ANNUAL REVIEWS

Annual Review of Economics

The Theory and Empirics of the Marriage Market

Pierre-André Chiappori[1,2]

[1]Department of Economics, Columbia University, New York, NY 10027, USA; email: pc2167@columbia.edu

[2]Toulouse School of Economics, 31000 Toulouse, France

Annu. Rev. Econ. 2020. 12:547–78

First published as a Review in Advance on May 1, 2020

The *Annual Review of Economics* is online at economics.annualreviews.org

https://doi.org/10.1146/annurev-economics-012320-121610

JEL codes: D1, D7, D4

Keywords

marriage, frictionless matching, search, transferable utility, imperfectly transferable utility

Abstract

This article reviews recent developments in the literature on marriage markets. A particular emphasis is put on frameworks based either on frictionless matching models with transfers or on search models.

1. INTRODUCTION

In economics, the mere notion of a market for marriage is (relatively) a newcomer. Becker (1973, 1974) was the first to point out that the tools of economic analysis (and in particular price theory) could be applied to the analysis of such demographic phenomena as marriage, divorce, or fertility—which until then had been left to sociologists and anthropologists.[1] Since his seminal contributions, an abundant literature has been devoted to these issues. Moreover, the topic has recently attracted renewed attention, in particular with the emergence of non-unitary models of household behavior.

The importance of the family as a social group cannot be overestimated.[2] Its crucial role in human evolution has been abundantly documented. Even within contemporary developed societies, its role remains central in many respects. Demographic issues such as fertility, aging, or immigration are at the core of many major challenges modern economies face. Investment in human capital (HC) has been known for a long time to be a key component of economic growth. HC accumulation obviously starts within the family; moreover, an abundant literature on child development has demonstrated that the first few years of life have a long-lasting influence on the whole process. A large fraction of the increase in inequality over the recent decades is likely to have been generated by demographic phenomena—not only because assortative matching in marriage tends to mechanically amplify inequality among individuals, but also, and in a much deeper way, because parents endowed with a larger stock of HC appear to invest more in their children, generating even more inequality of opportunity for the next generation. Finally, a common feature of many non-unitary models—starting with the now dominant collective approach—is the emphasis put on power relationships within the household. The idea that intra-household allocation should be analyzed as an endogenous phenomenon largely reflecting conditions of the marriage market dates back to Becker (1973). The fact that, more than four decades later, modern analysis should heavily rely on these insights is a clear tribute to Becker's legacy.

The next section is devoted to the main theoretical approaches to marriage markets. These raise a series of econometric challenges, which are described in Section 3. Several empirical applications are reviewed in the last section.

2. MODELS OF THE MARRIAGE MARKET: THEORY

Most of this survey is devoted to a class of models that emphasize the very specific nature of the market for marriage. While differing in a number of important dimensions, these models share a common feature: They consider individuals who are fundamentally heterogeneous. Following the standard approach of the hedonic literature, this heterogeneity can be described by a list of characteristics (or traits). An important aspect is that some of these traits, while visible to the individuals themselves (and their potential mates), are unobserved by the econometrician; this formulation has important consequences in terms of empirical applications. In particular, individuals may differ in tastes; for instance, the (relative and absolute) valuation of the observable characteristics of an individual will not be uniform over the set of potential mates—on the contrary, it will typically be mate specific.

Before presenting these models in some detail, one should mention two different (although by no means incompatible) perspectives. Grossbard's (1993) approach is based on a

[1] Pollak (2003) documents how many researchers (including well-known ones) were openly hostile to Becker's application of mathematical microeconomic tools to issues related to the family.

[2] Doepke & Tertilt (2016) provide a recent survey of the links between family economics and macroeconomic analysis.

standard partial equilibrium vision in which homogeneous spouses are producers and sellers of activities that benefit their partners, which she calls "work-in-household" (WiHo). In this model, individuals independently maximize their private utility under a budget constraint that includes prices for each individual's WiHos. Gersbach & Haller (2017) consider marriage markets from a general equilibrium perspective. In particular, they analyze the interactions between the efficient, but in general heterogeneous, collective decision processes used by households, on the one hand, and the constraints imposed by the competitive market environment in which these households operate on the other hand. A particular focus of their work is on the efficiency properties and decentralization possibilities of the allocation mechanism.

2.1. Models of the Marriage Market: A Taxonomy

The numerous models that have been used to analyze marital patterns can be classified into a few general categories.

2.1.1. Search versus frictionless matching. From a theoretical perspective, the analysis of marriage markets relies on a simple but fundamental intuition: Marriage generates a surplus, in the sense that, when married, two individuals can both achieve a higher level of well-being than they would as singles. The exact nature of this surplus is complex and may, depending on the issues under consideration, be described in different ways. Nonmonetary aspects, a vague notion that includes (among other things) what is usually called love, certainly do play a role, although economists are often quite reluctant to model them in any specific manner; most of the time, these features are summarized by some random variable that represents, in a parsimonious way, the "quality" of the match. Economic benefits, on the other hand, are generated in ways that are more familiar to economists—from the existence of commodities that are publicly consumed within the family (children's welfare being a crucial example) to gender specialization to risk sharing.[3]

A precise definition of the market for marriage thus requires two components: (*a*) a description of the two populations at stake (technically, this will take the form of a measure over the space of individual characteristics), and (*b*) an evaluation of the benefits that would be generated by the matching of any two potential spouses.

In turn, the economic analysis of the marriage will aim at answering two sets of questions: (*a*) Who marries whom? and (*b*) How are the benefits distributed between spouses? Mostly, these questions have been analyzed using either of two different frameworks: frictionless matching theory and search models. The basic distinction between the two is related to the emphasis that is put (or not) on frictions in the description of the market. In search models, frictions are paramount. Typically, each individual sequentially and randomly meets one person of the opposite gender; after such a meeting, both individuals must decide whether to settle for the current mate or continue searching. The latter option involves various costs, from discounting to the risk of never finding a better partner. If both individuals agree to engage in a relationship (which can be marriage or, in some models, cohabitation), then a negotiation begins on the way the surplus is shared.

[3] In this review, I concentrate on bilateral, one-to-one matching; that is, I consider the standard situation of marriage between one man and one woman. Needless to say, similar tools could be applied to analyze different contexts, such as same-sex marriages (the so-called roommate problem in matching theory) or polygamy. Interestingly, the theoretical properties of such problems can significantly differ from the standard ones. For instance, in both cases a stable matching may fail to exist, although existence typically obtains in large markets. The reader is referred to Azevedo & Hatfield (2015) for many-to-one matching, Chiappori et al. (2019c) for the roommate case, and Reynoso (2019) for a model of polygamy based on a matching approach.

Matching models, on the contrary, assume a frictionless environment. In the matching process, each woman (say) is assumed to have free access to the pool of all potential men, with perfect knowledge of the characteristics of each of them—and vice versa. In other words, matching models disregard the cost of acquiring information about potential matches as well as the role of meeting technologies of all sorts (from social media to head hunters and from dating sites to pure luck). As always, an assumption of this kind is but a simplification of the infinite complexity of real-life processes. The question is whether such a simplification is acceptable. While the answer can only be case specific, two general remarks can be made. First, the relevance of a frictionless setting largely depends on the question under consideration. A labor economist who is mostly interested in the dynamics of unemployment may be rightly reluctant to use a frictionless framework for modeling the matching of workers to jobs: There is a general consensus that a large fraction of unemployment results from various frictions on the labor market, and a search model should probably be preferred. If, however, the main issue is the allocation of specific individuals to specific tasks—for example, which type of CEO ends up heading which type of firms—then a frictionless context may be more acceptable: While it is probably true that firms may not have a perfect knowledge of the pool of available CEOs, the resulting bias may be of second order, and neglecting it may be fully appropriate.

Second, the size of the market, but also its structure, may have an impact on the relevance of a particular framework. Matching within a small community, in which individuals know each other, may be closer to the frictionless reference than matching within a very large market. Even with large populations (e.g., the market for marriage in a region of the United States), one may want to concentrate on a very specific aspect involving a small number of broad categories (e.g., matching by education, the latter being broadly defined by four or five possible levels); then one may well accept, for the sake of parsimony, to disregard other individual traits and consider that two spouses with the same education are perfect substitutes. In such a context, neglecting frictions may be admissible.[4]

In the end, the choice of a specific model (frictionless matching versus search or other) should be driven, at least in part, by empirical considerations: what are the main stylized features of the situation we want to investigate, what types of frictions are more likely to matter, how important they are likely to be, and, last but not least, which version fits the data best. It is therefore crucial to deeply understand the meaning and the implications of the various models under consideration, including in their auxiliary hypotheses and apparently mundane details, and to keep in mind the issues related to the empirical implementation of the various concepts at stake.

2.1.2. Frictionless matching models: formal structure. Within the family of frictionless matching frameworks, a second and crucial distinction relies on the role of transfers. The issue here is whether a technology exists that would allow to transfer utility between agents participating in a matching process. Whether such transfers are possible makes a fundamental difference: When available, they allow agents to "bid" for their preferred mate by accepting to reduce their own gain from the match in order to increase the partner's. The exact nature of these bids depends on the context and does not need to take the form of monetary transfers; in family economics, they may affect the allocation of time between paid work, domestic work, and leisure, or the choice between current and future consumption, or the structure of expenditures for private or public goods.

[4]A small but important literature considers the convergence of search equilibria to stable matchings of the frictionless framework when search frictions become negligible. The issue is quite complex, as one can find examples in which the limit of search equilibria fails to be efficient (see, for instance, Lauermann & Noldeke 2015 for the non-transferable utility case and Atakan 2006 for the transferable utility context).

Whatever the particular translation, the possibility of utility transfers enables agents to negotiate, compromise, and ultimately exploit mutually beneficial solutions.

An abundant literature considers the so-called non-transferable utility (NTU) case, where transfers are not possible: There is simply no technology enabling agents to decrease their utility to the benefit of a potential partner.[5] This framework has been applied to a host of important issues, from the allocation of residents to hospitals (Roth 1984) to kidney exchange (Roth et al. 2004, 2005) or the allocation of students to public schools (Abdulkadiroğlu & Sönmez 2003). In the case of marriage, however, the NTU case appears to be much less relevant; it is hard to imagine situations where any move along the Pareto frontier, whereby a spouse's welfare is decreased to the benefit of the partner, is simply ruled out. In this review, therefore, I concentrate on matching models involving transfers.[6]

When transfers are possible, the total benefit created by any possible matching has to be divided between partners. An equilibrium must therefore specify not only matching patterns—who is matched with whom—but also the supporting division of the surplus; the latter is now endogenous and determined (or at least constrained) by equilibrium conditions on the marriage market.[7]

When transfers are possible, a polar setting [based on transferable utility (TU)] postulates moreover that the transfer technology has a very striking property: It allows to transfer utility between agents at a constant exchange rate. This means that for a well-chosen cardinalization of individual utilities, increasing one's partner's utility by one util (i.e., unit of utility) has a cost of exactly one util for the individual, irrespective of the economic environment (prices, incomes, etc.). In that case, a given match generates some total gain, and moreover the division is such that individual utilities always add up to the total gain. Technically, the Pareto frontier, which represents the set of utility pairs that are just feasible given resource constraints, is a straight line with slope -1 irrespective of the constraints. Alternatively, a more general version [often called imperfectly transferable utility (ITU)] allows for transfers but recognizes that the exchange rate between individual utilities is not constant and is typically endogenous to the economic environment.

