage	0.22	0.25	0.45	0.39	0.18
	(0.27)	(0.20)	(0.18)	(0.21)	(0.18)
Wife's log-wage	0.77	-0.37	-0.05	-0.03	0.37
	(0.30)	(0.23)	(0.19)	(0.21)	(0.20)
Husband's log-wage (sq.)	-0.07	-0.16	-0.16	-0.12	-0.13
	(0.03)	(0.02)	(0.02)	(0.02)	(0.02)
Wife's log-wage (sq.)	-0.21	-0.17	-0.17	-0.12	-0.17
	(0.04)	(0.03)	(0.03)	(0.02)	(0.03)
Husband's log-wage × wife's log-wage	0.21	0.35	0.31	0.22	0.26
	(0.04)	(0.03)	(0.03)	(0.03)	(0.03)

Notes: The table reports estimates of the coefficients λ of the $M_{i,j}$ function, which gives the number of meetings between types $\{i, j\}$ and whose specification is given by (4.5). The estimates $\hat{\lambda}$ are obtained with the Matching Moment estimator described in section 4.4. Standard errors in parentheses.

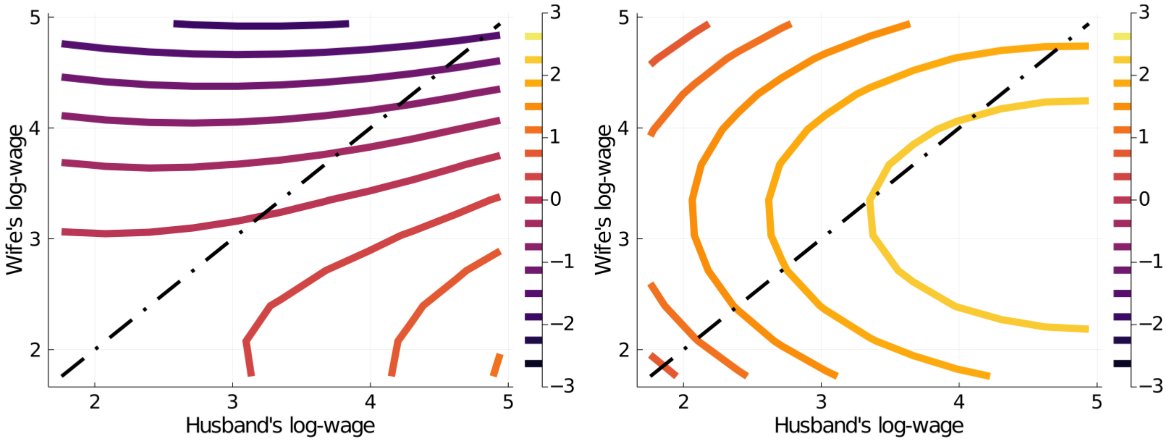

(a) 1971-1975

(b) 2011-2015

Figure 6: Estimated surplus $\hat{S}$ as a function of partners' wages

Notes: the picture shows how the estimated $\hat{S}_{i,j}$, the deterministic component of the surplus function common to all pairs $\{i, j\}$, changes with agents' wages. The higher $\hat{S}_{i,j}$, the larger the average gains from marriage for a couple $\{i, j\}$ (and the more stable the relationship). The average levels of $\hat{S}_{i,j}$ are normalized to compare the two data periods.

older, their match surplus increases and their separation risk declines, which is largely explained by the partners' reservation values decreasing with age. This pattern is stronger in the 1970s than in the 2010s, which is consistent with the increase in separation rates among older couples.

Importantly, we find that wages are neither complements nor substitutes in the surplus function; the interaction term is very close to zero in all periods. Hence, after controlling for other observed characteristics, the residual positive correlation between partners' wages among new matches is driven by search frictions only. Yet, Figure 6 shows that the match surplus is highest for couples formed by a high-wage husband and a low-wage wife. This gender-asymmetric pattern favors "traditional" specialized couples, where the husband has a higher wage than the wife. In the 2010s, we observe a new emerging pattern: $S_{i,j}$ increases in the wife's wage up to a certain threshold, which approximately corresponds to the third wage quartile, and then decreases. Compared with the past, matches where both partners have a high wage are now more successful. However, gender asymmetries persist; the most successful couple is the one formed by a top-wage husband and a wife with a medium-high wage. The origins of this gender asymmetry are further discussed in the next section.

5.4. Household production

Starting from our estimates of the match surplus, we recover $\log H_{i,j}$, the log-level of the Pareto frontier; the larger $\log H_{i,j}$, the greater the static component of the gains from marriage. While we cannot distinguish between its different components, $\log H_{i,j}$ encompasses the match gains coming from economies of scale in consumption (e.g., sharing the rent and the bills) and household production (e.g., domestic chores and childrearing), as well as non-economic

Table 2: Estimated coefficients of the surplus function $S^*_{i,j}$

Polynomial term $x^{\alpha}_{i,j}$	1971-1975 (1)	1981-1985 (2)	1991-1995 (3)	2001-2015 (4)	2011-2015 (5)
Constant	1.15	0.65	0.22	-0.34	-0.98
	(0.62)	(0.35)	(0.37)	(0.44)	(0.41)
Wife has a college degree	0.25	0.53	0.60	0.61	0.28
	(0.18)	(0.12)	(0.12)	(0.13)	(0.13)
Husband has a college degree	0.74	0.32	0.26	0.40	0.53
	(0.20)	(0.13)	(0.12)	(0.14)	(0.13)
Spouses have same education	0.29	0.38	0.21	0.23	0.24
	(0.17)	(0.12)	(0.11)	(0.13)	(0.12)
Wife has high family background	0.16	0.27	0.03	0.03	0.16
	(0.12)	(0.09)	(0.09)	(0.10)	(0.11)
Husband has high family background	0.29	0.20	0.01	0.17	-0.00
	(0.12)	(0.09)	(0.09)	(0.11)	(0.11)
Spouses have same family background	0.31	0.04	0.28	0.27	0.22
	(0.11)	(0.09)	(0.08)	(0.10)	(0.10)
Husband's age	0.24	0.14	0.03	0.03	-0.02
	(0.08)	(0.07)	(0.06)	(0.07)	(0.08)
Wife's age	-0.23	-0.02	0.14	0.05	0.11
	(0.09)	(0.07)	(0.06)	(0.07)	(0.07)
Husband's age (sq.)	-0.03	-0.03	-0.03	-0.03	-0.01
	(0.01)	(0.01)	(0.01)	(0.01)	(0.01)
Wife's age (sq.)	-0.01	-0.03	-0.04	-0.03	-0.02
	(0.01)	(0.01)	(0.01)	(0.01)	(0.01)
Husband's age × wife's age	0.05	0.07	0.06	0.06	0.03
	(0.02)	(0.02)	(0.01)	(0.02)	(0.02)
Husband's log-wage	0.02	0.13	0.07	0.16	0.30
	(0.14)	(0.10)	(0.10)	(0.11)	(0.11)
Wife's log-wage	0.14	0.02	0.01	0.31	0.34
	(0.20)	(0.12)	(0.11)	(0.13)	(0.12)
Husband's log-wage (sq.)	0.01	-0.01	-0.00	-0.01	-0.01
	(0.01)	(0.01)	(0.01)	(0.01)	(0.01)
Wife's log-wage (sq.)	-0.03	-0.01	-0.01	-0.02	-0.03
	(0.03)	(0.02)	(0.01)	(0.01)	(0.01)
Husband's log-wage × wife's log-wage	-0.01	-0.01	0.02	-0.01	-0.00
	(0.03)	(0.02)	(0.01)	(0.02)	(0.02)

Notes: The table reports estimates of the coefficients α of the $\alpha_{i,j}$ function, which gives the odds that a couple will continue a relationship. The estimates $\hat{\alpha}$ are obtained with the Matching Moment estimator described in section 5. Standard errors in parentheses.

components dependent on socioeconomic and demographic traits (e.g., non-economic benefits from age proximity or from sharing a similar socioeconomic background).

Table 3 shows that having a college degree results in a higher $\log H_{i,j}$, thereby expanding the couple's consumption set. On the other hand, sharing a similar educational level does not affect $\log H_{i,j}$ significantly. Yet, Table 2 indicates that $S_{i,j}$ is supermodular in education; where does this supermodularity come from? If the educational levels are not complements in $\log H_{i,j}$, then the supermodularity of $S_{i,j}$ comes from its dynamic component, and is explained by the mutual insurance mechanism implied by the marriage contract. In other words, if both partners have a college degree, they *both* benefit from better wage prospects. They anticipate that their match will be a stable one and will likely result in satisfactory household consumption levels; in the Bellman equations (2.13) and (2.14), this corresponds to the educational levels being complements in the match continuation value (the third and last term).

The last line of Table 3 shows that, consistent with the theory of Becker (1973), full-time wages are substitutes in household production, a finding that is robust across periods. Hence, couples do have an incentive to specialize. However, wages are substitutes in $\log H_{i,j}$, but not in $S_{i,j}$. Since wages are persistent, they are instead complements in the match continuation value, similarly to what was observed for education. This feature highlights the importance of mutual insurance in marriage. While spouses do benefit from household specialization, they also value the joint presence of two earners who can both contribute to keeping a high level of household consumption and thus insure the household against income risk.

In Table 3, we can see that $\log H_{i,j}$ is indeed increasing in both partners' wages almost everywhere, which is not surprising as a higher full-time wage usually results in better consumption opportunities for the family. Figure 8 illustrates this pattern, but also shows that *who* earns the wage matters. In the 1970s, there is a strong gender asymmetry in household production; couples with a high-wage husband and a low-wage wife enjoy much larger economies of scale than couples with a low-wage husband and a high-wage wife. This pattern is so strong that, when the husband's wage largely exceeds the wife's (in the bottom-right corner), an increase in the wife's wage can result in an output loss. This gender asymmetry has largely, but not completely, faded by the 2010s. While the figure looks more symmetric and couples with the highest household full-time labor income now enjoy the best consumption opportunities, switching wages between partners when the husband outranks the wife still results in an output loss. This pattern is reminiscent of the findings of Bertrand et al. (2015) on household gender norms; deviating from the behavioral prescription stating that the husband should earn more than the wife can result in a utility loss. Finally, in the 2010s, the surplus $S_{i,j}$ starts decreasing when the wife's wage is too high; however, $\log H_{i,j}$ is strictly increasing in her wage. Considering

Table 3: Joint household production as measured by the Pareto frontier log-level $\log H_{i,j}$

Polynomial term	1971-1975 (1)	1981-1985 (2)	1991-1995 (3)	2001-2015 (4)	2011-2015 (5)
Constant	4.81	4.38	4.41	4.24	3.67
	[3.70,5.91]	[3.80,4.94]	[3.79,5.04]	[3.62,4.98]	[2.99,4.33]
Wife has a college degree	0.18	0.23	0.22	0.33	0.23
	[-0.10,0.49]	[0.07,0.40]	[0.05,0.39]	[0.15,0.51]	[0.06,0.39]
Husband has a college degree	0.20	0.06	0.08	0.23	0.20
	[-0.01,0.41]	[-0.07,0.20]	[-0.09,0.27]	[0.08,0.37]	[0.01,0.37]
Spouses have same education	0.02	0.04	0.00	0.01	0.01
	[-0.01,0.06]	[0.01,0.08]	[-0.03,0.04]	[-0.02,0.04]	[-0.02,0.05]
Wife has high family background	-0.04	-0.14	-0.30	-0.55	-0.40
	[-0.25,0.18]	[-0.28,0.01]	[-0.43,-0.16]	[-0.70,-0.40]	[-0.54,-0.25]
Husband has high family background	-0.03	0.02	-0.17	-0.29	-0.21
	[-0.15,0.10]	[-0.08,0.11]	[-0.31,-0.07]	[-0.42,-0.16]	[-0.32,-0.10]
Spouses have same family background	0.03	0.00	0.02	0.00	0.01
	[0.01,0.05]	[-0.02,0.02]	[0.00,0.04]	[-0.03,0.03]	[-0.01,0.04]
Husband's age	0.10	0.11	0.06	0.05	0.02
	[0.03,0.16]	[0.05,0.16]	[0.01,0.10]	[0.01,0.11]	[-0.03,0.07]
Wife's age	-0.08	0.02	0.05	0.02	0.04
	[-0.13,-0.03]	[-0.03,0.06]	[0.01,0.09]	[-0.02,0.06]	[-0.01,0.08]
Husband's age (sq.)	-0.01	-0.01	-0.01	-0.01	-0.00
	[-0.01,-0.00]	[-0.01,-0.01]	[-0.01,-0.01]	[-0.01,-0.00]	[-0.01,-0.00]
Wife's age (sq.)	-0.00	-0.01	-0.01	-0.01	-0.00
	[-0.00,0.00]	[-0.01,-0.00]	[-0.01,-0.00]	[-0.01,-0.00]	[-0.01,-0.00]
Husband's age × wife's age	0.00	0.01	0.01	0.01	0.00
	[-0.00,0.01]	[0.00,0.02]	[0.00,0.01]	[0.00,0.01]	[-0.00,0.01]
Husband's log-wage	0.24	0.34	0.33	0.42	0.48
	[-0.01,0.49]	[0.17,0.50]	[0.17,0.49]	[0.25,0.59]	[0.32,0.65]
Wife's log-wage	0.54	0.35	0.36	0.48	0.57
	[0.19,0.89]	[0.18,0.51]	[0.19,0.51]	[0.29,0.67]	[0.39,0.75]
Husband's log-wage (sq.)	0.04	0.03	0.02	0.02	0.02
	[0.02,0.07]	[0.01,0.04]	[0.01,0.04]	[0.00,0.04]	[-0.00,0.03]
Wife's log-wage (sq.)	-0.01	0.01	0.01	0.01	-0.00
	[-0.05,0.04]	[-0.01,0.03]	[-0.00,0.03]	[-0.01,0.03]	[-0.02,0.02]
Husband's log-wage × wife's log-wage	-0.05	-0.05	-0.03	-0.05	-0.04
	[-0.10,-0.01]	[-0.07,-0.02]	[-0.05,-0.01]	[-0.08,-0.02]	[-0.06,-0.02]

Notes: The table reports the results from projecting the estimates of $\log H_{i,j}$ using the same polynomial terms in the surplus function $S_{i,j}$. We report two-sided 95% confidence interval.

the incentives for household specialization and gender asymmetries discussed in this paragraph, this can be explained as follows; after a certain threshold, the wife's reservation value grows faster (in her wage) than her contribution to the match in terms of consumption opportunities, thereby reducing the surplus and the match stability.