2.2. Matching Models Under Transferable Utility

Let me start with some notation. I consider two compact sets $\bar{X} \subset \mathbb{R}^n$ and $\bar{Y} \subset \mathbb{R}^m$, which respectively represent the space of female and male characteristics. Note that the spaces may be multidimensional—that is, $n \geq 1$ and $m \geq 1$ in general. The corresponding vectors of characteristics fully describe the agents; i.e., for any $\mathbf{x} \in \bar{X}$, two women with the same vector of characteristics $\mathbf{x}$ are perfect substitutes as far as matching is concerned (and similarly for men). These spaces are

---

[5]The interested reader is referred to Roth & Sotomayor's (1992) excellent monograph.

[6]NTU models have been applied to dating (see Hitsch et al. 2010, Banerjee et al. 2013). Regarding marriage, they may be relevant in contexts where transfers, although they do take place, are not endogenous to the matching mechanism but are instead determined outside of it (for example, in societies ruled by very rigid social norms). Another issue is that endogenous transfers typically require some minimum level of commitment between agents. In the (somewhat extreme) bargaining in marriage (BIM) framework by Lundberg & Pollak (2009), no commitment is possible at all. In a BIM world, any promise one may make before marriage can (and therefore will) be reneged upon just after the ceremony; there is just no way spouses can commit beforehand on their future behavior. Moreover, upfront payments, whereby an individual transfers some money, commodities, or property rights to the potential spouse conditional on marriage, are also excluded. Then the intra-household allocation of welfare is decided after marriage, irrespective of the commitment made before. Marriage decisions will therefore take the outcome of this yet-to-come bargaining process as given, and we are back to a NTU setting in which each partner's share of the surplus is fixed and cannot be altered by transfers decided ex ante.

[7]This is exactly Becker's (1973, p. 813) original intuition: "Theory does not take the division of output between mates as given, but rather derives it from the nature of the marriage market equilibrium."

endowed with measures F and G respectively; both $F(X)$ and $G(Y)$ are finite. In order to capture the case of persons remaining single within this framework, a standard trick is to augment the spaces by including an isolated point in each: a dummy partner $\emptyset_X$ for any unmatched man and a dummy partner $\emptyset_Y$ for any unmatched woman. Therefore, from now on I consider the spaces $X := \bar{X} \cup \{\emptyset_X\}$ and $Y := \bar{Y} \cup \{\emptyset_Y\}$, where the point $\emptyset_X$ (resp. $\emptyset_Y$) is endowed with a mass measure equal to the total measure of $\bar{Y}$ ($\bar{X}$). In particular, it is possible (although not efficient in general) to consider a matching in which all women (or all men) remain single by posing that they are all matched with $\emptyset_Y$ (or $\emptyset_X$).[8]

To define a matching, one must first answer question (*a*) above: Who marries whom? This defines a measure h on $X \times Y$; intuitively, one can think of $h(x, y)$ as the probability that x is matched to y. Note that this definition allows for randomization: A woman x may be matched with positive probability to several different men.[9] Also, this formulation implies conditions on the measure h, namely that the marginals of h on X and Y are F and G, respectively; formally, we obtain

$$\int_{\mathbf{y}\in Y} dh(\mathbf{x},\mathbf{y}) = F(\mathbf{x}) \text{ and } \int_{\mathbf{x}\in X} dh(\mathbf{x},\mathbf{y}) = G(\mathbf{y}). \qquad 1.$$

Note that this constraint is linear in h, a point that will become very important later on.

2.2.1. The transferable utility assumption in family economics.

Next, let us consider the TU case. The assumption here is that, for a well-chosen cardinalization of individual utilities, a potential match between Mrs. $\mathbf{x}$ and Mr. $\mathbf{y}$ generates a joint surplus[10] $S(\mathbf{x}, \mathbf{y})$, and that individual utilities add up to that surplus. Obviously, such a property requires specific assumptions on individual preferences;[11] a simple example will be provided below. Then a matching is defined by a measure h on $X \times Y$ satisfying Equation 1 and two functions $u(\mathbf{x})$ and $v(\mathbf{y})$ such that

$$h(\mathbf{x},\mathbf{y}) > 0 \Rightarrow u(\mathbf{x}) + v(\mathbf{y}) = S(\mathbf{x},\mathbf{y}). \qquad 2.$$

Here, $u(\mathbf{x})$ [resp. $v(\mathbf{y})$] denotes the utility reached by $\mathbf{x}$ for that matching, and Equation 2 simply states that matched people share the resulting surplus. Singles' utility is normalized to zero, implying that $S(\mathbf{x}, \emptyset_Y) = S(\emptyset_X, \mathbf{y}) = 0$ for all $\mathbf{x}$, $\mathbf{y}$.

Lastly, the equilibrium condition is stability (Gale & Shapley 1962, Shapley & Shubik 1971). A matching is said to be stable if (*a*) no matched individual would prefer being single, and (*b*) no pair of individuals would both prefer being matched together (for a well-chosen distribution of the surplus) over their current situation. The second condition is often referred to as divorce at will: Whenever it is violated, the corresponding individuals will each divorce their current spouse (or abandon their current singlehood) to form a new union, implying that the initial matching was not stable.

[8]Throughout this article, a match is defined as the association of two specific individuals, whereas a matching is the collection of matches (or individual singlehood) over the entire population.

[9]Allowing for randomization is crucial: It is easy to generate examples in which the unique equilibrium matching requires randomization for an open subset of characteristics (see, for instance, Chiappori et al. 2010). In the absence of randomization—i.e., when each x has a unique match $y = \phi(x)$ and vice versa—the matching is said to be pure.

[10]I use the standard definition of surplus, according to which it is the sum of utilities that the spouses can reach when matched, minus the sum of their individual utilities if singles. In particular, the surplus generated by singlehood (i.e., a match with the dummy partner $\emptyset_X$ or $\emptyset_Y$) is normalized to zero.

[11]Specifically, individual preferences must be such that, for a well-chosen cardinalization, individual conditional indirect utilities are affine in private expenditures, with an identical coefficient for all agents (see Chiappori & Gugl 2020 for a precise statement).

One can readily see that stability requires the following inequalities:

$$u(\mathbf{x}) + v(\mathbf{y}) \geq S(\mathbf{x}, \mathbf{y}) \;\; \forall (\mathbf{x}, \mathbf{y}) \in X \times Y. \tag{3}$$

Indeed, if there exists a pair $(\mathbf{x}, \mathbf{y}) \in X \times Y$ such that $u(\mathbf{x}) + v(\mathbf{y}) < S(\mathbf{x}, \mathbf{y})$, then $\mathbf{x}$ and $\mathbf{y}$, who are not matched with positive probability (otherwise we would have an equality), could both benefit from being matched together—a violation of stability. An equivalent statement is the following: If a matching is stable, the corresponding functions u and v, from X to $\mathbb{R}$ and from Y to $\mathbb{R}$, respectively, are such that

$$u(\mathbf{x}) = \max_{\mathbf{z} \in Y} \{S(\mathbf{x}, \mathbf{z}) - v(\mathbf{z})\} \;\; \text{and} \tag{4}$$

$$v(\mathbf{y}) = \max_{\mathbf{z} \in X} \{S(\mathbf{z}, \mathbf{y}) - u(\mathbf{z})\}, \tag{5}$$

the maximum being reached in each case for potential spouses (possibly including the dummy one) to whom the individual is matched with positive probability. Note that Equation 4 has a natural interpretation in hedonic terms: $v(\mathbf{y})$ is the price (in utility terms) that Mrs. $\mathbf{x}$ would have to pay should she choose to marry Mr. $\mathbf{y}$, and then she would keep what is left of the surplus, namely $S(\mathbf{x}, \mathbf{y}) - v(\mathbf{y})$. Obviously, the same argument applies, mutatis mutandis, to Equation 5.

2.2.2. The marriage market under transferable utility: basic results. The TU assumption, while strong, generates important properties.

2.2.2.1. Surplus maximization and duality. A crucial property of matching models under TU is their intrinsic relationship with a class of linear maximization problems called "optimal transportation."[12] Consider the following question: Find a measure h on $X \times Y$, the marginals of which are F and G, respectively, that maximizes the integral

$$\mathcal{S} = \int_{X \times Y} S(\mathbf{x}, \mathbf{y}) dh(\mathbf{x}, \mathbf{y}). \tag{6}$$

From an economic perspective, this problem has a straightforward interpretation; just think of a benevolent dictator who can match people at will and is trying to maximize total welfare. In a TU framework, where individual utilities can all be measured in the same units, the natural measure of total welfare is the sum of all surpluses generated; that is exactly the meaning of the right-hand-side integral in Equation 6.

Now, this problem is linear in h; one can thus apply the results of duality theory (see Galichon 2016, Chiappori 2017 for a precise description). The dual problem is: Find two functions u and v, respectively defined on X and Y, that minimize the sum

$$\tilde{\mathcal{S}} = \int_X u(\mathbf{x}) dF(\mathbf{x}) + \int_Y v(\mathbf{y}) dG(\mathbf{y}) \tag{7}$$

under the constraints

$$u(\mathbf{x}) + v(\mathbf{y}) \geq S(\mathbf{x}, \mathbf{y}) \;\; \forall (\mathbf{x}, \mathbf{y}) \in X \times Y.$$

Note that these are exactly the stability constraints in Equation 3.

[12]This class of problems can also be called "Monge-Kantorovitch." It was initially introduced by Gaspar Monge, who studied a soil-transport problem; the use of linear programming techniques to solve it is due to Kantorovitch.

The main result is that, under mild conditions, if h [resp. (u, v)] is a solution to the primal (dual) problem, then (h, u, v) define a matching, and this matching is moreover stable; conversely, if (h, u, v) is a stable matching, then h maximizes the surplus S and (u, v) solve the dual problem in Equation 7. In other words, finding a stable matching boils down to the resolution of a linear maximization problem.

This equivalence has several interesting implications. The existence of a stable matching is equivalent to the existence of a solution to a linear maximization problem that obtains under mild continuity and compactness conditions. The corresponding measure h will in general be unique—in the sense that while examples with multiple stable matchings can be constructed, they are typically not robust to small perturbations. Finally, and from a more practical perspective, a matching model under TU can be approached from two different perspectives: finding a measure h and two functions u and v satisfying the set of equalities and inequalities described in the definition, or maximizing total surplus (and deriving u and v as either the dual variables of the primal problem or the solutions to the dual problem). In many situations, particularly when numerical simulations are involved, the second approach is more tractable than the first, since it boils down to linear programming.

2.2.2.2. *Supermodularity.* Lastly, in the one-dimensional case $m = n = 1$, an important notion is the supermodularity of the surplus, defined by the following property: For all x, x', y, y' such that $x \leq x'$ and $y \leq y'$, we have

$$S(x, y) + S(x', y') \geq S(x, y') + S(x', y). \qquad 8.$$

When S is twice continuously differentiable, this is equivalent to a standard Spence-Mirrlees condition,

$$\frac{\partial^2 S}{\partial x \partial y}(x, y) \geq 0 \ \ \forall x, y.$$

One can readily see that when S is supermodular, then the only stable matching must be assortative; that is, for any two matched couples (x, y) and (x', y') such that $x \leq x'$, we must have that $y \leq y'$. Then matching patterns follow, among married couples, a simple rule: x is matched to y if and only if the total mass of matched women above x equals the total mass of matched men above y. Formally, the matching is pure, and (assuming atomless distributions and an identical mass of men and women) we obtain

$$1 - F(x) = 1 - G(y) \Rightarrow y = \phi(x) = G^{-1} \circ F(x).$$

In particular, matching patterns do not depend on the surplus function, provided the latter is supermodular.

Lastly, if Equation 8 holds with the opposite inequality, then the surplus function is submodular, and the stable matching is now negative assortative (i.e., larger x matching with smaller y and vice versa).

2.2.3. A simple example. To illustrate the previous discussion, consider the following example, borrowed from Chiappori (2017). The couple consumes two commodities: a private good C_i and a public good Q, which is domestically produced from parental time according to a Cobb-Douglas function $Q = (t_1 t_2)^\alpha$—a typical example being children's welfare or HC. Individual preferences are also Cobb-Douglas,

$$u_i(C_i, Q) = C_i Q,$$

where $i = 1, 2$. Agents only differ by their HC H_i; the time not devoted to children, $1 - t_i$, is spent on market work for a wage $w_i = WH_i$. The budget constraint is therefore given by

$$C_1 + C_2 + w_1 t_1 + w_2 t_2 = w_1 + w_2.$$

One can show that the model satisfies the TU conditions: Any (interior) efficient allocation must maximize the sum of utilities. Efficiency thus requires

$$C_1 + C_2 = \frac{w_1 + w_2}{1 + 2\alpha} \text{ and } w_1 t_1 = w_2 t_2 = \frac{\alpha(w_1 + w_2)}{1 + 2\alpha},$$

giving a surplus of

$$S(H_1, H_2) = \frac{\alpha^{2\alpha}}{(1 + 2\alpha)^{1+2\alpha}} W (H_1 + H_2)^{1+2\alpha} H_1^{-\alpha} H_2^{-\alpha}.$$

In particular, we obtain

$$\frac{\partial^2 S(H_1, H_2)}{\partial H_1 \partial H_2} = -\frac{\alpha^{2\alpha+1} W (H_1 + H_2)^{2\alpha-1}}{(2\alpha + 1)^{2\alpha+1} H_1^{\alpha+1} H_2^{\alpha+1}} (H_1^2 + H_2^2 + \alpha(H_1 - H_2)^2) < 0.$$

S is submodular, and the stable matching is negative assortative: High-HC men marry low-HC women (and vice versa). This reflects Becker's idea of (partial) specialization.[13] The optimal allocation of time requires low-wage people to devote much time to domestic production. Then a high-wage spouse can concentrate on market work, generating a large income and a high level of private consumption while benefiting from the large supply of public good provided by the partner's domestic work.

However, let us now change the production function by assuming that parental HC is also an input of the production function—which would typically be the case for children's development and education. If $Q = (H_1 t_1)^{\alpha/2} (H_2 t_2)^{\alpha/2}$, the model is still TU, but now the surplus becomes

$$S(H_1, H_2) = W \frac{\alpha^{2\alpha} T^{1+2\alpha}}{(1 + 2\alpha)^{1+2\alpha}} (H_1 + H_2)^{1+2\alpha},$$

which is supermodular (since $\partial^2 S / \partial H_1 \partial H_2 > 0$), generating positive assortative matching. This shows how matching patterns deeply depend on the precise nature of the technology of HC production.

2.2.4. The marriage market under transferable utility: extensions.

A few extensions of this basic framework are briefly discussed below.

2.2.4.1. Multidimensional matching.

While most applied matching models are (still) one-dimensional, the previous tools can be extended to multidimensional settings. An empirically important case is the so-called index model, in which the various dimensions enter the surplus only

[13] Full specialization obtains, for instance, when the time inputs in domestic production are perfect substitutes, as noted by Becker himself. But even complementary inputs may give negative assortative matching, as this example indicates.

through some one-dimensional index,

$$S(x_1, \ldots, x_n, \mathbf{y}) = \Sigma(I(x_1, \ldots, x_n), \mathbf{y}), \tag{9}$$

for some functions Σ and I. While a woman is represented by a vector of characteristics $\mathbf{x} = (x_1, \ldots, x_n)$, these only matter through the index I, which fully reflects her attractiveness on the marriage market: Two women with different vectors $\mathbf{x}$, $\mathbf{x}'$ but the same index ($I(\mathbf{x}) = I(\mathbf{x}')$) are perfect substitutes. The underlying intuition is the following. In a multidimensional setting, there exist trade-offs between the various traits that characterize a woman, which are described by the ratio (formally equivalent to a marginal rate of substitution)

$$M_{ij}(x, y) = \frac{\partial S/\partial x_j}{\partial S/\partial x_i}.$$

In words, M_{ij} represents the infinitesimal amount by which the j-th trait must be increased to compensate for an infinitesimal reduction in the i-th ("compensation" meaning that the surplus is not changed). In general, M_{ij} is y-specific: Two different males would disagree on the ratio. An index framework, on the contrary, postulates that the compensation is evaluated in exactly the same way by all men.

In practice, empirical applications often use a double index framework,

$$S(x_1, \ldots, x_n, \mathbf{y}) = \Sigma(I(x_1, \ldots, x_n), J(y_1, \ldots, y_m)). \tag{10}$$

Such applications impose strong restrictions on the model. But these restrictions are testable, in particular from regressions of female on male characteristics (and vice versa); moreover, if the corresponding restrictions are satisfied, then the corresponding indices are identified (up to a normalization) (see Chiappori et al. 2012, 2017b, 2020).

Another interesting situation obtains when dimensions m and n differ. Assume, for instance, that $m = 1$ while $n \geq 2$. In that case, a husband with a given characteristic y will marry with positive probability any of a continuum of different women $\mathbf{x}$, thus defining "iso-husband" curves in the space of female characteristics. Theory generates testable predictions relating the surplus function to the shape of iso-husband curves; moreover, these curves are (in principle) identifiable from data on matching pattern (see Chiappori et al. 2017b).

2.2.4.2. Premarital investments. So far, the characteristic spaces have been taken as given. Quite often, however, they result from some prior investment (think, for instance, of HC). A standard representation of the investment decision takes the form of a two-stage game: Individuals first (non-cooperatively) invest in their own HC, then they enter a matching game based on HC. Can we expect these investments to be socially efficient, given the noncooperative nature of stage one? Surprisingly, the answer is yes, as showed by Cole et al. (2001) and Nöldeke & Samuelson (2015). The latter contribution can be summarized as follows. Consider an auxiliary game in which the timing is reversed—i.e., people first match and then (cooperatively) invest in HC. Clearly, the outcome of the auxiliary game will be socially efficient. The result states that the stable matching of the auxiliary game can always be implemented as a Nash equilibrium of the initial game; that is, the initial, noncooperative game always admits one Nash equilibrium that generates socially efficient investments.[14] This result is general and applies to ITU models as well; it has many applications, including for the econometric estimation of structural matching models.

[14]Other equilibria may, however, exist. For instance, if the surplus is supermodular and all agents but one underinvest, the last person's incentives to invest are suboptimal, which may lead to a coordination failure equilibrium.

Iyigun & Walsh (2007) build on this intuition to develop a theoretical model in which agents first invest in HC (therefore, in their future income), then they match on the marriage market. They derive the equilibrium sharing rule, which determines intra-household allocation, as a function of the distribution of premarital endowments and the sex ratios in the market; they show that all such sharing rules support efficient outcomes.

2.2.4.3. Risk sharing. It must be emphasized that TU is an ordinal property. As such, it is compatible with concave Von Neumann–Morgenstern (VNM) utilities, that is, with risk aversion, as illustrated by the following example. Assume that spouses consume a private and a public good [C_i ($i = 1, 2$) and Q, respectively], and individual utilities are

$$u_i(C_i, Q) = (1-\alpha)\ln C_i + \alpha \ln Q.$$

Individuals are each endowed with a random income; once married, they can make ex ante efficient contracts, involving in particular risk sharing. Let $\tilde{x}_1$ and $\tilde{x}_2$ respectively denote her and his income; for any realizations (x_1, x_2), the couple's budget constraint is given by

$$C_1 + C_2 + PQ = x_1 + x_2. \qquad 11.$$

Note, first, that any ex post efficient allocation satisfies

$$C_1 + C_2 = (1-\alpha)(x_1 + x_2), PQ = \alpha(x_1 + x_2).$$

Let $C(x_1, x_2)$ denote 1's private consumption [so that 2's is $(1-\alpha)(x_1 + x_2) - C(x_1, x_2)$]. By the mutuality principle, C only depends on $X = x_1 + x_2$; ex ante efficiency requires that $C(X)$ maximizes some weighted sum of individual expected utilities,

$$\max Eu_1 + \mu E u_2 \text{ (for some } \mu > 0),$$

under the resource constraints in Equation 11. First-order conditions imply that

$$\frac{\partial u_1}{\partial C} = \frac{1-\alpha}{C} = \mu \frac{\partial u_2}{\partial C} = \mu \frac{1-\alpha}{(1-\alpha)X - C},$$

which gives

$$C_1 = C = \frac{1-\alpha}{1+\mu} X \text{ and } C_2 = (1-\alpha)X - C = \mu \frac{1-\alpha}{1+\mu} X.$$