Finally, Figure 33 in Appendix compares the levels of household production by spouses' ages and decade. In the 1970s, couples where the wife is very young and has an older husband enjoy comparatively larger economies of scale. In the 2010s, couples formed by two 30-year-old spouses have, *ceteris paribus*, the largest economies of scale. This evolution is consistent with the changes in the dynamics of women's reproductive capital, and particularly by the improvements in birth control and assisted reproductive technology (Low, 2014).

(a) 1971-1975

(b) 2011-2015

Figure 8: Estimated household production $\widehat{\log H}$ as a function of partners' wages

Notes: The figure shows how the estimated $\widehat{\log H}_{i,j}$, the deterministic component of the household Pareto frontier log-level for a couple $\{i, j\}$, changes with agents' wages. The higher $\widehat{\log H}_{i,j}$, the better the consumption opportunities for a couple $\{i, j\}$. The average levels of $\widehat{\log H}_{i,j}$ are normalized to compare the two data periods.

5.5. Model's fit

For every time period, the estimation of the main parameters of the model targets marriage and divorce flows over a 5-year time span (e.g., from 1971 to 1975) without relying on any stationarity assumption, and only using the identifying restrictions implied by the stock-flow equations. For every time period, we then compute the steady-state distribution of couples and singles implied by our estimated parameters. Hence, whether it can correctly predict the cross-sectional distribution of couples and singles observed in the data also depends on how close to its steady-state equilibrium the marriage market actually is.

Figure 10 shows the model's fit for aggregate separation rates, which have increased by about 50% over the last 50 years and have historically been higher for individuals without a college degree, a gap that has been widening over time. Figure 35 shows that the model fits the age profile of separation risk extremely well in all periods. In particular, it replicates the rise of separation rates at different ages; rates for the 30-year-old increased from 6% to 10%, while for the 50-year-old, they increased from only 1% to 5%. Finally, Figure 36 shows that a higher husband's wage is associated with a lower separation rate both in the fitted model and the data. On the other hand, the wife's wage profiles of separation rates are comparatively flat in the 1970s and the 1980s. In more recent decades, separation rates are instead decreasing in the wife's wage.

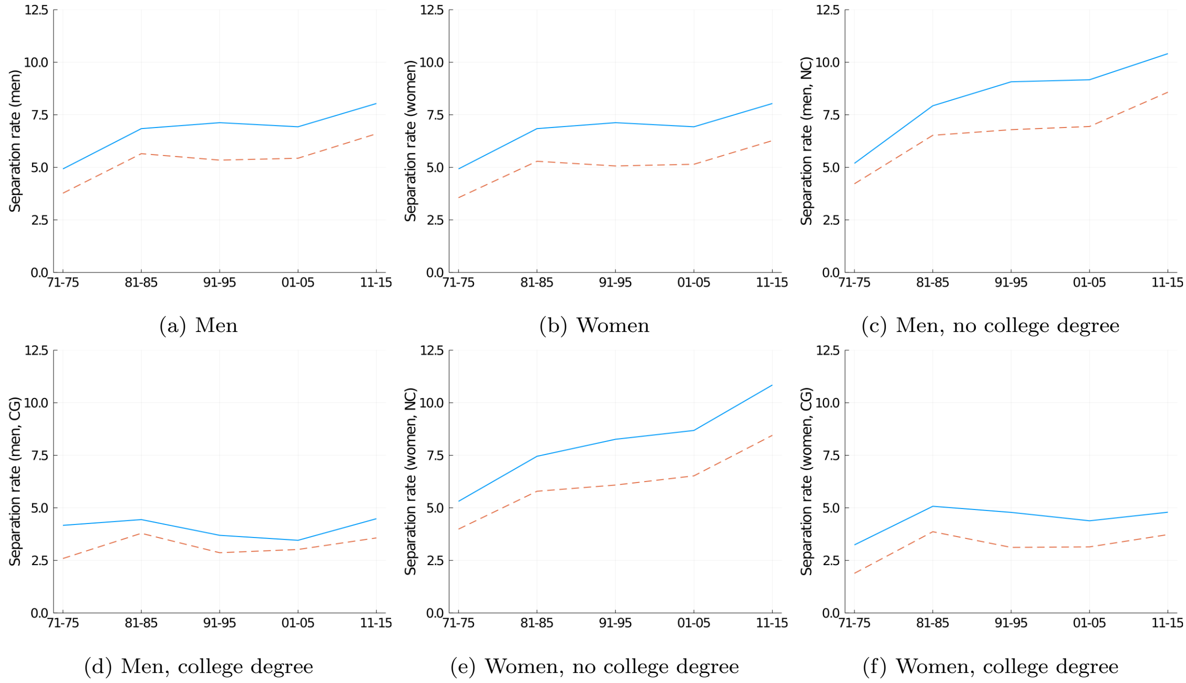

(a) Men (b) Women (c) Men, no college degree

(d) Men, college degree (e) Women, no college degree (f) Women, college degree

Figure 10: Separation rates by gender and education; model vs data

Notes: Solid blue lines correspond to data moments, dashed red lines to simulated moments. Separation rates are calculated as the ratio of the number of separations over the number of matched working-age individuals.

At the same time, match rates have also declined in the aggregate, as shown in Figure 12. Differences between individuals with and without a college degree are small, suggesting that separation rates are more important in explaining the differences in stocks across educational groups. Figure 38 shows that women have higher match rates than men when very young, while men have slightly higher match rates starting from their early 30s. The figure also clarifies that the aggregate decline is mainly due to a decrease in match rates among the youth. On the other hand, the gender gap in match rates among elder individuals has shrunk.

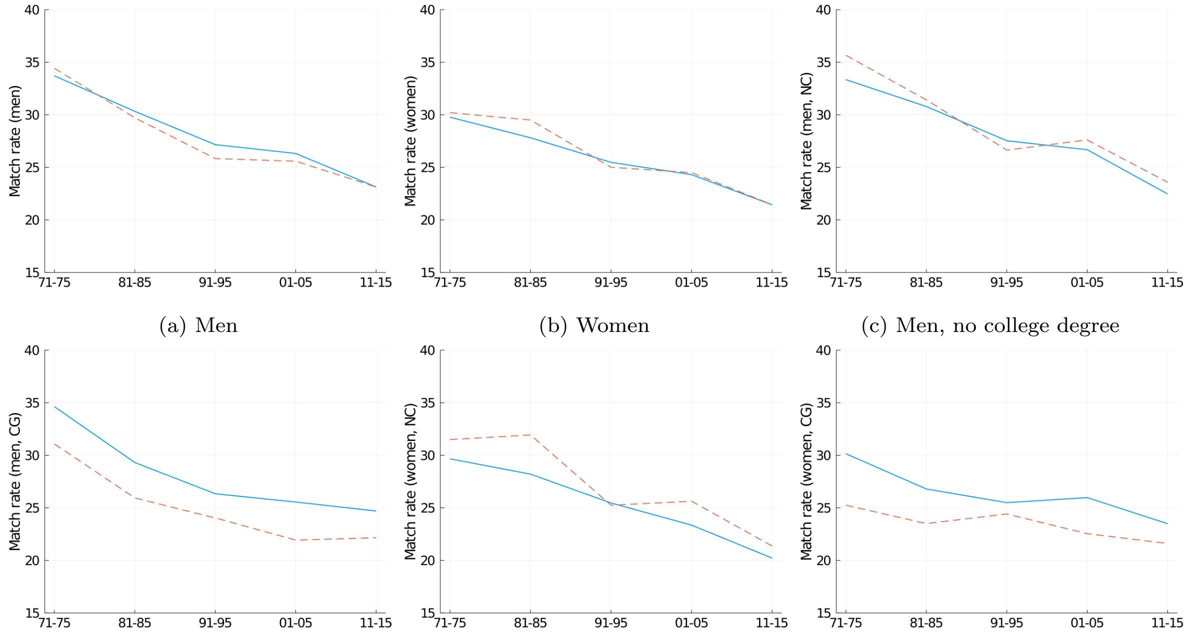

(a) Men (b) Women (c) Men, no college degree

(d) Men, college degree (e) Women, no college degree (f) Women, college degree

Figure 12: Match rates by gender and education; model vs data

Notes: Solid blue lines correspond to data moments, dashed red lines to simulated moments. Match rates are calculated as the ratio of the number of new matches over the number of working-age singles.

Figure 14 compares between the factual and predicted evolution of the share of matched individuals by gender and education. The share of individuals living with a partner has decreased, with the decrease being much steadier and stronger among both men and women without a college degree. Although these moments are not targeted in the estimation, the model's fit is excellent in more recent decades, suggesting that the marriage market outcome is well represented by the steady-state search equilibrium defined in Section 2. The fit is slightly worse for the 1970s, but Figure 40 shows that the model's fit of the age profiles of the share of matched individuals is good also in earlier decades. The only feature that the model struggles to replicate in the 1970s and 1980s is the rapid decline in the share of matched women after the age of 50. This could be explained by specific non-stationary demographic patterns characterizing those periods, such as higher male mortality among older cohorts due to the mid-century wars and the baby boom. In Figure 14, we also see that the estimated model tends to predict a higher share of matched college graduates in the 1970s. Yet, Figures 37 and 39 show that the age profiles of the transition rates are very well fitted in both the 1970s and in the 2010s. The data vs model mismatch in the shares of matched college graduates is thus likely due to non-stationary dynamics of the population composition, such as stark differences in the monetary and non-monetary cost of college across cohorts. In Appendix, we present another comparison to validate our model. Figure 42 shows that, when the model is fitted with data on legal marriages only,[25] we are able to replicate the age profiles of men's and women's marital histories, and the model correctly predicts the share of never married, married-once and remarried individuals by age.

Figure 16 shows that the model can replicate sorting patterns with respect to partners' education. Couples where both partners are college graduates more than doubled in 50 years, and couples where only the wife holds a college degree more than tripled. The model also replicates the very high correlation rate between partners' ages. On the other hand, while it can replicate the increase in the correlation between partners' wages across decades, it predicts a lower level in all periods. Blundell et al. (2016b) find that partners' wage shocks might display a low positive correlation (in the 1999-2009 period), a feature that we rule out in our model, but could potentially explain why our predicted wage correlation is lower than the observed.