Finally, individual expected utilities are equal to

$$Eu_1 = (1-\alpha)\ln\left(\frac{1-\alpha}{1+\mu}\right) + \alpha \ln \frac{\alpha}{P} + \int \ln X dF_X(X) \text{ and}$$
$$Eu_2 = (1-\alpha)\ln\left(\mu\frac{1-\alpha}{1+\mu}\right) + \alpha \ln \frac{\alpha}{P} + \int \ln X dF_X(X),$$

where F_X denotes the distribution of X. We see that

$$\exp\left(\frac{1}{1-\alpha}Eu_1\right) = \left(\frac{1-\alpha}{1+\mu}\right)\left(\frac{\alpha}{P}\right)^{\frac{\alpha}{1-\alpha}} \exp\left(\frac{1}{1-\alpha}\int \ln X dF_X(X)\right) \text{ and}$$
$$\exp\left(\frac{1}{1-\alpha}Eu_2\right) = \left(\mu\frac{1-\alpha}{1+\mu}\right)\left(\frac{\alpha}{P}\right)^{\frac{\alpha}{1-\alpha}} \exp\left(\frac{1}{1-\alpha}\int \ln X dF_X(X)\right),$$

so that the set of ex ante Pareto efficient allocations satisfies

$$\exp\left(\frac{1}{1-\alpha}Eu_1\right)+\exp\left(\frac{1}{1-\alpha}Eu_2\right)=(1-\alpha)\left(\frac{\alpha}{P}\right)^{\frac{\alpha}{1-\alpha}}\exp\left(\frac{1}{1-\alpha}\int \ln X dF_X(X)\right),$$

which is a TU form for $\tilde{u}_i=\exp\left(\frac{1}{1-\alpha}Eu_i\right),\ i=1,2.$

A general result was derived by Mazzocco (2004) and Schulhofer-Wohl (2006). They show that if VNM utilities belong to the Identical Shape Harmonic Absolute Risk Aversion (ISHARA) family, then the couple behaves as a single decision maker, and the resulting preferences satisfy a TU property. The ISHARA property requires (*a*) that each absolute risk aversion index be a harmonic function of income,

$$-\frac{u_i''(x)}{u_i'(x)}=\frac{1}{a_i x+b_i},$$

and (*b*) that the income coefficient a_i be the same for all individuals. In particular:

1. If individual utilities all belong to the Constant Absolute Risk Aversion (CARA) family, the property is satisfied (with $a_i=0$ for all i).
2. In the Constant Relative Risk Aversion (CRRA) case, which corresponds to $b_i=0$ for all i, TU obtains if and only if all agents have the same coefficient of relative risk aversion.

2.2.4.4. Asymmetric information. A recent and interesting literature extends the concept of stability to contexts involving asymmetric information. This is not an easy task, because from any on- or off-path event (e.g., the existence of a blocking pair), individuals will try to infer information about the individual's unobserved characteristics. Indeed, a consistent stability criterion requires, in addition, the Bayesian consistency of three probabilistic beliefs: exogenously given prior beliefs, off-path beliefs conditional on counterfactual pairwise blockings, and on-path beliefs for stable matchings in the absence of such blockings. The interested reader is referred to Liu et al. (2014) and Liu (2018) for a detailed analysis.

2.2.4.5. Divorce. While matching models are mainly used to model household formation and the resulting allocation of power within the couple, the same technology can, at little cost, be extended to discuss divorce—a point initially made by Weiss & Willis (1993, 1997). The standard assumption is that, once married, individuals observe the realization of a random shock θ (i.e., quality of the match) that enters the surplus generated by marriage over and above the economic benefits discussed above. Should the shock be negative enough, the spouses would have to choose between preserving the existing economic gains and remaining single, thus avoiding the negative shock. Technically, the realized surplus, normalizing at zero the utility of remaining single, is thus given by

$$S=s(\mathbf{x},\mathbf{y})+\theta,$$

and agents divorce if and only if

$$\theta\leq -s(\mathbf{x},\mathbf{y}),$$

which occurs with probability $F_\theta[-s(\mathbf{x},\mathbf{y})]$. The ex ante expected surplus is thus represented by

$$ES(\mathbf{x},\mathbf{y})=(1-F_\theta[-s(\mathbf{x},\mathbf{y})])[s(\mathbf{x},\mathbf{y})+E(\theta\mid\theta\geq -s(\mathbf{x},\mathbf{y}))],$$

which determines the equilibrium on the marriage market.[15] It can be stressed that, in this framework, a one-to-one correspondence exists between divorce probability and the size of the surplus (for instance, these models typically predict a lower divorce rate among wealthier couples). From an empirical perspective, this suggests that divorce probability can be a good proxy for marital gains. Note, finally, that the previous argument implicitly assumes that divorced individuals remain single. If remarriage can occur, then the corresponding (remarriage) market must also be modeled (see, for instance, Chiappori & Weiss 2006, 2007).

2.3. Alternative Models: Theory

Two main families of alternative models can be found in the literature.

2.3.1. Imperfectly transferable utility models. TU models rely on a highly specific property—namely, that for a well-chosen cardinalization of individual preferences, the Pareto frontier is an hyperplane orthogonal to the unitary vector for all prices and incomes. In the most general framework, this property is not satisfied, and the model must be reformulated as follows. The Pareto set is defined by an equation of the form

$$U \leq \Phi(\mathbf{x}, \mathbf{y}, V), \tag{12.}$$

where $U(V)$ is her (his) utility and Φ is nondecreasing in v. A matching is still defined as a 3-tuple (h, u, v) where the marginals of measure h are F and G, respectively, and Equation 12 is satisfied with equality whenever $h(\mathbf{x}, \mathbf{y}) > 0$. Stability requires, moreover, that

$$u(\mathbf{x}) \geq \Phi(\mathbf{x}, \mathbf{y}, v(\mathbf{y})) \ \ \forall\, x, y, \tag{13.}$$

with the same interpretation as in the TU case. In particular, we have that

$$u(\mathbf{x}) = \max_y \Phi(\mathbf{x}, \mathbf{y}, v(\mathbf{y})) \ \ \forall\, x, \text{ therefore } u'(\mathbf{x}) = \frac{\partial \Phi}{\partial x}(\mathbf{x}, \mathbf{y}, v(\mathbf{y})). \tag{14.}$$

However, stability is no longer equivalent to surplus maximization; in fact, the notion of total surplus is not defined in this framework. It follows that existence of a stable matching cannot be established as simply as in the TU case (see Greinecker & Kah 2019 for a general result). Moreover, in the one-dimensional case, supermodularity (in the sense that $\partial^2\Phi/\partial x \partial y > 0$) is no longer sufficient to guarantee assortative matching; the conditions also involve another cross derivative, namely, $\partial^2\Phi/\partial x \partial v$. Intuitively, $\partial\Phi/\partial v$ represents the exchange rate between his and her utility; unlike the TU case, this rate is not constant, and the sign of $\partial^2\Phi/\partial x \partial v$ indicates how it changes with the wife's characteristics (see Legros & Newman 2007, Chiappori 2017).

2.3.2. Search models. Search models play a crucial role in labor economics; modern approaches recognize that unemployment is due, at least in part, to search frictions on the labor market. More recently, they have been extensively applied to marriage.[16] The basic framework

---

[15]By the envelope theorem applied to $ES(x, y)$, we have

$$\frac{\partial ES}{\partial x} = (1 - F_\theta[-s(x, y)])\frac{\partial s}{\partial x}.$$

It follows that if s is increasing in x and y and supermodular, then ES is also supermodular.

[16]Note, however, that one of the early and seminal papers in the search literature, by Mortensen (1988), explicitly mentions the marriage market as a possible application of search models.

is generally similar to the one developed above. Women (resp. men) are each characterized by a vector of characteristics $\mathbf{x} \in X$ ($\mathbf{y} \in Y$), and a matching between $\mathbf{x}$ and $\mathbf{y}$ generates a surplus $s(\mathbf{x}, \mathbf{y})$ that must be shared between spouses in a TU framework. The main difference, however, is the dynamics induced by the search frictions. Meetings between men and women occur randomly; when a meeting takes place, each partner must choose between marrying or continuing to search, the latter option involving a cost related to the individual's discount factor. In the standard version, a fixed percentage of individuals disappear ("die") in each period, and/or some couples may divorce, whereas new individuals continuously emerge (as singles). Technically, most investigations concentrate on the steady state. (For a detailed presentation see, for instance, Pissarides 1990, Shimer & Smith 2000, Smith 2011.)

Equilibrium equations reflect (*a*) a stationarity condition, according to which the flows in and out of each category of agents and marital status cancel off; and (*b*) an optimization condition, reflecting the fact that individuals' marriage decisions always maximize (expected) utility. In practice, the system is described by a set of value functions, reflecting the current value of future expected utility of individuals of a given type conditional on their marital status; dynamic optimality translates into Bellman equations on these functions. A meeting between Mrs. $\mathbf{x}$ and Mr. $\mathbf{y}$ ends up in marriage if and only if the generated surplus can be shared in such a way that both partners are strictly better off than they would if they remained single. Given the TU assumption, this is equivalent to saying that the sum of value functions when married is larger than the sum of value functions when single.

A few remarks can be made at this point. First, the TU assumption is typically made both before and after marriage; in practice, this may require specific assumptions regarding individual preferences, particularly in the presence of public goods. Second, should a meeting end up in marriage, there exists in general a continuum of intra-household equilibrium allocations; indeed, since the sum of utilities is generally strictly larger when married than when single, agents are de facto in a situation of bilateral monopoly. As a result, the model must include assumptions on the bargaining process used to allocate welfare within the couple. Note also that, in this context, commitment issues become important, particularly when divorce is allowed for. Indeed, whether, following an exogenous shock, agents will decide to remain married typically depends on (their expectations regarding) the future of the relationship—which in turn depends, among other things, on the agents' ability to (fully or partly) commit. Third, the meeting technology can be formulated in different ways. While most models assume that only singles search, some versions allow for "on-the-job search." In some frameworks, agents may affect the meeting process by varying their search intensity—the latter then reflecting the expected gain from new meetings. In general, the matching technology is assumed to exhibit constant returns; in some cases, though, increasing returns may be considered (then multiple equilibria may coexist; see, for instance, Diamond 1982, 1984).

3. MODELS OF THE MARRIAGE MARKET: ECONOMETRICS

Let me now present the main tools that have been applied to the econometric analysis of the marriage market (a detailed analysis can be found in Chiappori & Salanié 2016 and Galichon & Salanié 2017). A first task of any empirical exploration of marriage markets is to reconcile the somewhat simplistic predictions of basic theoretical models with the complex patterns we find in real data. Consider, for instance, matching on education (or HC), a question that will be abundantly discussed below. Theory strongly suggests that matching should be positive assortative. However, positive assortative matching (PAM), in a theoretical sense, would imply very specific—and, as a matter of fact, very unrealistic—matching patterns. Assume, to keep things simple, that there exist

		Women	
		U	E
Men	U	$1 - n$	$n - m$
	E	0	m

Figure 1

Matching patterns under pure positive assortative matching (total population size normalized to 1). Abbreviations: E, educated; *m*, proportion of educated men; *n*, proportion of educated women; U, uneducated.

only two levels of education, educated (E) and uneducated (U). If the number of educated people is the same among men and women, then PAM implies that in all couples, husband and wife have the same education. If the numbers differ—for example, the proportion among married couples of educated women, *n*, is larger than that of men, *m*—then there would also be couples in which she is educated and he is not; but an educated man could not possibly marry an educated woman at a stable matching. In other words, the matrix of marital patterns that describes the distribution of married couples by education of the spouses could have at most one non-diagonal term different from zero (see **Figure 1**).

Obviously, actual matching patterns never exhibit such an extreme form. The number of couples located on the diagonal of the matrix—i.e., the proportion of homogamous marriages—is typically higher than what would obtain under random matching, suggesting some assortativeness; but no cell remains empty.

The literature has followed two paths to capture the discrepancy between theory and observation. One invokes frictions on the marriage market; for instance, the presence of couples where he is educated and she is not may stem from the fact that, after having (randomly) met an uneducated potential spouse, educated men simply stop searching for a better mate. An alternative approach maintains the frictionless structure but recognizes that matching is typically multidimensional: Real-life matches are never exclusively based on education, as other characteristics are also involved (e.g., tastes or physical attractiveness). These additional traits, however, are observed by the partners but not in general by the econometrician; the latter must then take into account the resulting unobserved heterogeneity in the estimation process. I present both approaches below, starting with the frictionless version.

3.1. Frictionless Matching Under Transferable Utility

The econometrics of TU matching models follow two main approaches.

3.1.1. The separable extreme value approach and its extensions. The first comprehensive investigation of frictionless matching with unobservable characteristics is due to Choo & Siow (2006). The data they use exclusively consist of matching patterns. Specifically, agents are assumed to belong to a finite (and actually small) set of categories; the market is fully summarized by the matching table indicating, for each gender and each category (e.g., level of education), the proportion of individuals married to a spouse in each category plus the proportion of singles. For convenience, I will stick to the 2×2 example just mentioned (generalization to any number of categories is straightforward); the resulting table is shown in **Figure 2**.

The question, now, is to provide a theoretical interpretation of these data. The basic idea is quite simple: adding to the surplus a random term reflecting unobserved characteristics. With

		Women		
		U	E	Singles
Men	U	μ^{UU}	μ^{EU}	$\mu^{\emptyset\mathrm{U}}$
	E	μ^{UE}	μ^{EE}	$\mu^{\emptyset\mathrm{E}}$
	Singles	$\mu^{\mathrm{U}\emptyset}$	$\mu^{\mathrm{E}\emptyset}$	

Figure 2

Empirical matching patterns. Abbreviations: E, educated; μ^{XY}, proportion of (X,Y) couples; $\mu^{\mathrm{X}\emptyset}$, proportion of singles with education X; U, uneducated.

the previous notations, thus, the surplus created by the match of Mrs. i with education I and Mr. j with education J, (I, $J \in \{\mathrm{U, E}\}$), is the sum of two terms: a deterministic component $z(I, J)$ that summarizes the impact of education and an additively separable stochastic term ε_{ij} reflecting unobserved heterogeneity. That is, we obtain

$$s(i, j) = z(I, J) + \varepsilon_{ij}. \tag{15}$$

Similarly, the utility of singles takes the form

$$s(i, \emptyset) = z(I, \emptyset) + \varepsilon_{i\emptyset} \text{ and } s(\emptyset, j) = z(\emptyset, J) + \varepsilon_{\emptyset j},$$

where $\varepsilon_{i\emptyset}$ and $\varepsilon_{\emptyset j}$ respectively represent i's and j's idiosyncratic preferences for singlehood. Clearly, $z(I, \emptyset)$ and $z(\emptyset, J)$ can be normalized to zero.

A form like the one in Equation 15 raises, however, several difficulties. One relates to the correlation structure for the ε's; this is crucial, since the properties of the resulting stable matching will in part depend on the super- (or sub)modularity introduced by the stochastic terms. A general correlation structure for the ε's cannot be reasonably estimated from the data; yet, independence is probably a poor choice (see Chiappori & Salanié 2016 for a precise discussion). Moreover, from a theoretical perspective, the surplus maximization problem in Equation 6 becomes stochastic; as a result, the dual variables in Equation 7 are also stochastic, and very little is known about their distribution (see Chiappori 2017 for an extensive discussion and Chiappori et al. 2019d for an evaluation of potential biases).

To address these issues, Choo & Siow (2006) introduce the following assumption, directly borrowed from Dagsvik (2000):

Assumption 1 (separability). The random term ε_{ij} is of the form

$$\varepsilon_{ij} = \alpha_i^J + \beta_j^I, \tag{16}$$

where the random variables $\alpha_i^J, J \in \{\mathrm{U, E}\}$ and $\beta_j^I, I \in \{\mathrm{U, E}\}$ are independent.

The random term ε_{ij} is thus additively separable in two terms, α_i^J and β_j^I. The former is woman specific but only depends on the husband's education; the latter is man specific but only depends on the wife's education. It may be that each woman has idiosyncratic preferences on her husband's education (and vice versa). Alternatively, a large α_i^J may reflect the fact that Mrs. i has some specific quality that is particularly valued by husbands from category J (for example, she is a very good classical pianist, a talent that educated men value more).

Anyhow, this assumption implies an important property, discussed below.

Proposition 1 (Choo & Siow 2006, Chiappori et al. 2017). Assume the surplus has the form given in Equation 15, where the ε's satisfy the separability property in Equation 16.

Then there exists 2×4 numbers, U^{IJ} and V^{IJ}, with $I, J \in \{\text{U, E}\}$, such that for all I, J, we obtain

$$U^{IJ} + V^{IJ} = z(I, J). \qquad 17.$$

If $i \in I$ is married to $j \in J$ at the stable matching, then the payoffs are given by

$$u_i = U^{IJ} + \alpha_i^J \text{ and } v_j = V^{IJ} + \beta_j^I.$$

Proof. See Chiappori et al. (2017). ■

In other words, the payoff of Mrs. i is the sum of two components: a deterministic effect U^{IJ}, which only depends on the categories of i and i's spouse, and the stochastic, i-specific term α_i^J. In particular, we have exactly characterized the stochastic structure of the dual variables of the surplus maximization program: It is precisely that of the stochastic terms in the separable structure, up to a shift in mean. Moreover, Proposition 1 also applies to singles under the normalization $U^{I\emptyset} = V^{\emptyset J} = 0$ for all I, J; then, we obtain $u_i = \alpha_i^{\emptyset}$ and $v_j = \beta_j^{\emptyset}$.

From an empirical perspective, Proposition 1 has a very important consequence: The matching problem can be decomposed into a series of independent maximization problems, as stated by the following result.

Corollary 1. Under the assumptions of Proposition 1, a woman i is matched with a spouse in J if and only if

$$U^{IJ} + \alpha_i^J \geq U^{IL} + \alpha_i^L \text{ for } L \in \{\emptyset, \text{U}, \text{E}\}, \qquad 18.$$

and i remains single if and only if

$$\alpha_i^{\emptyset} \geq U^{IL} + \alpha_i^L \text{ for } L \in \{\emptyset, \text{U}, \text{E}\}. \qquad 19.$$

A similar conclusion holds for men.

Here, u_i can be seen as the utility that i will receive on the matching market; i's choice can simply be modeled as selecting the spousal category for which the utility is highest (remember that under separability, i only cares about her spouse's category, not his precise identity). In particular, and unlike the general case (as described by Equations 4 and 5), the marital choice of a woman with a given education can be modeled independently of the (endogenous) equilibrium utilities of her potential spouses. This turns out to be extremely useful for empirical applications, because it allows to directly transpose the whole machinery of discrete choice models to the (a priori more complex, since two-sided) matching problem. The thresholds U^{IL} and V^{LJ} can be econometrically identified from a series of standard discrete choice models. In practice, Choo & Siow (2006) assume that the α's and β's are extreme value distributed, and they run a series of multilogit regressions (for each gender/education). Then one can compute the sum $U^{IL} + V^{LJ} = z(I, J)$ and evaluate the respective importance of economic and noneconomic factors (as respectively summarized by the deterministic and stochastic components in Equation 15).

If we assume that the distributions of the α's and β's are fully known ex ante, the model is exactly identified. That is, even under these very strong parametric assumptions, the model cannot generate testable restrictions on matching patterns; and any generalization (e.g., allowing for heteroskedasticity) leads to non-identifiability. Galichon & Salanié (2015) show how the extreme value assumption can be relaxed. In fact, the same approach can be used with basically any distributions for the random terms, provided, again, that these are perfectly known ex ante—and again, the model is then exactly identified.

Two paths have been followed to generate testable restrictions. The multi-market approach simultaneously analyzes several marriage markets ($t = 1, \dots, T$) that share common structural characteristics; in general, it is assumed that (*a*) the random terms are drawn from the same distribution irrespective of the market, and (*b*) the deterministic components $z_t(I, J)$ exhibit a similar tendency to assortativeness, in the sense that the supermodular cores $z_t(I, I) + z_t(J, J) - z_t(I, J) - z_t(J, I)$ are identical across markets. Alternatively, one can use direct information on the surplus function. The idea here is that, once married, the household maximizes the sum of utilities (a consequence of the TU assumption); using standard techniques, one can recover the maximand—which is exactly the surplus—from observed behavior. Then, identification obtains from the joint observation of matching patterns and postmarital behavior. In practice, this typically requires the joint identification of the various parameters of the model from data on matching patterns and postmarriage behavior (typically labor supply); a convenient method relies on simulated moments.

Both paths are illustrated in the next section.