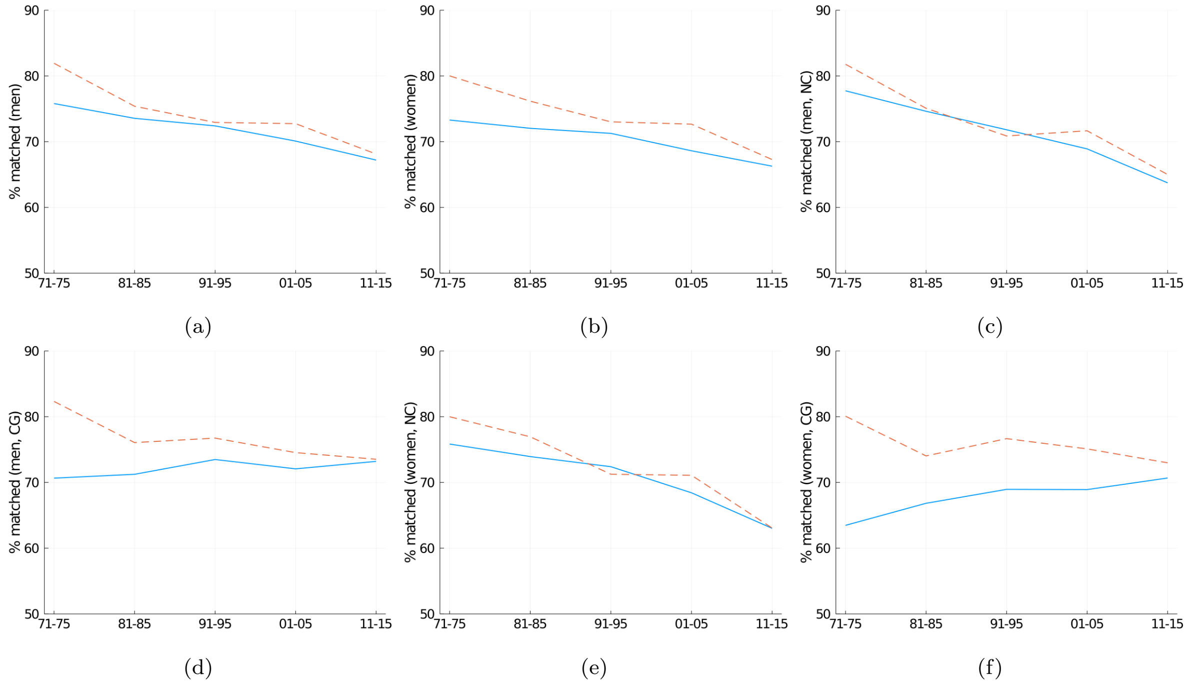

Figure 14: Share of matched individuals by gender and education; model vs data

Notes: Solid blue lines correspond to data moments, dashed red lines to simulated moments. The shares of matched individuals are calculated as the ratio of the number of matches involving a person of a given type over the total number of people of the same type in the population.

[25]Starting from 1985, the PSID survey has collected retrospective information on marital histories; these data do not cover cohabiting relationships, but only legal marriages. Therefore, we are not able to determine the number of past relationships of an individual, but only the number of her past legal marriages.

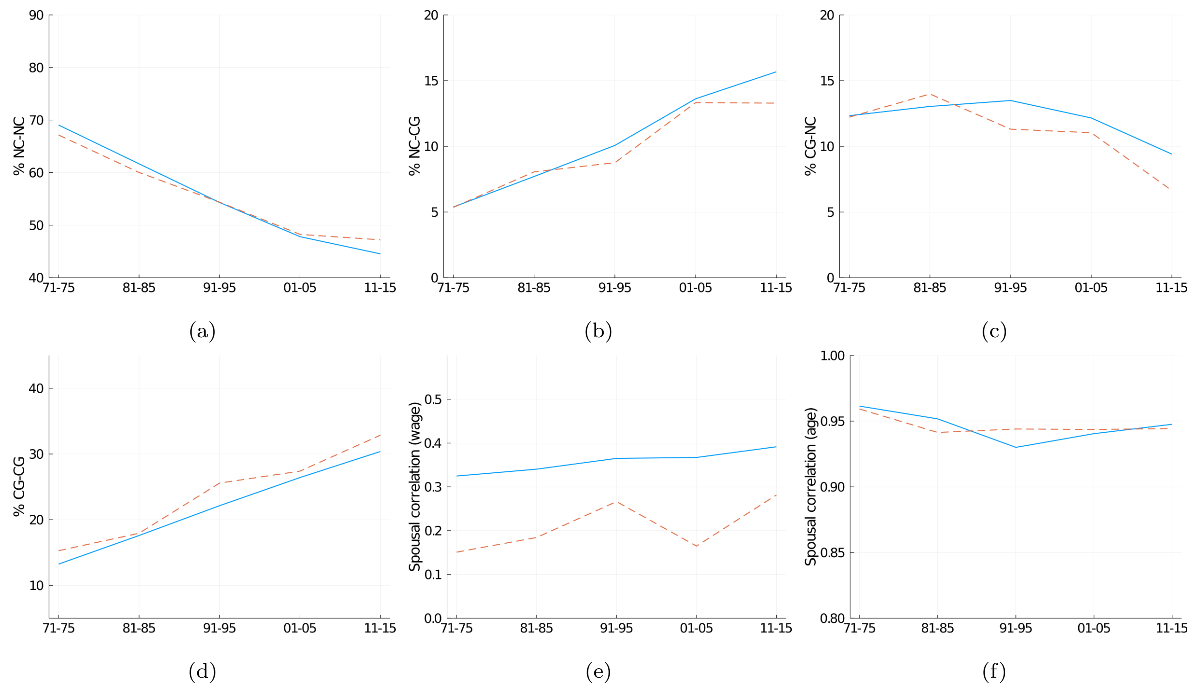

Figure 16: Sorting patterns with respect to education, age and wages; model vs data

Notes: Solid blue lines correspond to data moments, dashed red lines to simulated moments. Panels (a) to (d) report the proportion of working-age married couples with a certain educational level: NC stands for "no college" and CG for "college degree". Panels (e) and (f) report the age and wage correlation rates respectively.

Finally, our estimated model can be used to derive predictions about the sharing rule, which can be calculated using expression (2.6). Since we do not estimate the bargaining parameter θ, but we set it to 0.5 beforehand, we use time-use data from the PSID in order to show that this calibration does not produce implausible predictions for the sharing rule. From the PSID survey, we compute the average number of leisure hours for the head and his/her partner as the difference between the total hours in a week and the sum of hours worked outside and inside the household.[26] We then compute the monetary value of leisure time by multiplying idle hours by the hourly wage. In Figure 18, we compare the distribution of the sharing rule among the married population as predicted by the model with the empirical distribution of the ratio of the male partner's monetary value of leisure time to the sum of the total monetary value of leisure time in the household. The model tends to predict a distribution of the sharing rule that is more skewed in favor of men, particularly in the 1970s and the 1980s, while in more recent waves, the predicted distribution is remarkably close to the empirical, with extremely well fitted first-order moments. On one hand, the discrepancies between the two distributions

[26]Information on hours spent on domestic chores is available yearly for the head and the partner for all waves except for 1975. Leisure time is the only assignable private consumption item that can be consistently measured across all waves. Starting from 1997, consumption data provide a detailed description of household expenditure but do not assign consumption to family members.

in earlier decades suggest that θ could actually be slightly lower than 0.5; i.e., women might have had more bargaining power than men. On the other hand, alternative empirical measures of the sharing rule accounting for other private consumption items (clothing, personal hobbies, smoking, etc.) could also explain the discrepancy.

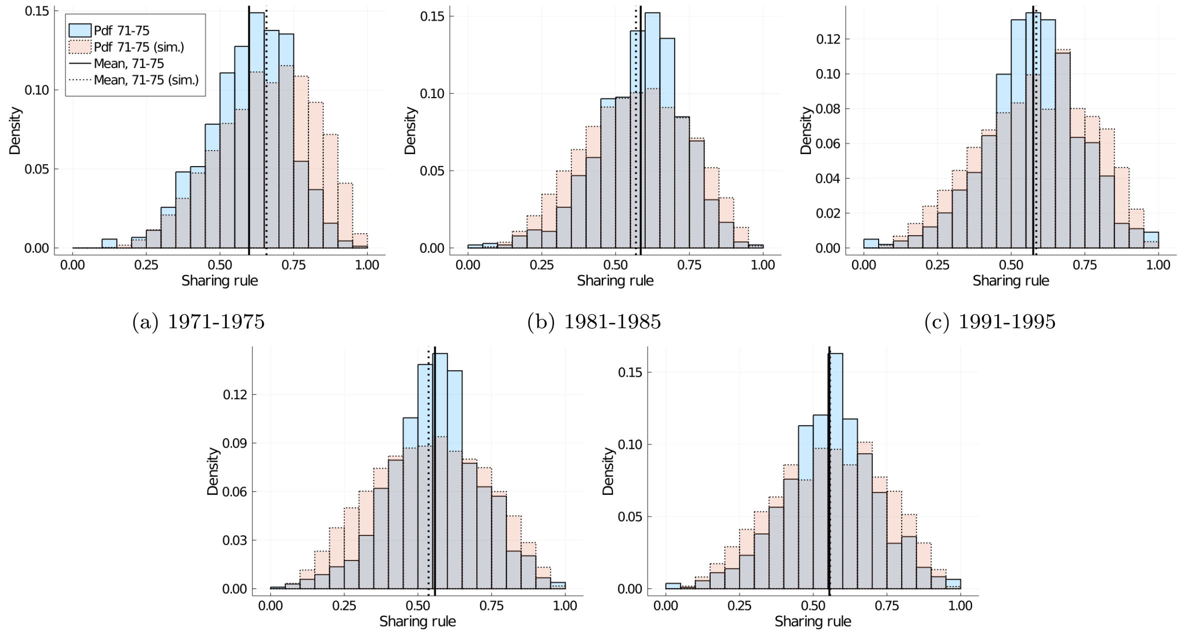

(a) 1971-1975 (b) 1981-1985 (c) 1991-1995

(d) 2001-2005 (e) 2011-2015

Figure 18: Sharing rule; model vs data

Notes: The blue histogram depicts the distribution of the sharing rule among married couples observed in the data, while the orange histogram depicts the distribution implied by the fitted model steady-state equilibrium. The solid black line corresponds to the average observed sharing rule in a data period, while the dashed black line corresponds to the average sharing rule predicted by the model.

5.6. Returns to college

In this section, we take advantage of the life-cycle dimension present in our model to quantify changes in lifetime consumption value between the 1970s and the 2010s. Worse marriage prospects can entail welfare losses due to foregone economies of scale. Moreover, risk-averse agents will dread a higher separation risk. This aspect is particularly critical because low-wage men typically experience more unstable relationships. While they wish to be insured against wage risk through marriage, some large wage shocks will be particularly hard to insure and will lead to their partners quitting the match.

Figure 20 shows how the present discounted value of expected lifetime log-consumption for agents aged 17 has changed between the 1970s and the 2010s for different educational groups. Female college graduates are the only group that has experienced welfare gains (+8.5%). Male college graduates and female non-graduates have experienced relatively small losses (-2% and

-4.5%, respectively), while men without a college degree have suffered from large losses (-10%). The difference in lifetime consumption between educational groups yields the overall returns to college, also defined in equation (2.17). Figure 21 shows that, on average, returns to college have increased by 82% for men and 138% for women.

We decompose these welfare changes into two components and compute the present discounted value of being single throughout the entire life-cycle, conditionally on family background x and educational level h, as follows:

$$\mathbb{E}\left[\sum_{a=19}^{84} \beta^{a-19} \log w_a \middle| h, x\right] \tag{5.1}$$

where w_a denotes the wage at age a. Changes in this welfare component reflect changes in the wage structure and mortality rates and are represented by the red bars in Figure 20. All groups have experienced improvements, except for men without a college degree, whose wage prospects have worsened dramatically between the two periods.[27]

At the same time, the expected welfare from matching on marriage markets, represented by the blue bars in Figure 20, has decreased for all categories. Individuals without a college degree have experienced the largest marriage market losses due to delayed household formation and increased marital instability. These foregone gains from marriage have compounded men's worsened wage prospects and resulted in men suffering from large overall welfare losses. Their impoverishment also has important implications for women without a college degree, whose typical match will be a man from the lower part of the wage distribution, since meetings are positive assortative (also with respect to wages) and traits that positively correlate with wages (age and education) are complements in the surplus function. On one hand, in the 2010s, better wages allow less educated women to be better insured against separation risk, demand a more favorable sharing rule within their household, and stabilize their relationships where they are now more likely to be the primary earner. On the other hand, many of these women will live in households where the economies of scale can hardly justify the high income risk associated with their male partners' worsened wage prospects. In many cases, the husband's low wage levels and the threat of a future negative wage shock will scare women away from the current match.

In contrast, both male and female college graduates have suffered from relatively smaller marriage market losses. While they have futher delayed household formation, most of them eventually engage in long-lasting relationships also in more recent periods, as their separation

[27]As we move from the 1970s to the 2010s, a better life expectancy had a positive impact on the welfare of all agents. However, it does not differ across educational groups in our calibration, so we rarely mention it in the discussion.

risk has stayed relatively low. In the case of female college graduates, improved wage prospects have generated welfare gains that are sufficiently large to exceed losses due to foregone marriage gains. Single women with a college degree are now better prepared to deal with singlehood, both before the first match and following a separation. Moreover, weaker gender asymmetries also mean that women can actively increase their families' consumption opportunities and provide insurance against income risk through labor market participation. Hence, investments in human capital pay off on marriage markets in the form of more stable and productive matches. As a result, in Figure 21, a non-negligible part of the increase is explained by higher marriage market returns to college for both men (19 out of a total 82pp increase) and women (24 out of 138pp).