3.1.2. The rank-order approach. An influential approach has been proposed by Jeremy Fox in several papers (Fox 2008, 2010; Fox & Bajari 2013.) A common feature of these papers is that they rely on a rank-order property, although the property slightly differs across papers. The basic intuition relies on the duality results described above, whereby stable matchings maximize total surplus; the econometric translation postulates a direct relationship between the likelihood of observing a given matching and the surplus this matching generates. Specifically, assume, as above, that the total surplus is the sum of an observable component (involving only observable traits) and a random term. Fox (2010) considers matchings from a collection of nonoverlapping, finite-size markets and assumes that the mean surplus function is the same in all of these markets. Consider, now, two markets in which the distributions of observable characteristics are identical (while those of unobserved characteristics may differ). The rank-order property postulates that the distributions of the random terms are such that a matching is more likely to be observed if and only if the observable component of the surplus is larger. Under this rank-order property, if the distribution of observed characteristics has continuous support, then comparing such markets identifies the supermodular core of the surplus function (up to an increasing transform); interestingly, this approach lends itself very well to an estimation method based on inequalities. Thus, the rank-order approach seems to be a natural extension of standard discrete choice models, with the observable component of the surplus playing the role of the latent variable. The problem, however, is that, due to the two-sided nature of the model, this intuition works only in very specific contexts.

Fox & Bajari (2013) exploit a related but different version of the rank-order property, initially introduced by Fox (2008, section 3.1), that also applies to data on a single large matching market. Let $\eta(x, y)$ denote the number of matches between women of type x and men of type y on the market. The property states that

$$s(x,y) + s(x',y') \geq s(x,y') + s(x',y)$$

if and only if

$$\eta(x,y)\eta'(x',y') \geq \eta(x,y')\eta'(x',y).$$

Unlike the initial notion, this version can be derived from more primitive assumptions; in particular, the corresponding rank-order property applies in the separable extreme value (SEV) model.

3.2. Other Matching Models

Finally, several alternative approaches do not rely on the TU assumption.

3.2.1. Imperfectly transferable utility matching. Although the econometric analysis of ITU models is recent, it is possible to present some interesting advances. Just as in the TU case, the main difficulty is to account for unobserved heterogeneity; in practice, this requires introducing a random component in the matching process, while keeping as much theoretical and empirical tractability as possible. As it happens, this can be done through a direct extension of the SEV method, as recently demonstrated by Galichon et al. (2019). The intuition underlying their work can be summarized as follows. Assume, as before, that agents belong to a finite set of categories (e.g., education classes). Take a given woman i, belonging to category I, and a given man j, belonging to category J. The feasibility constraint limiting the pair $(u(i), v(j))$ of utilities that i and j could reach if matched together is now of the form

$$u(i) = F(I, J, v(j) - \beta_j^I) + \alpha_i^J, \qquad 20.$$

where α_i^J and β_j^I are random coefficients, the interpretation of which is exactly the same as in the TU case. Note that, as in the SEV model, the random terms enter additively and satisfy a strong separability property: α_i^J only depends on the category J of i's spouse, not on his identity, and the same holds for β_j^I. Also, the TU case is obviously nested in this framework.

Proposition 1 can be directly extended to that framework. There exists $2 \times K^2$ numbers, U^{IJ} and V^{IJ}, with $I, J = 1, \ldots, K$, such that for all I, J, we have that $U^{IJ} = F(I, J, V^{IJ})$, and such that if $i \in I$ is married to $j \in J$ at the stable matching, then the payoffs are $u_i = U^{IJ} + \alpha_i^J$ and $v_j = V^{IJ} + \beta_j^I$. Thus the matching model still boils down to a series of standard discrete choice models at the individual level. In practice, the authors describe the Pareto set through a distance function $D(u, v)$ that measures the signed distance of any pair of utilities (u, v) from the efficient frontier. They then show how the corresponding model can be estimated in different ways, including maximum likelihood.

Finally, the practical estimation of models of this type either exploits the theoretical results above or relies on simulated moments.

3.2.2. Search models. The econometrics of search models applied to marriage markets largely reflects two background influences, namely macroeconomics and labor. The macro approach does not formally describe a stochastic structure for the model under consideration. Instead, it constructs a theoretical model that is then simulated using specific values for the relevant parameters. In practice, some parameter values (e.g., discount rate, life expectancy, individual preferences, etc.) are picked up from existing studies. The remaining parameters (if any) may then be calibrated to best approximate a set of data targets. Technically, the calibrated parameters minimize the distance between the targets predicted by the model and the actual ones, although the definition of the corresponding metrics does not typically refer to the stochastic properties of the series under consideration.[17]

Models belonging to the labor tradition, on the other hand, are more explicitly structural. While identification strategies tend to be more model specific than in the frictionless case, some common features emerge. In general, the characteristics of the meeting process, as well as marriage and divorce probabilities, can be directly recovered from observed flows in and out of marriage, and parameters governing preferences and/or domestic production are identified from observed behavior in a standard way. The remaining parameters (e.g., those governing the bargaining process or the variance of match-specific shocks) can then be derived from moment estimators.

[17] In particular, they may not be minimum distance estimators in the statistical sense.

4. MODELS OF THE MARRIAGE MARKET: APPLICATIONS

Matching models have been applied to a host of specific issues related to the marriage market. An exhaustive presentation exceeds by far the scope of the current article. Here, I will only try to describe some recent and significant contributions.

4.1. Frictionless Matching Under Transferable Utility: The One-Dimensional Case

One-dimensional matching remains the most common framework in the empirical literature. The choice of the relevant characteristic is, however, far from obvious. While data reveal, in most countries, a significant level of assortative matching on income, the economic interpretation of this pattern is complex: Realized (postmarriage) individual incomes are largely endogenous, if only because they reflect labor supply decisions, which in turn depend on the spouses' characteristics and behaviors. HC appears to be a more sensible choice; yet, one must keep in mind that HC investments, while often decided before marriage, are partly influenced by their expected consequences on the marriage market.

4.1.1. Reduced form approaches. An early, mostly sociological literature has documented a strong tendency to assortative matching on the marriage market in terms of education, income, and/or HC.[18] Several authors (Burtless 1999, Greenwood et al. 2014), in particular, have argued that increases in assortative matching had a direct impact on the evolution of inequality over the last decades, although the mere definition of "increased assortativeness" is far from obvious, even in the simplest frameworks.[19] Analyzing French data, Frémeaux & Lefranc (2017) estimate that the effect of assortative matching on household potential earnings amounts to 10–20% of observed inequality. Using a large data set of Italian tax records, Fiorio & Verzillo (2018) argue that assortative matching concentrates mostly at the top of the distribution.

Some contributions consider other characteristics. For instance, Dokko et al. (2015) document substantial positive assortative matching with respect to credit scores, even when controlling for other socioeconomic and demographic characteristics; moreover, the couples' match quality in credit scores, measured at the time of relationship formation, is highly predictive of subsequent separations. Harris & Cronin (2014) examine how a single woman's investment in healthy body weight is affected by the quality of single men in her marriage market. They find that when potential mate quality in a marriage market decreases, single black women invest less in healthy body weight, as would be suggested by the supermodularity of the surplus function.

While these works are mostly descriptive, other contributions develop a one-dimensional matching model to generate testable predictions that can be taken to data in reduced form. Iyigun & Lafortune (2016) propose a multi-period frictionless matching model where educational and marriage decisions are endogenous, under two assumptions: Marriage requires a fixed cost, and married couples cannot study simultaneously. They show that exogenous delays in marriage age caused by minimum age laws decreased the educational difference within a couple while increasing the spouses' educational attainment; in particular, this mechanism explains the patterns

[18]Recent contributions include work by Mare (1991), Schwartz & Mare (2005), and Schwartz (2010). These studies, however, mainly use log-linear models for contingency tables, the theoretical foundations of which are unclear.

[19]Distributions with identical marginals are easy to compare (intuitively, an increase in assortativeness translates into a higher number of people on the diagonal). Things are, however, more complex, even with two categories (e.g., educated and non-educated) for each gender and without singles, when marginals differ (see, e.g., Chiappori et al. 2019a for a precise discussion).

followed by age at first marriage and why the gender education gap tracked an inverted-U path in the United States over the twentieth century.

Ashraf et al. (2020) analyze a particular feature of some marriage markets, namely the existence of bride price. In their model, individuals differ in innate ability and match on education, as the parents' decision to educate their children is directly related to the expected impact on the marriage market. Revisiting Indonesia's school-construction program, they find that bride price has large positive effects on girls' schooling: A daughter's education, by increasing the amount of money parents receive at marriage, generates an additional incentive for parents to educate their daughters.

In a recent contribution, Bisin & Tura (2019) study the cultural integration of immigrants in Italy. They estimate a model in which individuals exclusively match along ethnic dimensions and invest into children's socialization. The model generates strong predictions (e.g., heterogamous households invest more when divorced than when married, whereas the opposite is true for homogamous households), which are tested on rich administrative data on Italian marriages.

4.1.2. Quasi-natural experiments. Other contributions exploit policy reforms that affect the marriage market, resulting in changes in matching patterns. Chiappori et al. (2017a) study the effect of a reform granting alimony rights to cohabiting couples in Canada. The key theoretical insight is that the reform should affect existing couples and couples-to-be differently. In existing couples, the intended beneficiary unambiguously gains; for couples not yet formed, however, equilibrium on the marriage market requires offsetting intra-household transfers. The empirical analysis confirms these predictions.

Persson (2020) analyzes a 1989 reform in Sweden, which essentially suppressed the "survivor insurance" that had been an implicit part of marriage contracts until then (i.e., upon the death of one household member, the surviving spouse was granted a lifetime annuity). Persson shows that the reform, the financial impact of which would not be felt before several decades, immediately affected marital patterns. Specifically, the consequent reduction in marital surplus decreased the number of marriages and increased the steady-state rate of separation from cohabiting unions. Some cohabiting couples, however, were given the opportunity to benefit from the old system if they married before a given deadline. This resulted in a considerable surge in marriages for this subgroup; by a selection effect, the corresponding divorce rate is expected to be higher than for the rest of the population, a prediction that is confirmed empirically. Overall, the response to the reform is also found to increase with the expected gain, whether the latter is estimated from observable characteristics at marriage or from realized (ex post) mortality. Finally, survivors insurance constituted de facto a public subsidy on matches with highly unequal earnings (capacities); as predicted, the long-term impact of its suppression has been an increase in assortative matching.