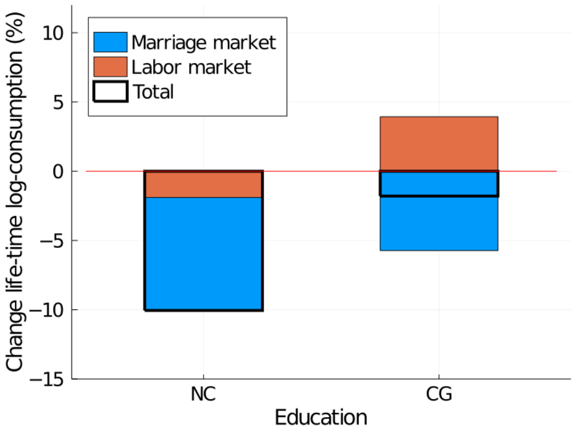

(a) Men

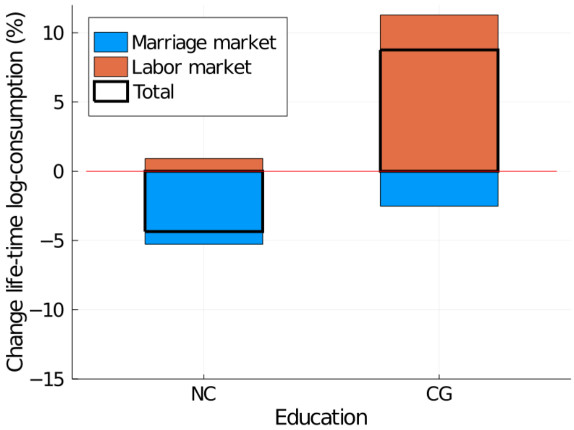

(b) Women

Figure 20: Changes in lifetime log-consumption between the 1970s and the 2010s

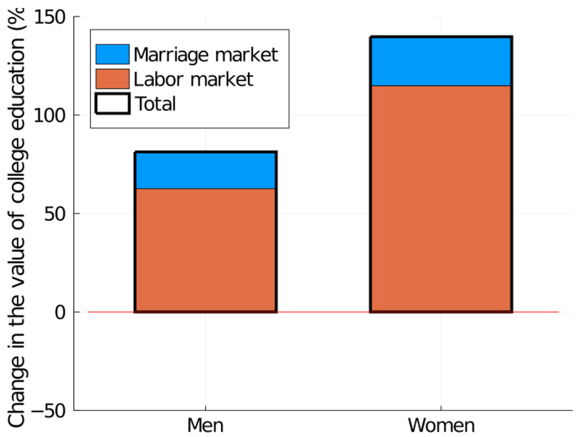

Figure 21: Changes in returns to college between the 1970s and the 2010s

6. Counterfactuals

6.1. Isolating the impact of changing wages

After estimating the model with samples from different periods, we perform several comparative statics exercises in order to understand how the marriage market and the returns to college education react to changes in parameters that are exogenous in our model. We start by isolating the impact of changes in the parameters of the wage process between the 1970s to the 2010s. We simulate a counterfactual steady-state equilibrium where all parameters are taken from our estimates for the 1970s except for those of the wage process, which are taken from the 2010s.

Table 4: Counterfactual experiment, the 1970s steady-state equilibrium with wages from the 2010s

	Men			Women		
Changes in...	All (1)	NC (2)	CG (3)	All (4)	NC (5)	CG (6)
Share college graduates (pp)	21.42			20.31		
Share matched (pp)	-2.38	-3.70	-1.24	-2.12	0.65	-6.17
Share matched, aged 21 to 30 (pp)	-6.10	-6.01	-3.41	-4.03	-2.04	-3.06
Share matched, aged 31 to 40 (pp)	-0.67	-3.26	-0.18	-0.47	0.36	-4.64
Share matched, aged 41 to 50 (pp)	-0.69	-2.87	-0.49	-0.83	2.04	-7.94
Share matched, aged 51 to 60 (pp)	-1.07	-2.62	-1.07	-1.19	2.97	-9.84
Separation rate (pp)	-0.15	0.47	-0.05	-0.15	-0.06	0.69
Match rate (pp)	-4.32	-4.55	-2.15	-3.40	1.30	-5.24

Notes: We report changes in different moments of the matching distribution and educational achievements between two steady-state equilibria. The first steady-state is the one implied by the parameters estimated with data from the 1971-1975 period. The second is the counterfactual steady-state equilibrium implied by all the parameters for the 1971-1975 period, except for those of the wage process, which take on the values estimated with data from the 2011-2015 period. NC means "no college"', "CG" means "college graduates".

Table 4 shows that changes in wages induce a large expansion of college education, although, contrarily to the data, the increase is of similar size for both men and women (resp. 21.42 and 20.31 percentage points). While women have enjoyed a larger increase in the labor market returns to college, changes in wages alone cannot explain why women's college graduation rates have outpaced men's. Figure 22 shows that changes in wages would result in decreasing women's marriage market returns to college when isolated from changes in parameters capturing household technology, gender norms, and marriage market frictions. In other words, as long as higher wages do not pay off on marriage markets, women face a strong trade-off between family and career, and female college graduation rates cannot grow as fast as observed in the data.

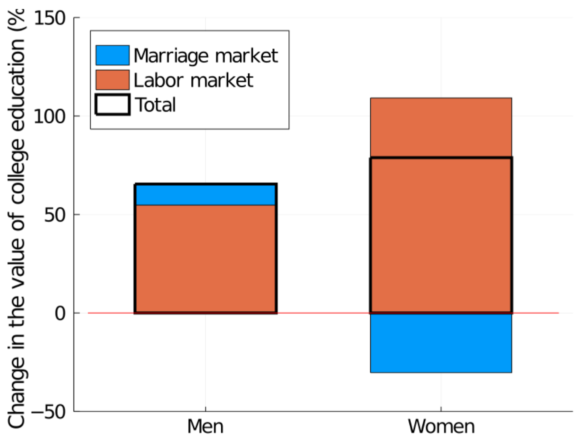

Figure 22: Counterfactual experiment, the 1970s steady-state equilibrium with wages from the 2010s

In the counterfactual, female college graduates face stronger competition on marriage markets, as their numbers grow while their typical profile becomes less attractive due to their increased wages. Table 4 shows that their match rates drop while their separation risk increases. Meanwhile, the correlation between partners' wages increases by 11 percentage points. Changes in the marginal distributions lead to more meetings (and matches) among people with similar wages, thus amplifying the relative importance of the mutual insurance aspect of marriage. On the other hand, these more "egalitarian" matches are not as attractive as the "traditional" specialized ones (i.e., with a male breadwinner) due to the gender asymmetries discussed in Section 5.4. Men still seek women with a wage lower than theirs and often "marry down"; as a result, the share of couples where the husband is more educated than the wife increases from 12.2% to 16.6% of all couples, even though non-graduate women are scarcer.[28]

Changes in wages lead to a 2.4pp increase in the share of adult singles, corresponding to about 17% of the total increase observed between the 1970s and the 2010s. The increase in the share of singles is concentrated among the youth and household formation is delayed. While this can partly be explained by steeper wage profiles, agents also have incentives to search longer due to fewer opportunities for "traditional" household specialization. Similarly to what was suggested by Stevenson and Wolfers (2007), when the economic motives for marriage are

[28]We conduct additional counterfactual exercises where we shift only some parameters of the wage process; these findings are available on request. Consistent with the mechanism explained in the main text, changes in men's mean wages result in better match prospects for male and female college graduates, but worse prospects for the less educated. On the other hand, changes in women's mean wages result in much worse match prospects for female college graduates, whose improved wages constitute a hindrance to marriage. The implications of increased wage dispersion are more nuanced; it leads to higher separation rates as wage shocks are more sizable, but also to higher match rates as agents are more willing to insure themselves against income risk. Finally, we do not find that changes in wage mobility, as measured by the Spearman's wage (rank) autocorrelation, had any large impact.

weaker, more couples will match based on love and companionship only; finding a good match (i.e., drawing a high match quality z) can take time. Our results are also qualitatively similar to those of Greenwood et al. (2016), who find that, in absence of any change in household technology, changes in wage structure can partly explain the decline of marriage and the increase in cross-sectional income inequality. However, without changes in the nature of the gains from marriage, they cannot explain why college graduation rates have increased faster for women and specialized couples have become less frequent.

6.2. From the 1970s to the 2010s

Table 5: Counterfactual experiment, decomposition of changes between the 1970s and the 2010s

	1970s	Counterfactuals (see notes)				2010s
	(1)	(2)	(3)	(4)	(5)	(6)
% college graduates (men)	27.37	48.79	29.31	32.35	31.10	36.62
% college graduates (women)	20.57	40.88	21.82	23.46	32.35	42.53
% matched (men)	81.92	79.54	80.31	81.51	80.93	68.12
% matched (women)	80.01	77.89	78.63	81.71	79.96	67.27
% matched (men, NC)	81.77	78.07	79.92	81.19	81.16	65.00
% matched (men, CG)	82.32	81.08	81.25	82.18	80.41	73.50
% matched (women, NC)	79.99	80.64	80.48	83.40	81.29	63.06
% matched (women, CG)	80.07	73.90	72.00	76.18	77.19	72.96
Separation rate (men)	3.77	3.62	3.84	3.27	2.34	6.60
Separation rate (women)	3.56	3.41	3.62	3.06	2.23	6.27
Match rate (men)	34.43	30.11	31.97	31.68	27.68	23.13
Match rate (women)	30.20	26.80	28.39	30.69	25.48	21.45

Notes: We report different moments of the matching distribution and educational achievements for different steady-state equilibria. Column (1) is the steady-state equilibrium implied by the parameters estimated with data from the 1971-1975 period. Column (2) corresponds to the counterfactual steady-state equilibrium implied by the parameters for the 1971-1975 period with wages from the 2011-2015 period (details in Table 5). Column (3) corresponds to the steady-state equilibrium implied by the parameters for the 1971-1975 period with wages and the distribution of the cost of college education from the 2011-2015 period. Column (4) corresponds to the steady-state equilibrium implied by the parameters for the 1971-1975 period with wages, college costs and mortality rates from the 2011-2015 period. Column (5) corresponds to the steady-state equilibrium implied by the parameters for the 1971-1975 period with wages, college costs, mortality rates and search frictions from the 2011-2015 period. Column (6) corresponds to the steady-state equilibrium implied by the parameters estimated with data from the 2011-2015 period. Residuals changes between columns (5) and (6) are due to changes in the shape of the Pareto frontier ($\log H_{i,j}$). NC means "no college"', "CG" means "college graduates".

In Table 5, we progressively shift all primitive factors from their 1970s to their 2010s levels. Column (2) describes the counterfactual equilibrium seen in the previous section. In column (3), the rising cost of college balances the increased labor market college premium and leads to a slight increase in the number of marriages by curbing graduation rates. The mechanism is exactly the opposite of what discussed in the previous section. In column (4), longer life expectancy results in higher college graduation rates and more marriages. The number of couples dissolved due to the death of a partner sharply decreases, which particularly benefits women, who often became widows before turning 60 in the 1970s. In summary, changes in the wage process, cost of college, and mortality rates can jointly explain about 53% of the rise of men's college graduation rates, but only 13% of the rise of women's rates.

In column (5), we show that the changes in marriage market frictions observed between the 1970s and the 2010s have led to lower match rates, higher separation rates, and a lower number of matches overall. In other words, frictions have become stronger, which might sound counter-intuitive given the rise of cohabitation and technological progress in matchmaking technology thanks to the Internet (e.g., online platforms and dating apps). Decreased match rates are partly due to a change in the age structure of frictions. In Figure 31, we have shown that search frictions have become weaker for older individuals, which relieves the pressure on the search of young individuals, who can now take their time before making a decision. A second explanation is that changes in frictions could reflect an increase in the sunk cost associated with starting a family. This interpretation, in line with the analysis of Lundberg et al. (2016) and Chiappori et al. (2017), suggests that the relative importance of static gains (e.g., sharing the rent and domestic chores) might have decreased, while the relative importance of match-specific investments (e.g., starting a relationship to have a child together) might have increased. In Table 1, we show that female college graduates face stronger frictions in more recent periods: an explanation is that they demand a higher commitment and let coincide household formation with an initial match-specific investment, such as a costly wedding, the birth of a child or the purchase of real estate.[29]

Finally, the last column of Table 5 shows that changes in household technology and norms, captured by the shape of the household Pareto frontier $\log H_{i,j}$, explain the remaining differences in outcomes between the two periods. The diffuse decrease in the value of marriage leads to lower match rates, higher separation rates, and a decrease in the number of overall matches. While changes in $\log H_{i,j}$ negatively affect the marriage market outcomes of all categories, their consequences are particularly strong for both men and women without a college degree. The interaction between the changes in the wage structure and in $\log H_{i,j}$ are partly responsible for both these marriage patterns and the increased marriage market returns to college. As we move from a matching based on gender-asymmetric household specialization to a matching motivated by love, companionship, mutual insurance, the relative importance of wages (and

[29] Lafortune and Low (2017) stress that sharing assets, including real estate, through a legal marriage contract imply a strong commitment. Laws regulating the division of property upon divorce can change the incentives to commit.

wage expectations) increases, leading agents, particularly women, to invest more on human capital.

7. Conclusion

In this paper, we present a new search-and-matching model of marriage and divorce over the life-cycle with endogenous education decisions. Our work is motivated by the increasing differences in marriage market outcomes between human capital groups. On one hand, highly educated individuals delay household formation to search for a good match, but eventually engage in long-lasting relationships. On the other hand, individuals with low human capital are exposed to a higher separation risk and more often engage in short-term relationships.

Our empirical investigation aims to understand the determinants of marriage, divorce, and educational choices in the U.S. over the last 50 years. In our model, we let marriage and divorce patterns be determined by both economies of scale in household production and consumption, as well as search frictions. We discuss the identification of both objects with panel data on marriages and divorces: our identification strategy can be seen as an extension of Choo and Siow (2006), Choo (2015), and Galichon and Salanié (2015). Our empirical findings show that assortative mating is partly explained by search frictions being weaker among likes. We also emphasize that economies of scale are stronger for specialized couples, although the joint presence of two high-wage partners is still beneficial because of the mutual insurance mechanism implied by the marriage contract. We also find that gender roles in the household are asymmetric; couples where the husband is the primary earner tend to be more productive. Since the 1970s, the college wage premium has increased, driving up labor market returns to college. At the same time, women's human capital investments have started paying off on marriage markets in the form of more stable relationships and high-wage partners. While wage inequality between genders has decreased and gender asymmetries in household production have been slowly fading away, fewer couples rely on specialization. On the other hand, more couples benefit from mutual insurance. Women are more often the primary earner and can sustain household consumption in case of negative income shocks. As a result, marriage market returns to college have increased.

While our model helps us distinguish between the different sources of assortative mating on marriage markets, our analysis does not directly look into the microeconomic household behavior behind our results. In particular, there are few if any works analyzing patterns of partner search behavior and providing direct empirical evidence of marriage market segmentation.[30]

[30] Big data produced by dating apps are certainly an interesting pathway to explore (Hitsch et al., 2010; Belot and Francesconi, 2013), although there are likely many steps between selecting a first date on an app and

On the other hand, our findings on search frictions can partly reflect systematic differences in match-specific investment across population groups. Economists have long been interested in human capital accumulation patterns resulting from household specialization, and Shephard (2018) is the first to combine this analysis with a dynamic marriage market equilibrium framework. However, there can be additional reasons explaining why some individuals - typically those with higher human capital - succeed in committing to long-term relationships, while others are not as deeply invested or do not even try to start a family. Match-specific capital is arguably the key factor to understand these patterns. Concretely, such investments could take the form of joint savings, home ownership (Lafortune and Low, 2017), childbearing, and child-rearing (Lundberg et al., 2016; Chiappori et al., 2017).

♭ *Department of Economics, KU Leuven.*
Address: KU Leuven, Faculty of Economics and Business, 69 Naamsestraat, 3000 Leuven, Belgium. Email: edoardo.ciscato@kuleuven.be.

starting a family together.

A. Technical appendix

A.1. Recovering the Pareto frontier

After estimating both the surplus function $S_{i,j}$ and the meeting function $M_{i,j}$, we aim to decompose the static component of the surplus, as represented by the household Pareto frontier level $H_{i,j}$, and its dynamic component. A Bellman equation for the surplus function $S_{i,j}(z)$ is obtained by combining the Bellman equations for husbands and wives: (2.13) and (2.14), respectively. We can link the match surplus with the household Pareto frontier level $H_{i,j}$ using the Pareto frontier equation (2.5). Under the specific restrictions we put on utility functions in our parameterization, namely $\psi^g(.) = \exp(.)$, we can actually derive a closed-form recursive formulation of $S_{i,j}$:[31]

$$\begin{aligned} S_{i,j} =& \log H_{i,j} + \beta \mathbb{E}_{i',j'|i,j} \bar{S}_{i',j'} + \\ & - 2 \log \left(\exp \left(V_i^m - \beta \mathbb{E}_{i'|i} V_{i',0}^m \right) + \exp \left(V_j^f - \beta \mathbb{E}_{j'|j} V_{j',0}^f \right) \right) . \end{aligned} \tag{A.1}$$

We first recover $V_{i,0}$ and $V_{0,j}$ for any i and j using the Bellman equations (2.15) and (2.16). At terminal age $\bar{a}$, the reservation values correspond to the per-period payoffs of single agents. Solving backwards starting from $\bar{a}$, one can recover the reservation values for all types given the estimated surplus function $S_{i,j}$ and meeting probabilities from $M_{i,j}$. Once the reservation values are solved for, it is straightforward to recover $H_{i,j}$ from the Bellman equation (A.1) for any $\{i, j\}$.

A.2. Plackett's copula and transition matrices

The $AR(1)$ wage process is modeled through a copula that links the wage rank of an individual across two consecutive periods. The joint CDF of an agent's current wage rank r and future wage rank r' is given by the Plackett's copula:

$$C^g(r, r'|h, a) = \frac{1 + \rho(h,a)(u+v) - \left[1 + \rho(h,a)(u+v)^2 - 4\rho(h,a)(\rho(h,a)+1)uv\right]^{1/2}}{2\rho(h,a)} \tag{A.2}$$

where the parameter $\rho(h, a)$ is such that the higher $\rho(h, a)$, the lower the mobility. In particular, Nelsen (2007, Chapter 5) shows that ρ is a monotonically increasing function of the Spearman's rank correlation coefficient. Dropping the arguments (h, a) for the sake of clarity, the two are related as follows,

$$SpearmanCorrRate = \frac{2\rho + \rho^2 - 2(1+\rho)\log(\rho+1)}{\rho^2}. \tag{A.3}$$

[31] Under assumption the Transferable Utility 2, S always exists and is unique. However, a closed-form recursive formulation of S can only be derived under specific assumptions on the curvature of the utility functions; e.g., if ψ^g is linear or exponential.

A.3. Meeting process

In section 4.4, we assume the meeting function takes the form $M_{i,j}(\tilde{n}) = \exp(\lambda_{i,j} + c_i(\tilde{n}) + d_j(\tilde{n}))$. Galichon and Salanié (2015) show that this function is the unique solution of the following optimal transport problem:

$$\min_{M} \left\{ -\sum_{i,j} M_{i,j}\lambda_{i,j} + \mathcal{E}(M) \right\} \text{ s.t. (2.7) holds for any } i \text{ and (2.8) holds for any } j \tag{A.4}$$

where $-\lambda_{i,j}$ can be interpreted as the cost of setting up a meeting between two singles i and j. The second term of the objective function, $\mathcal{E}(M)$, is the entropy associated with the optimal transport problem; we use $\mathcal{E}(M) = \sum_{i,j} M_{i,j} \log M_{i,j}$, as per Choo and Siow (2006). This characterization allows us to calculate $M_{i,j}$ quickly by finding two vectors c and d so that the marginal constraints (2.7) and (2.8) are respected. As suggested by Galichon and Salanié (2015), we use an Iterative Projection Fitting Procedure (IPFP) to calculate c and d.

A.4. Matching moment estimator

The matching moment estimator can be thought of as the global minimum of the following program:

$$\min_{(\lambda,\alpha)} \sum_{k}^{K} (\mu_k(\lambda,\alpha) - \hat{\mu}_k)^2 \tag{A.5}$$

where $\mu(\lambda,\alpha)$ is a vector of K moment predictions obtained with parameters (λ,α) and $\hat{\mu}$ is the vector of corresponding empirical moments. As detailed in Section 4.4, the number of moments is equal to the number of parameters to estimate, so that the First Order Conditions yield $\mu(\lambda,\alpha) = \hat{\mu}$ and the global minimum, if it exists, is zero. More precisely, the number of moments of the distribution MF is equal to the number of parameters in λ, while the number of moments of the distribution DF is equal to the number of parameters in α. The way the moments are constructed is given by equations (4.9) and (4.10) in the main text.

Under the specific parameterization used in this paper, we could derive the Jacobian matrix $J(\lambda,\alpha)$ analytically. The popular Limited-memory Broyden–Fletcher–Goldfarb–Shanno algorithm (LBFGS) always converges to the same global minimum - where the objective function is zero - as long as the initial values are not too far from it. The objective function is clearly not convex everywhere, and the LBFGS can get stuck in local minima; although this is impractical, these local minima can be easily ruled out if the objective function is still greater than zero. In practice, we run a derivative-free optimization algorithm to look for the initial values over a vast region and then run the LBFGS algorithm to obtain the global minimum. For a given sample, we have never found more than one global minimum so that the objective function is zero (i.e., so that all First Order Conditions hold).

After obtaining estimates $(\hat{\lambda}, \hat{\alpha})$, we compute the Jacobian matrix $J^* \equiv J(\hat{\lambda}, \hat{\alpha})$. The covariance matrix of $(\hat{\lambda}, \hat{\alpha})$ corresponds to:

$$(J^*)^T \Sigma J^* \tag{A.6}$$

where Σ is the covariance matrix of $\hat{\mu}$; Σ is computed with 5,000 bootstrap replications from the relevant PSID sample.

B. Additional Tables

Table 6: Average cost of college

Population group	1971-1975 (1)	1981-1985 (2)	1991-1995 (3)	2001-2015 (4)	2011-2015 (5)
Men (L)	1.91	-0.37	5.54	2.42	7.21
Men (H)	7.30	5.35	14.93	5.53	10.37
Women (L)	15.42	12.06	23.15	11.61	16.49
Women (H)	18.53	16.21	30.92	15.78	18.93
Ave. monetary fees (log)	7.93	6.73	6.79	7.05	7.34

Notes: The table reports point estimates of the average total cost of college, whose expression is given by (4.11). The notation (L) indicates agents with a less favorable family background (i.e., less educated parents and worse socioeconomic background), whereas (H) indicates agents with a more favorable background. In the last line, we report the average in-state undergraduate log-fees at public college institutions.

C. Additional Figures

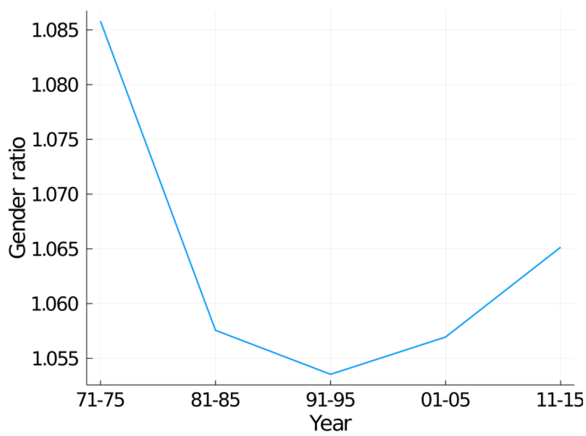

(a) Gender ratio in the working-age population

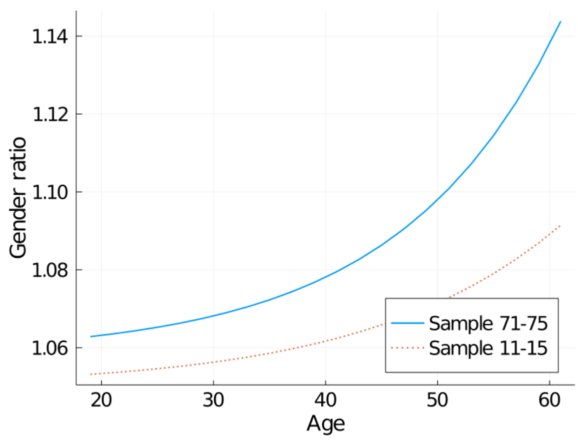

(b) Gender ratio by age and decade

Figure 23: Changes in gender ratio over time.

Notes: The gender ratios are calcualted as the overall number of women over the overall number of men. They are implied by the mortality rates estimated with Center for Disease Control and Prevention. The gender ratio at the age of 19 is calibrated to match the overall gender ratio in the PSID sample.