4.1.3. Investment in human capital, assortative matching, and inequality: structural approaches. Finally, several recent contributions use an explicitly structural approach to investigate how the marriage market interacts with HC investments. A well-known puzzle in the United States is the spectacular asymmetry between genders in the demand for higher education. Over the last decades, the college premium has surged, boosting the returns to investments in college education and beyond. Not surprisingly, the proportion of women with a graduate degree has vastly increased in response; however, the proportion of men has stagnated at best. Chiappori et al. (2009) suggest that this asymmetry might be related to the marriage market. They define the "marital college premium" as the additional gain perceived by educated people on the marriage market (with respect to less educated ones) over and above the benefits received on the labor market. Using a one-dimensional matching model, they show how the evolution of marital patterns over the period is compatible with a decrease (resp. increase) in the male (female) premium. Intuitively, in a world where few women are educated, the cost, in terms of marital prospects, of not acquiring an education is low, since many uneducated women will "marry up" at equilibrium. Not so, however, when most women go to college (or beyond): Nurses can hardly expect to marry a surgeon when faced with a fierce competition from female doctors.[20]

This idea is taken to data by Chiappori et al. (2017b), who estimate a semistructural model of the US marriage market over 30 years. The identifying assumptions, as discussed in subsection 3.1.1, are that the distribution of stochastic terms does not change over the periods under consideration, and that the supermodular core of the surplus function varies in a specific manner. In practice, they first test the null of constant supermodularity by assuming that the deterministic component of the surplus for cohort c takes the form

$$Z_c^{IJ} = \zeta_c^I + \xi_c^J + Z^{IJ}. \tag{21}$$

This version is strongly rejected for the white population (although it is not for blacks). Next, they allow for a linear trend,

$$Z_c^{IJ} = \zeta_c^I + \xi_c^J + a^{IJ} + b^{IJ} \times c, \tag{22}$$

where b^{IJ} is an (education-specific) supermodular trend. The fit with actual patterns is considerably improved; moreover, the matrix $B = ((b^{IJ}))$ is supermodular, indicating stronger preferences for assortative matching over time.[21]

Why such an evolution? In the authors' explanation, HC investments play a central role. Given the spectacular rise of college and postcollege premiums in the United States, investing in children's HC now constitutes a major component of marital gains; time use data indeed reveal that the time spent by parents with their children has significantly increased, particularly among more educated households. If, as empirical works clearly suggest, the parents' own HC is an important input in the process, theory predicts that assortative matching should increase in response. In the long run, this mechanism may mechanically amplify inequality cohort after cohort, generating what Chiappori (2017) calls an "inequality spiral."

Chiappori et al. (2018a) analyze the British marriage market using a structural model based on a three-stage game: Individuals independently invest in their HC, then they match on the marriage market, and the resulting households (singles or couples) finally consume and supply labor over several subperiods while facing various exogenous shocks. The model introduces three innovations. First, the last stage, which describes the household after marriage, is a dynamic household labor supply model with stochastic fluctuations in wages, where marital surplus is generated by risk sharing and the joint consumption of some unspecified public good. This framework allows, using standard methods, to recover the household's total utility from observed behavior; therefore, the (second-stage) matching game can be estimated using both matching patterns and the surplus function thus identified. Second, HC is modeled as the product of two components, namely education and some unobserved ability, the joint distribution of which can be estimated from the observed wage dynamics at the individual level. Lastly, a by-product of the matching game is the derivation of individual (lifetime, expected) utilities as a function of education; this, in turn, leads to a structural estimation of the demand for education. The model is estimated using simulated

[20]Note that the argument does not require multiple equilibria; in fact, in matching models under TU of this type, equilibria are generically unique. Yet the qualitative features of the unique equilibrium may exhibit large responses to minor changes in the fundamentals, an effect reminiscent of the social interaction literature.

[21]Eika et al. (2019) reach the opposite conclusion; however, their indicator for assortativeness has been criticized by Chiappori et al. (2019a).

moments; the technique relies in particular on the Nöldeke and Samuelson model of premarital investments described in subsection 2.2.4. The model is then used for policy simulations, including the impact of cheaper access to higher education.

4.2. Multidimensional Matching Under Transferable Utility

The empirical analysis of multidimensional matching models, although relatively recent, has evolved at a fast pace.

4.2.1. Index models. The simplest approach to multidimensional models relies on an index assumption, according to which an individual's multiple characteristics matter for the marriage market only through some one-dimensional index, as in Equation 9 above. Chiappori et al. (2012, 2020) show that this assumption, clearly restrictive, is however testable. Specifically, they consider a SEV model[22] in which (*a*) the deterministic component is quadratic and satisfies Equation 9, and (*b*) normally distributed stochastic terms enter through the scalar product with observable characteristics. They derive testable restrictions over the coefficients of the regressions of male-to-female characteristics (and vice versa); moreover, should these conditions be satisfied, the indices could be recovered (obviously up to an increasing transform). They apply this methodology to data from the Panel Study of Income Dynamics (see **https://psidonline.isr.umich.edu/**), with socioeconomic status being proxied by income or education and physical attractiveness being proxied by an individual's body mass index (BMI). The restrictions are not rejected; moreover, the estimated indices are linear.

4.2.2. Discrete/continuous characteristics. Convenient as they may be, index models only apply to very specific situations. Most of the time, the trade-offs between an individual's characteristics (as perceived by a potential spouse) vary with the spouse's characteristics, invalidating the index assumption. This fact has led to the development of "true" multidimensional models.

4.2.2.1. Smoking. A specific but often relevant case obtains when one characteristic is discrete while the other is continuous. In an early contribution, Chiappori et al. (2018b) consider a bidimensional model in which individuals are characterized by their (continuous) income and their (discrete) smoking status, under three assumptions: (*a*) Surplus is quadratic in income; (*b*) when one spouse (at least) is a smoker, total surplus is deflated by a fixed factor $\lambda < 1$; and (*c*) smoking prevalence is higher for men. They show that stable matchings typically involve randomization: While people at the top of the income distribution match in a perfectly assortative way, below some threshold nonsmoking women randomize between a nonsmoking partner with lower income and a richer smoker. The model generates several predictions that can be tested in reduced form; they appear to be well supported by the data.

4.2.2.2. Migrations. A particularly interesting application of these models relates to migrations: The status of migrant, a discrete characteristic, interacts with standard economic factors (income, education, etc.) in potentially interesting ways. Ahn (2017) studies the consequences of a sudden expansion of "marital migrations" from Vietnam to Taiwan. Between 1995 and 2000, due to the emergence of matchmaking intermediaries, the number of Vietnamese women marrying a Taiwanese husband surged from a few hundred to almost 15,000 per year; after a restrictive visa policy was implemented in 2004, however, the yearly number dropped to less than

[22]Their approach, however, can readily be extended to more general frameworks, including search models.

5,000. Ahn shows that the marital patterns closely follow theoretical predictions: The dominant form of cross-border marriage is Taiwanese men with Vietnamese women; cross-marrying Taiwanese men are selected from the middle of the socioeconomic status distributions, while Vietnamese women come from the top; and changes in costs of cross-matching affect this selection in the predicted manner. Even more interesting is the impact on the Vietnamese marriage market. The size of the cross-marriage phenomenon remains minor in Vietnam, a country of almost 100 million people. Yet, its impact on the marriage market in the most affected regions is large. Using a difference-in-difference approach, Ahn documents a significant decrease (resp. increase) in private consumption of male-exclusive (resp. female-exclusive) goods (e.g., smoking versus jewels) that affects all households. Ahn concludes that the integration of marriage markets may significantly affect the allocation of power within all couples, including those who do not cross-marry.

Using a similar theoretical framework, Nocera (2019) studies the immigration waves to the United States during the first half of the twentieth century, focusing in particular on the relationship between marital segregation and the labor market outcomes of immigrants. The choice between endo- and exogamy involves a trade-off between ethnic similarity and socioeconomic status among potential spouses. Theory predicts specific patterns of selection into exogamous marriage, of sorting in the marriage market, of intermarriage income gap, and of comparative statics with respect to the immigrant gender imbalance that are all supported by the data.

4.2.2.3. Women's demand for higher education. Finally, two contributions by Low (2014, 2019) investigate women's demand for higher education. The key insight is clear: Investing in a postgraduate degree by delaying marriage and childbearing has a higher cost for women, whose biological clock runs much faster. In her model, men differ by one characteristic (which can be interpreted as income, education, or social status), whereas women differ by two traits: innate ability and fertility. Women's innate ability can be boosted by access to graduate education, resulting in a larger stock of HC (and higher income) but lower fertility (which is modeled as discrete). In particular, if returns to HC are smaller and fertility loss is costly, stable matching may exhibit non-monotonic patterns, with the most able women investing in higher education but being matched to husbands of intermediate "quality," as top-income men prefer less skilled but more fertile women. When returns to HC increase and/or desired fertility is reduced, the stable matching switches to assortative matching on HC. Low argues that the evolution of the US marriage market over the last decades can be interpreted as a shift of this type, a claim well supported by the data.

4.2.3. Quadratic models. A specific version of bidimensional matching models, dating back at least to Tinbergen (1956), has recently attracted renewed attention. It relies on a basic Choo-Siow framework with continuous characteristics and two additional assumptions: (*a*) All marginal distributions are normal, and (*b*) the surplus is quadratic. In that case, the joint distribution of couples' characteristics corresponding to a stable matching is normal;[23] moreover, this joint distribution can be computed in closed form from the model primitives—and, conversely, the quadratic surplus function can be identified (also in closed form) from the optimal matching distribution.[24] This framework has been applied by Dupuy & Galichon (2014) to a unique survey of Dutch households

[23] Chiappori et al. (2017b) show that the opposite is also true: If the joint distribution of couples' characteristics corresponding to a stable matching is normal (implying in particular that marginal distributions are normal), then the surplus must be quadratic.

[24] In order to deal with continuous characteristics, a specific assumption on the matching technology is needed, which essentially extends the standard, multinomial logit choice model to a continuous framework (see, e.g., Dupuy & Galichon 2014, Bojilov & Galichon 2015).

containing information about education, height, BMI, health, attitude toward risk, and personality traits of the spouses. The estimates of the affinity matrix that defines the quadratic surplus lead to three important empirical conclusions. First, sorting does occur on several dimension; individuals face trade-offs between the attributes of their spouses, and these trade-offs depend on the individual's characteristics (implying that the game is not amenable to a one-dimensional index). Second, education explains the largest share of a couple's observable joint utility; yet personality traits explain another important part. Finally, different personality traits matter differently for men and for women.

Two recent contributions apply the Dupuy-Galichon technology to different data sets. Ciscato & Weber (2019) investigate how the nature of the gains from marriage has evolved in the United States from 1964 to 2017. Using Current Population Survey data, they find that the importance of education as a sorting dimension has increased since the 1960s, while the importance of age has decreased; their findings thus confirm the conclusions of Chiappori et al. (2017c) mentioned above. All in all, they estimate that, had marital preferences not changed, the 2017 Gini coefficient between married households would be lower by 6%. Finally, racial segregation on marriage markets, which used to be extremely strong in the 1960s, has much decreased in the 1970s but is nowadays slowly increasing. Ciscato et al. (2020) discuss the differences between same-sex and different-sex couples using a representative sample of Californian households for the period from 2008 to 2012. They find that same-sex couples are less segregated with respect to ethnic background. Moreover, education is the most important dimension of sorting for both male and female gay couples; on the contrary, age is the most important dimension of sorting for different-sex couples (education being second).