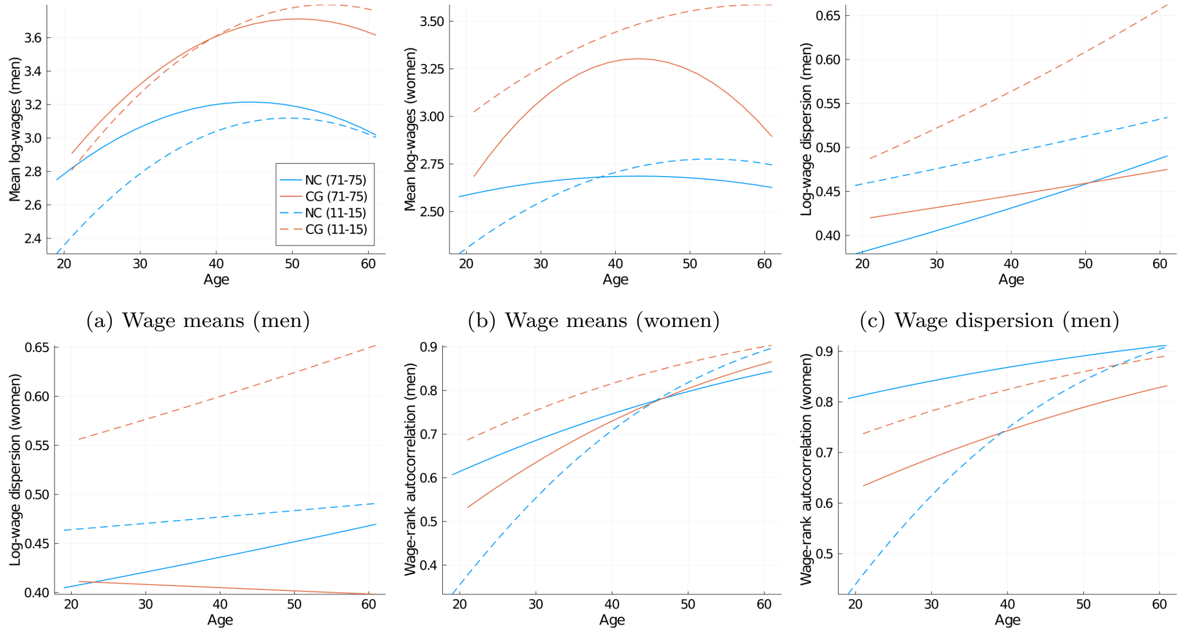

(a) Wage means (men) (b) Wage means (women) (c) Wage dispersion (men)

(d) Wage dispersion (women) (e) Wage mobility (men) (f) Wage mobility (women)

Figure 25: Changes in wage structure, 1970s vs 2010s.

Notes: Estimates obtained with PSID data as detailed in Section 4.3. Wage dispersion is measured by the cross-sectional wage variance and wage mobility is measured by the Spearman's wage (rank) autocorrelation.

Figure 27: Changes in the gender wage gap by wage percentile.

Notes: Men's and women's wage percentiles refer to the distribution of full-time log-wages. These are estimated after accounting for selection as explained in Section 4.3. The gender wage gap at different percentiles is measured as the woman's log-wage percentile over the man's.

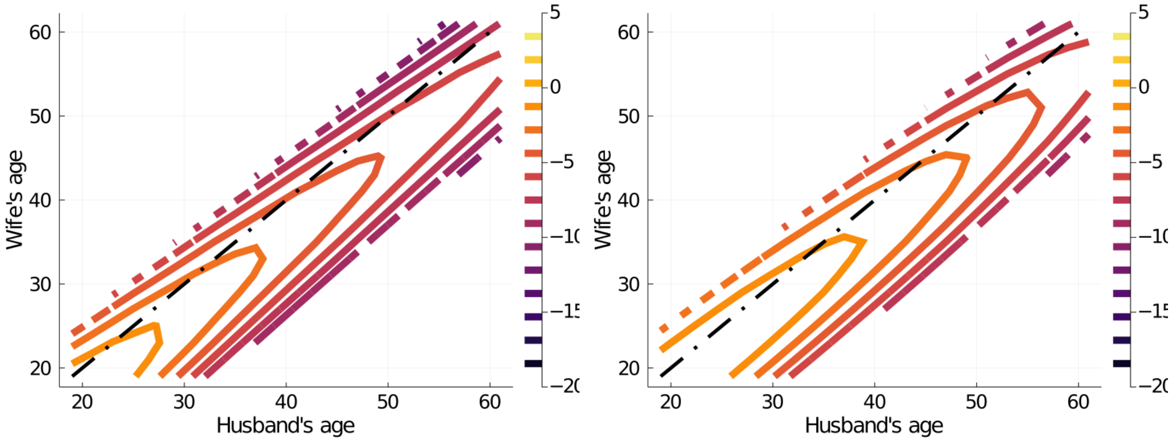

(a) 1971-1975

(b) 2011-2015

Figure 31: Estimated meeting rates $\hat{\lambda}$ as a function of partners' ages

Notes: The figure shows how the estimated $\hat{\lambda}_{i,j}$, the exogenous parameter of the meeting function $M_{i,j}$ that measures the strength of search frictions, changes with agents' ages. The higher $\hat{\lambda}_{i,j}$, the more likely a pair of single agents $\{i, j\}$ is to meet on the marriage market, all else constant. In white areas, meetings are ruled out by assumption. The average levels of $\hat{\lambda}_{i,j}$ are normalized to compare the two data periods.

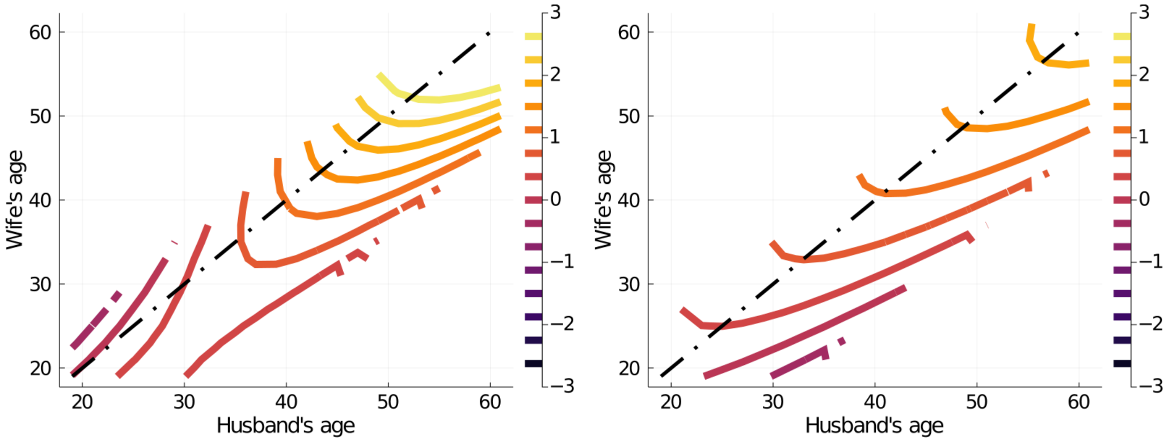

(a) 1971-1975

(b) 2011-2015

Figure 29: Estimated surplus $\hat{S}^*$ as a function of partners' ages

Notes: The figure shows how the estimated $\hat{S}_{i,j}$, the deterministic component of the surplus function common to all pairs $\{i, j\}$, changes with agents' ages. The higher $\hat{S}_{i,j}$, the larger the average gains from marriage for a couple $\{i, j\}$ (and the more stable the relationship). The average levels of $\hat{S}_{i,j}$ are normalized to compare the two data periods.

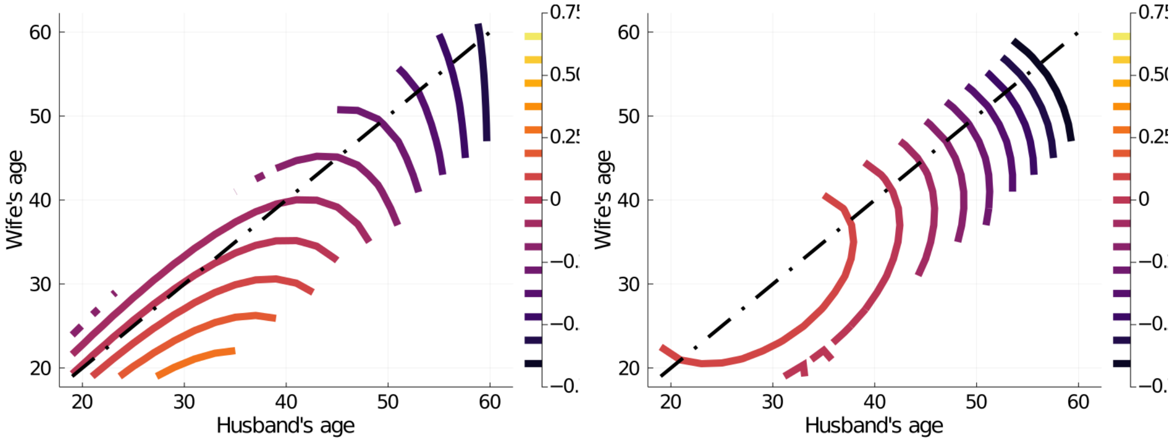

(a) 1971-1975

(b) 2011-2015

Figure 33: Estimated household production $\widehat{\log H_{i,j}}$ as a function of partners' ages

Notes: The figure shows how the estimated $\widehat{\log H_{i,j}}$, the deterministic component of the household Pareto frontier log-level for a couple $\{i, j\}$, changes with agents' ages. The higher $\widehat{\log H_{i,j}}$, the better the consumption opportunities for a couple $\{i, j\}$. The average levels of $\widehat{\log H_{i,j}}$ are normalized to compare the two data periods.

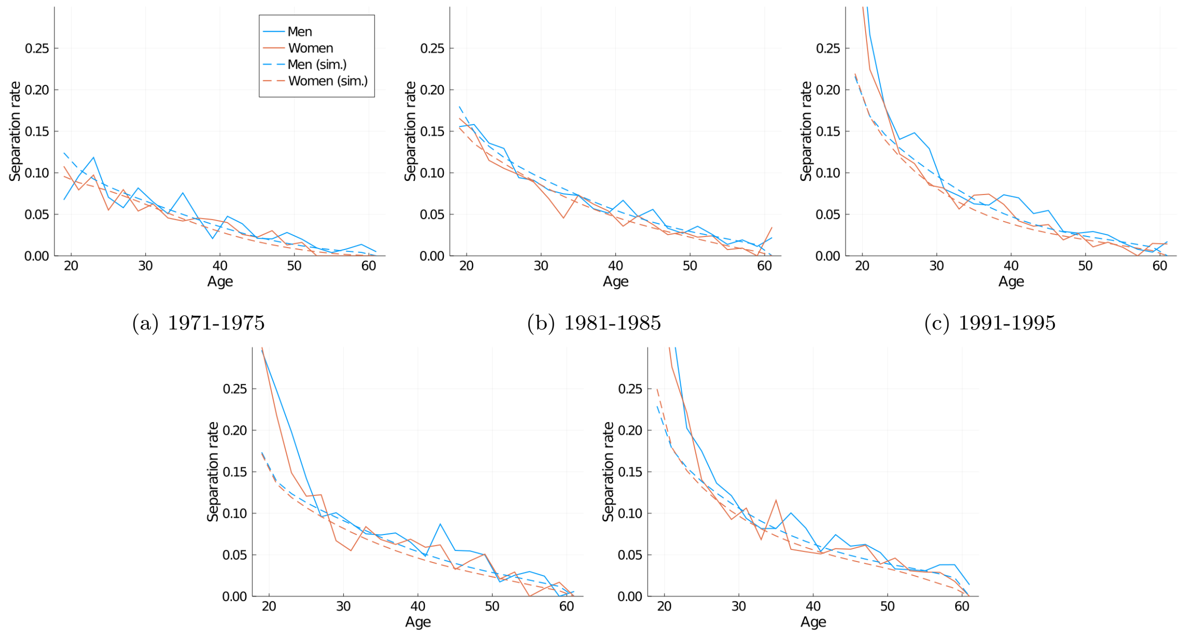

(a) 1971-1975

(b) 1981-1985

(c) 1991-1995

(d) 2001-2005

(e) 2011-2015

Figure 35: Separation rates by gender, age and period; model vs data

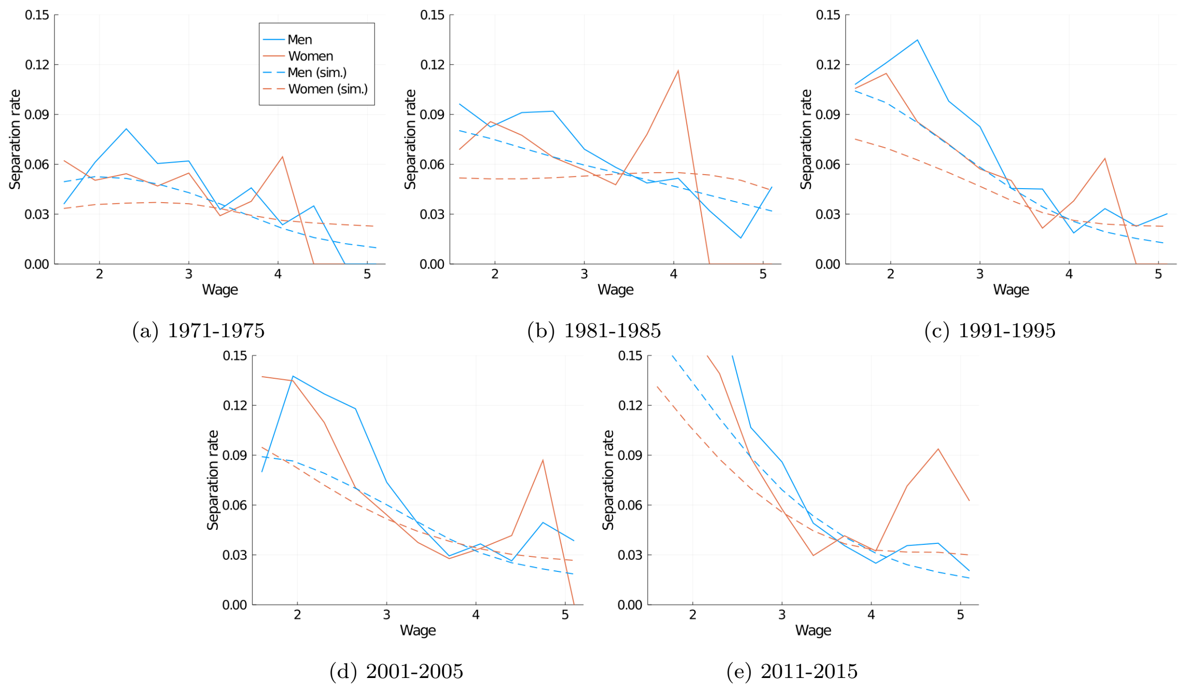

(a) 1971-1975

(b) 1981-1985

(c) 1991-1995

(d) 2001-2005

(e) 2011-2015

Figure 36: Separation rates by gender, wage and period; model vs data

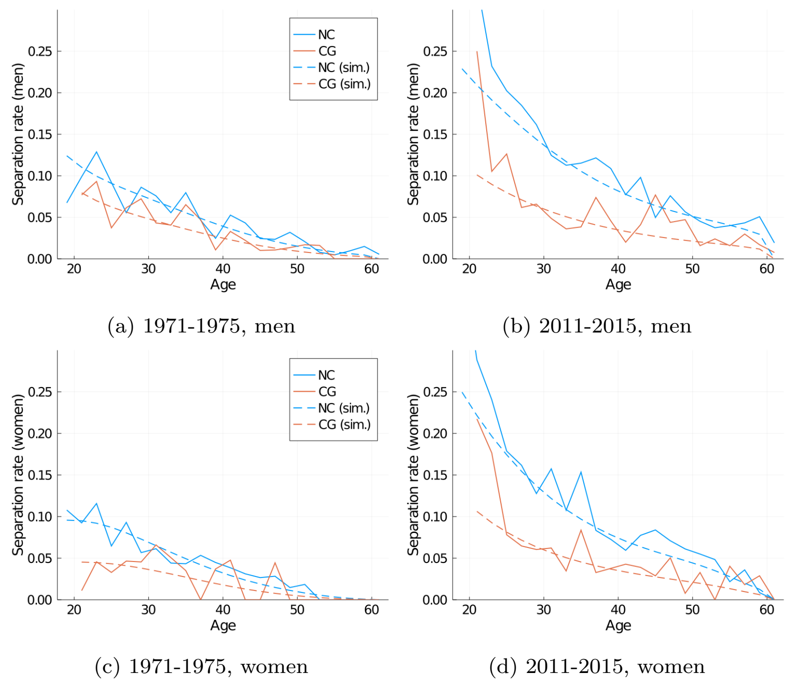

(a) 1971-1975, men

(b) 2011-2015, men

(c) 1971-1975, women

(d) 2011-2015, women

Figure 37: Separation rates by gender, age, education and period; model vs data

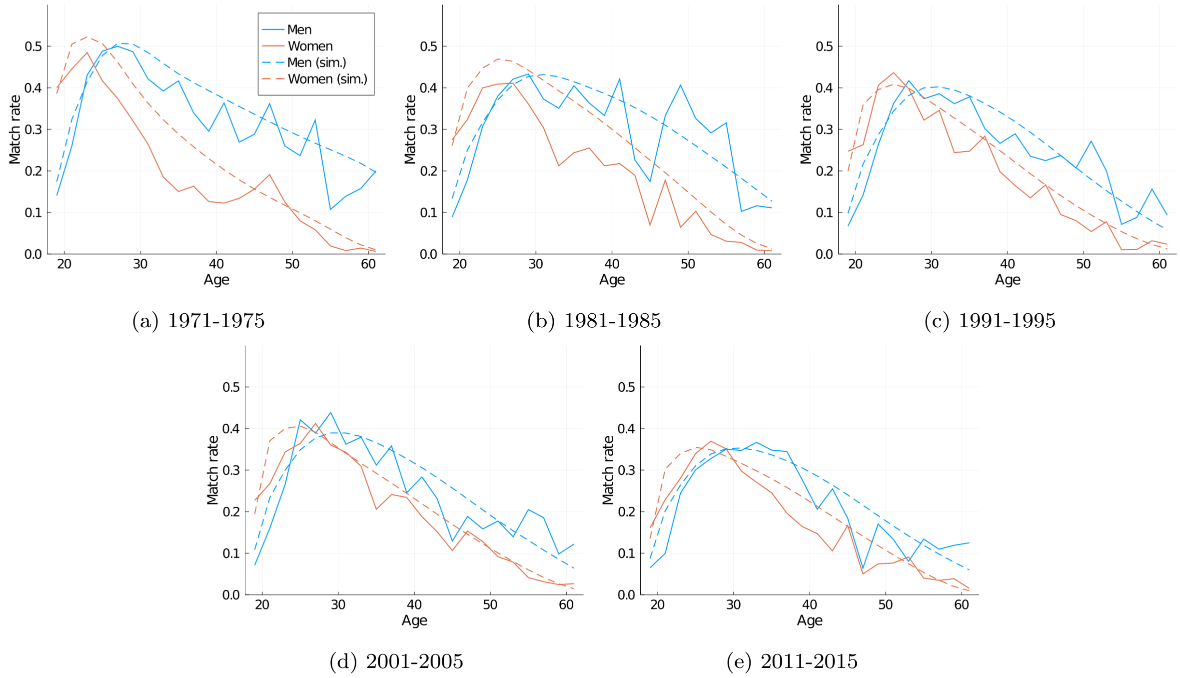

(a) 1971-1975 (b) 1981-1985 (c) 1991-1995

(d) 2001-2005 (e) 2011-2015

Figure 38: Match rates by gender, age and period; model vs data

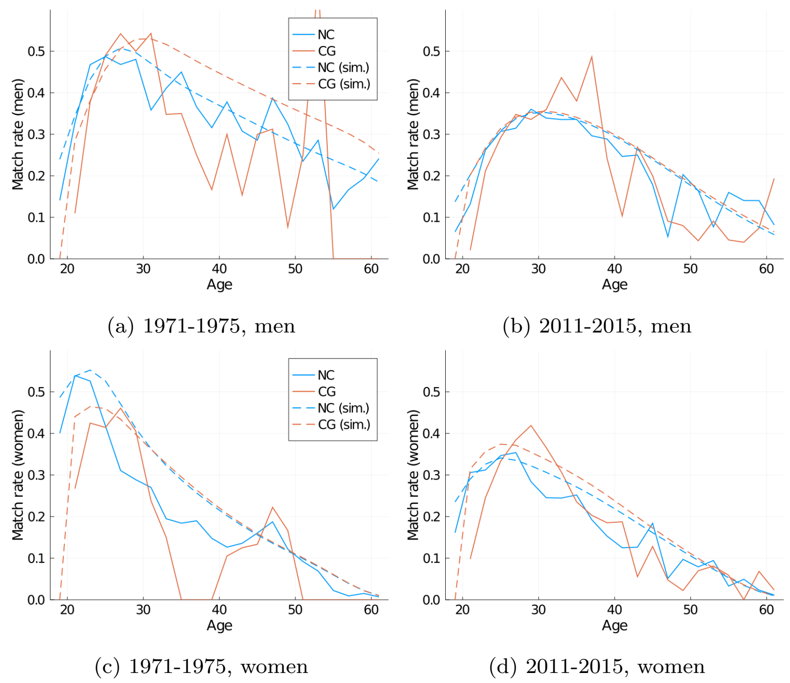

(a) 1971-1975, men (b) 2011-2015, men

(c) 1971-1975, women (d) 2011-2015, women

Figure 39: Match rates by gender, age, education and period; model vs data

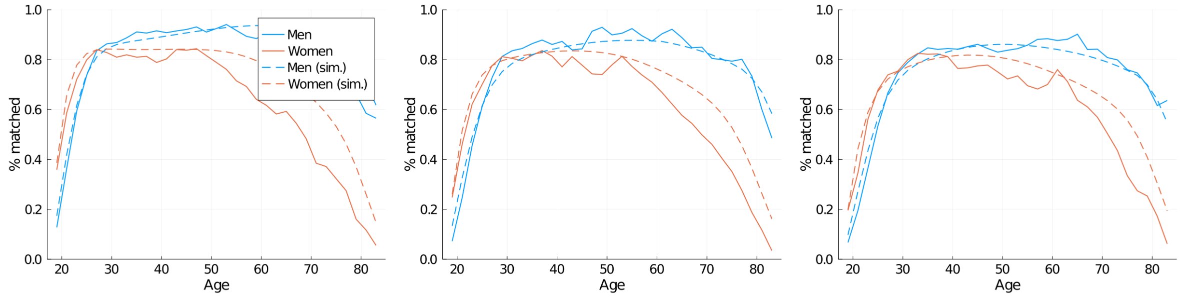

(a) 1971-1975 (b) 1981-1985 (c) 1991-1995

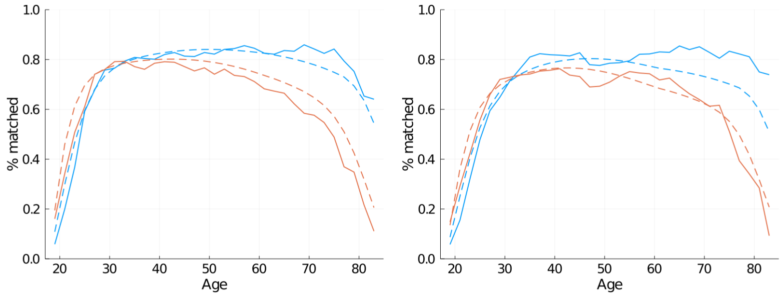

(d) 2001-2005 (e) 2011-2015

Figure 40: Share of matched individuals by gender, age and period; model vs data

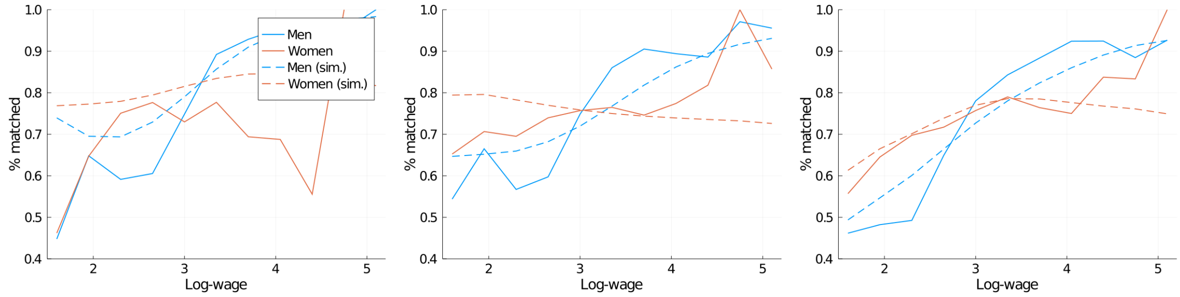

(a) 1971-1975 (b) 1981-1985 (c) 1991-1995

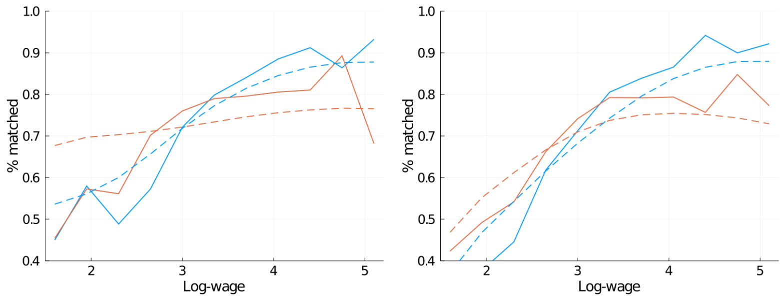

(d) 2001-2005 (e) 2011-2015

Figure 41: Share of matched individuals by gender, wage and period; model vs data

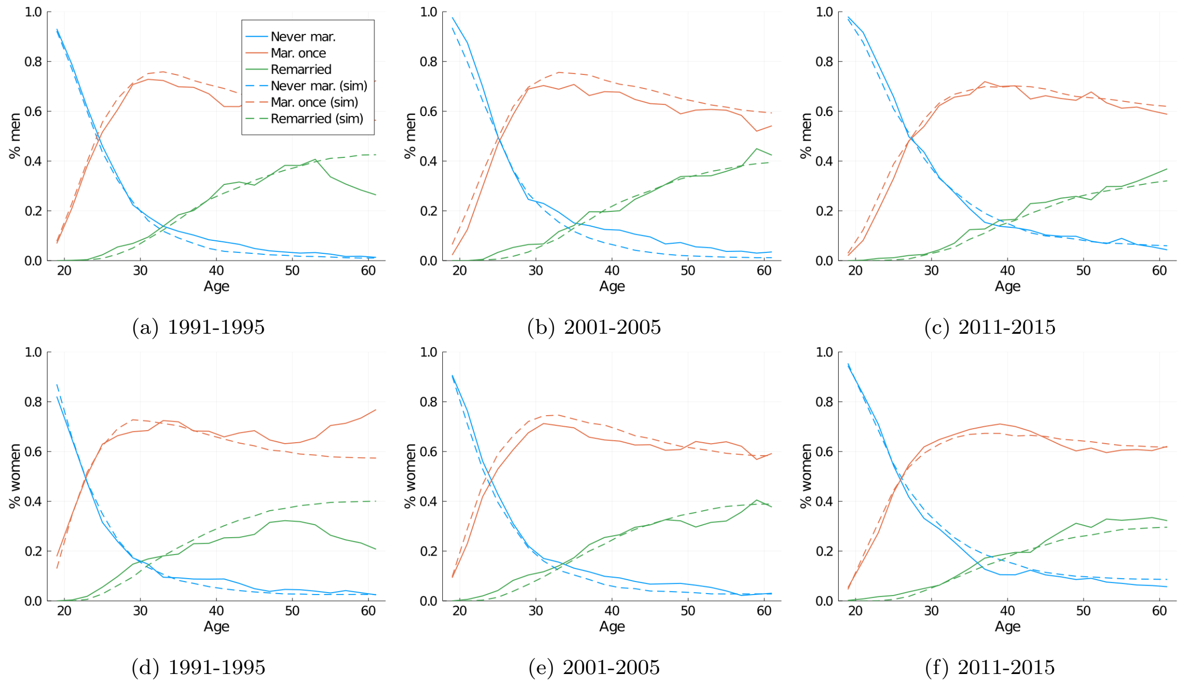

(a) 1991-1995 (b) 2001-2005 (c) 2011-2015

(d) 1991-1995 (e) 2001-2005 (f) 2011-2015

Figure 42: Share of married and remarried individuals by gender, age, and period; model vs data

Notes: The simulated moments in this figure are obtained with parameters estimated with an alternative sample, where we define matches as legal marriages and treat cohabiting individuals as singles. It is only possible to build these samples starting from the 1990s, while no difference was made between legal marriage and cohabitation earlier. The empirical moments in the figure are calculated using the retrospective marital history data collected from 1985; these data allow us to calculate the number of legal marriages of a respondent up to a certain date. Not all respondents filed the marital history questionnaire; marital histories are missing for about 25% of observations in our sample.

References

S Rao Aiyagari, Jeremy Greenwood, and Nezih Guner. On the state of the union. *Journal of Political Economy*, 108(2):213–244, 2000.

Andrew Beauchamp, Geoffrey Sanzenbacher, Shannon Seitz, and Meghan M Skira. Single moms and deadbeat dads: The role of earnings, marriage market conditions, and preference heterogeneity. *International Economic Review*, 59(1):191–232, 2018.

Gary S. Becker. A theory of marriage: Part i. *Journal of Political Economy*, 81(4):813–846, 1973.

Gary S Becker, Elisabeth M Landes, and Robert T Michael. *Journal of political Economy*, 85 (6):1141–1187, 1977.

Michèle Belot and Marco Francesconi. Dating preferences and meeting opportunities in mate choice decisions. *Journal of Human Resources*, 48(2):474–508, 2013.

Marianne Bertrand, Emir Kamenica, and Jessica Pan. Gender identity and relative income within households. *The Quarterly Journal of Economics*, 130(2):571–614, 2015.

Richard Blundell, Monica Costa Dias, Costas Meghir, and Jonathan Shaw. Female labor supply, human capital, and welfare reform. *Econometrica*, 84(5):1705–1753, 2016a.

Richard Blundell, Luigi Pistaferri, and Itay Saporta-Eksten. Consumption inequality and family labor supply. *American Economic Review*, 106(2):387–435, 2016b.

Stéphane Bonhomme and Jean-Marc Robin. Assessing the equalizing force of mobility using short panels: France, 1990–2000. *The Review of Economic Studies*, 76(1):63–92, 2009.

Mary Ann Bronson. Degrees are forever: Marriage, educational investment, and lifecycle labor decisions of men and women. 2014.

Charles Brown, Greg J Duncan, and Frank P Stafford. Data watch: The panel study of income dynamics. *Journal of Economic Perspectives*, 10(2):155–168, 1996.

Gustaf Bruze, Michael Svarer, and Yoram Weiss. The dynamics of marriage and divorce. *Journal of Labor Economics*, 33(1):123–170, 2015.

Ken Burdett and Melvyn G Coles. Marriage and class. *The Quarterly Journal of Economics*, 112(1):141–168, 1997.

Hector Chade and Gustavo Ventura. Taxes and marriage: a two-sided search analysis. *International Economic Review*, 43(3):955–985, 2002.

Pierre-André Chiappori. Collective labor supply and welfare. *Journal of Political Economy*, 100(3):437–67, 1992.

Pierre-André Chiappori, Bernard Salanié, and Yoram Weiss. Partner choice, investment in children, and the marital college premium. *American Economic Review*, 107(8):2109–67, 2017.

Pierre-André Chiappori, Monica Costa Dias, and Costas Meghir. The marriage market, labor supply, and education choice. *Journal of Political Economy*, 126(S1):S26–S72, 2018a.

Pierre-André Chiappori, Natalia Radchenko, and Bernard Salanié. Divorce and the duality of marital payoff. *Review of Economics of the Household*, 16(3):833–858, 2018b.

Eugene Choo. Dynamic marriage matching: An empirical framework. *Econometrica*, 83(4): 1373–1423, 2015.

Eugene Choo and Aloysius Siow. Who marries whom and why. *Journal of Political Economy*, 114(1):175–201, 2006.

Edoardo Ciscato. Marriage, divorce and wage uncertainty along the life-cycle. Technical Report 2018-046, Human Capital and Economic Opportunity Working Group, 2018. https://hceconomics.uchicago.edu/research/working-paper/marriage-divorce-and-wage-uncertainty-along-life-cycle.

Edoardo Ciscato and Simon Weber. The role of evolving marital preferences in growing income inequality. *Journal of Population Economics*, 33(1):307–347, 2020.

Melvyn G Coles and Marco Francesconi. On the emergence of toyboys: The timing of marriage with aging and uncertain careers. *International Economic Review*, 52(3):825–853, 2011.

Thomas Demuynck and Tom Potoms. Weakening transferable utility: the case of non-intersecting pareto curves. *Journal of Economic Theory*, 188:105035, 2020.

Javier Díaz-Giménez and Eugenio Giolito. Accounting for the timing of first marriage. *International Economic Review*, 54(1):135–158, 2013.

Lasse Eika, Magne Mogstad, and Basit Zafar. Educational assortative mating and household income inequality. *Journal of Political Economy*, 127(6):2795–2835, 2019.

Raquel Fernández and Richard Rogerson. Sorting and long-run inequality. *Quarterly Journal of Economics*, 116(4):1305–1341, 2001.

Raquel Fernández and Joyce Cheng Wong. Free to leave? a welfare analysis of divorce regimes. *American Economic Journal: Macroeconomics*, 9(3):72–115, 2017.

Raquel Fernández, Nezih Guner, and John Knowles. Love and money: A theoretical and empirical analysis of household sorting and inequality. *Quarterly Journal of Economics*, 120 (1):273–344, 2005.

Alfred Galichon and Bernard Salanié. Cupid's invisible hand: Social surplus and identification in matching models. Working paper, 2015.

Pilar Gonalons-Pons and Christine R Schwartz. Trends in economic homogamy: changes in assortative mating or the division of labor in marriage? *Demography*, 54(3):985–1005, 2017.

Marion Goussé, Nicolas Jacquemet, and Jean-Marc Robin. Marriage, labor supply, and home production. *Econometrica*, 85(6):1873–1919, 2017.

Jeremy Greenwood, Nezih Guner, and John A Knowles. More on marriage, fertility, and the distribution of income. *International Economic Review*, 44(3):827–862, 2003.

Jeremy Greenwood, Ananth Seshadri, and Mehmet Yorukoglu. Engines of liberation. *The Review of Economic Studies*, 72(1):109–133, 2005.

Jeremy Greenwood, Nezih Guner, Georgi Kocharkov, and Cezar Santos. Technology and the changing family: A unified model of marriage, divorce, educational attainment, and married female labor-force participation. *American Economic Journal: Macroeconomics*, 8(1):1–41, 2016.

Fatih Guvenen and Michelle Rendall. Women's emancipation through education: A macroeconomic analysis. *Review of Economic Dynamics*, 18(4):931–956, 2015.

Jonathan Heathcote, Kjetil Storesletten, and Giovanni L Violante. The macroeconomic implications of rising wage inequality in the united states. *Journal of Political Economy*, 118(4): 681–722, 2010.

James Heckman. Shadow prices, market wages, and labor supply. *Econometrica*, 42(4):679–94, 1974.

Gunter J Hitsch, Ali Hortaçsu, and Dan Ariely. Matching and sorting in online dating. *American Economic Review*, 100(1):130–63, 2010.

Sonia Jaffe and Simon Weber. The effect of meeting rates on matching outcomes. *Economic Theory*, 67(2):363–378, 2019.

Michael P Keane and Kenneth I Wolpin. The role of labor and marriage markets, preference heterogeneity, and the welfare system in the life cycle decisions of black, hispanic, and white women*. *International Economic Review*, 51(3):851–892, 2010.

Wojciech Kopczuk, Emmanuel Saez, and Jae Song. Earnings inequality and mobility in the united states: evidence from social security data since 1937. *The Quarterly Journal of Economics*, 125(1):91–128, 2010.

Jeanne Lafortune and Corinne Low. Tying the double-knot: The role of assets in marriage commitment. *American Economic Review*, 107(5):163–67, 2017.

Ethan Ligon, Jonathan P Thomas, and Tim Worrall. Informal insurance arrangements with limited commitment: Theory and evidence from village economies. *The Review of Economic Studies*, 69(1):209–244, 2002.

Corinne Low. Pricing the biological clock: Reproductive capital on the us marriage market. Working paper, 2014.

Shelly Lundberg, Robert A Pollak, and Jenna Stearns. Family inequality: Diverging patterns in marriage, cohabitation, and childbearing. *Journal of Economic Perspectives*, 30(2):79–102, 2016.

Mihai Manea. Steady states in matching and bargaining. *Journal of Economic Theory*, 167: 206–228, 2017.

Maurizio Mazzocco. Saving, risk sharing, and preferences for risk. *American Economic Review*, 94(4):1169–1182, 2004.

Roger B Nelsen. *An introduction to copulas*. Springer Science & Business Media, 2007.

Ferdinando Regalia, Jose-Victor Rios-Rull, and Jacob Short. What accounts for the increase in the number of single households? Working paper, 2001.

Ana Reynoso. The impact of divorce laws on the equilibrium in the marriage market. Working paper, 2017.

Michael J Rosenfeld and Reuben J Thomas. Searching for a mate: The rise of the internet as a social intermediary. *American Sociological Review*, 77(4):523–547, 2012.

Christine R Schwartz. Trends and variation in assortative mating: Causes and consequences. *Annual Review of Sociology*, 39:451–470, 2013.

Lloyd S Shapley and Martin Shubik. The assignment game i: The core. *International Journal of Game Theory*, 1(1):111–130, 1971.

Andrew Shephard. Marriage market dynamics, gender, and the age gap. Working paper, 2018.

Michelle Sheran. The career and family choices of women: A dynamic analysis of labor force participation, schooling, marriage, and fertility decisions. *Review of Economic Dynamics*, 10 (3):367–399, 2007.

Robert Shimer and Lones Smith. Assortative matching and search. *Econometrica*, 68(2):343–369, 2000.

Betsey Stevenson and Justin Wolfers. Marriage and divorce: Changes and their driving forces. *The Journal of Economic Perspectives*, 21(2):27–52, 2007.

Wilbert Van der Klaauw. Female labour supply and marital status decisions: A life-cycle model. *The Review of Economic Studies*, 63(2):199–235, 1996.

Alessandra Voena. Yours, mine, and ours: Do divorce laws affect the intertemporal behavior of married couples? *American Economic Review*, 105(8):2295–2332, 2015.

Linda Y Wong. Structural estimation of marriage models. *Journal of Labor Economics*, 21(3): 699–727, 2003.