4.3. Matching Under Imperfectly Transferable Utility

Some recent contributions tackle the general ITU framework (see, for instance, Galichon et al. 2019). The main drawback of this model is that the equivalence between stability and surplus maximization is lost, which significantly complicates both the theory and the empirical analysis. Yet, TU models exhibit limitations that may in some contexts appear excessive. For instance, if the marital surplus stems from the presence of public goods, a direct consequence of TU is that household demand for public goods cannot possibly depend on the spouses' respective powers; thus, empowering women cannot possibly affect expenditures on children—a feature that seems unduly restrictive.

A recent and important contribution by Gayle & Shephard (2019) encompasses several crucial aspects of the ITU framework. In particular, the economic gain from marriage stems from the existence of a public good that is jointly produced by the parents (the estimation thus refers to time use data), and the ITU setting implies that household behavior depends on the intra-household allocation of power, itself an endogenous outcome of the matching game. This highly sophisticated construct is used to analyze economic policy issues, and in particular the design of optimal taxation.

Chiappori et al. (2019b) consider a three-stage model similar in spirit to the one presented in subsection 4.1.3, with a few twists. The marital gains (beyond risk sharing) are explicitly described as investments in children's HC; in particular, the corresponding production function is estimated, and the ITU structure implies that changes in intra-household allocation of decision power may affect these investments. Moreover, the model allows for divorce, which may be triggered by a combination of external factors (e.g., wage fluctuations) and internal shocks to the quality of the relationship. Individuals are assumed to agree, at the date of marriage, on some (explicit or implicit) prenuptial contract, which must be renegotiation-proof. In particular, at any date and in any state of the world, the continuation game generates Pareto-efficient outcomes and can therefore be fully summarized by one state variable describing the spouses' respective Pareto weights. Finally, the framework assumes limited commitment à la Mazzocco (2004), in the sense that the sequence of contracts is (second-best) optimal under the condition that individuals cannot commit not to divorce.

4.4. Search Models and Macroeconomic Applications

Finally, a recent but abundant literature applies search models to the marriage market. Chiappori & Weiss (2006, 2007) analyze a model of divorce and remarriage where the probability for a divorcee to meet a potential mate increases with the number of divorcees. Then multiple equilibria may coexist. When divorces are scarce, remarriage probability is low, increasing the cost of divorce; if, alternatively, many couples break up, then remarriage is easier, making separation less costly. The authors analyze the consequences on postdivorce transfers, particularly alimony, and argue that a higher aggregate divorce/remarriage rate can raise the welfare of children. Payments made by the divorced father typically stop when the mother remarries. Since frictions in the (re)marriage market leave room for bargaining between the new spouses, commitments from the ex-husband raise the bargaining power of the mother, which ultimately benefits the child as well. The higher the expected remarriage rate, the more willing each father will be to commit on such payments.

Bronson & Mazzocco (2018) investigate the strong and persistent negative relationship between changes in cohort size and marriage rates of both women and men in the United States over decades. They show that a standard matching model with search frictions cannot explain this pattern, since it produces a negative relationship for women but a positive relationship for men. If, however, the marriage surplus deteriorates with cohort size, then both relationships become negative and the model fits the observed qualitative patterns.

I will now briefly review search models belonging to the two approaches mentioned above.

4.4.1. The macro perspective. Introducing family-related considerations into macro models was, in itself, an important innovation; as Fernández et al. (2005, p. 273) argue, "The vast majority of macroeconomic models tend to assume the existence of infinitely lived agents (with no offspring) or a dynastic formulation of a parent with children." One of the first macro models devoted to family formation and dissolution and using a search technology is due to Aiyagari et al. (2000). They consider an overlapping generation (OLG) model in which individuals with heterogeneous productivity live for two periods as children, then two periods as adults—when they, in turn, have children. Women share their time between market work, time spent with children, and leisure, while men do not spend time with children (although both spouses enjoy the children's HC, which is produced from consumption and maternal time); consumption is also public in this framework, so that only leisure is privately consumed. Exogenous "love" shocks may trigger divorce, after which women remain single mothers and men stop caring about children. Finally, decisions within the household are noncooperative, which has two implications. First, couples typically overinvest in leisure and underinvest in children and consumption. Second, the intra-family allocation of welfare (which exclusively depends on individuals' choices of leisure) is not directly driven by equilibrium conditions on the marriage market, and it does not involve long-term considerations summarized by value functions and Bellman equations (as would be the case should decisions be made cooperatively); de facto, the model belongs to the NTU framework.[25] Specific as the model may be, however, when simulated with reasonable values for the parameters, it delivers conclusions that have been largely confirmed by the subsequent

[25] In particular, the model has several stable matchings.

literature. In particular, the economy exhibits significant inequality and a low level of intergenerational mobility. Children born in a low-productivity family will receive less investment, which depreciates their own ability distribution; their prospects on the marriage market are equally poor, implying less investment into their own children.

Greenwood et al. (2003) extend the previous model by relaxing the noncooperative decision assumption and considering Nash bargaining instead. The model generates interesting predictions regarding the interaction between income inequality and family size. Fertility declines with income, and single mothers are poorer and have the most children. Children raised by a single mother have a greater tendency (relative to other children) to grow up poor due to a lack of HC investment. Finally, the analysis of potential policy interventions is significantly affected by the general equilibrium features of the framework. For instance, child tax credits and child support payments improve children's welfare when fertility and divorce rates are held constant; however, the latter factors vary endogenously in response to the policy, essentially nullifying that impact.

Fernández et al. (2005) study the implications of marital sorting for household income inequality. In their framework, individuals make an early decision about becoming skilled or unskilled; skill acquisition is costly and requires borrowing on an imperfect capital market, with parents' income being used as a collateral. Then individuals meet randomly on a marriage market, generating some random match quality ("love") at each meeting. When a skilled individual meets an unskilled individual with high match quality, there will be a trade-off between forming a household with relatively lower consumption but high match quality and continuing the search. They show that several equilibria may coexist, some entailing a higher degree of sorting than others, and that high sorting is correlated with high inequality and low per capita income.

Greenwood & Guner (2004) and Greenwood et al. (2016) consider a search model aimed at explaining a number of stylized facts about the US marriage market over the last half-century: the decline in marriage rates and the increase in divorce rates, particularly for non-college-educated individuals, and the rise of positive assortative matching and income inequality (see also Fernández 2000, Fernández & Rogerson 2001). In their model, these evolutions are driven by two main forces. Technological progress in the household sector increases the reservation utility of singles while facilitating the labor supply of married women. Simultaneously, the returns to education surge, enticing more individuals to go to college, while a decline in the gender wage gap both encourages labor force participation by married women and makes singlehood more affordable for women. The model distinguishes the respective contributions of these two factors. Most of the rise in married women's labor force participation, the fall in marriages, and the rise in divorces is related to technological progress. Changes in the wage structure, on the other hand, explain the increase in assortative matching and educational attainment. Finally, both skilled-biased technical progress and the induced evolution of family structure explain the widening household income inequality.

4.4.2. Structural search models of the marriage market. A distinct, although closely related, line of research borrows most of its tools from search models of the labor market. Using British data, Goussé et al. (2017) construct and estimate a model of marriage formation and dissolution aimed at explaining the dynamics of household behavior, in particular labor supply and home production time inputs. Their model introduces two important innovations. First, divorce is explicitly modeled as a steady-state phenomenon. When hit by a negative "bliss shock," spouses may either separate (then they each get back to the marriage market) or renegotiate the way resources and duties are allocated within the household.[26] Second, the authors allow for behavior to be influenced

[26]The model assumes no commitment: Any (monetary or nonmonetary) shock affecting the household triggers renegotiation.

by family values, which are heterogeneous among individuals. In particular, they present counterfactual simulations of an economy where all individuals are liberal; they show that the marriage rate would decline, and married women would increase labor market participation very substantially.

Ciscato (2019) estimates a related model based on US data. In his framework, spouses are able to insure each other against wage shocks. However, in the absence of full commitment, both wage and love shocks can trigger divorce; then, agents are free to look for a new spouse, but their marriage prospects deteriorate as they get older. The model, estimated for two separate periods, the 1970s and the 2000s, replicates the cross-sectional marriage patterns (who gets married and with whom), the longitudinal marriage patterns (the odds of getting married and divorced at various stages of the life cycle), and the female labor supply patterns. Up to a third of the decline in the share of married adults between the 1970s and the 2000s appears to be due to changes in the wage distribution.

Finally, a recent contribution by Shephard (2019) departs from the previous body of work in several respects. He considers an overlapping generation model in which individuals, in each period, can meet at most one potential spouse from all marriageable cohorts; if marriage occurs, the quality of the match evolves stochastically. Importantly, Shephard assumes limited commitment à la Mazzocco (2004), a feature that allows for a significant amount of risk sharing within the couple: While agents cannot commit not to divorce, marriage contracts are second-best efficient ex ante, resulting in a renegotiation of Pareto weights only when the participation constraint of one spouse becomes binding. Finally, the presence of several cohorts that may inter-marry allows to analyze topics like the evolution of age at first marriage or of the marital age gap—aspects that are influenced by the economic environment and that may, in turn, affect household behavior. For instance, Shephard (2019) finds that the significant increase in women's relative earnings since the 1980s increased female employment and the age of first marriage for women, while reducing male employment and the marital age gap.

5. CONCLUSIONS

Although the economic analysis of the marriage market was initiated some 50 years ago by Becker's seminal 1973 contribution, it has attracted renewed attention over the last two decades, accompanying a reconsideration of microeconomic models of household behavior that started in the late 1980s. The notion that the intra-household allocation of resources should be influenced by the situation of the market for marriage dates back to Becker (at least), but it was hard to reconcile with a unitary representation of the family as maximizing a single, exogenously given utility. By explicitly recognizing the importance of members' respective power in the household decision process, and more importantly, by providing tools for both formalizing this intuition and identifying the corresponding structure from observed behavior, new approaches—and especially the collective model—have provided the missing link. This has opened the way for a joint formalization of the interactions between family behavior and the market for (re-)marriage. Moreover, the (static) efficiency assumption that lies at the heart of the collective approach is fully compatible with the frameworks typically used by market models (which commonly refer to bargaining and/or equilibrium considerations). All in all, the landscape, almost half a century after Becker's work, is one of general unification of the field around some core principles, as well as a spectacular development of the empirical tools needed to take these concepts to data in a highly rigorous way.

In practice, two main approaches, respectively based on frictionless matching and search, have generated a host of specific models addressing different issues and relying on diverse methodologies. Despite its diversity, however, this burgeoning literature tends to converge on a small number of common, important messages. One is the crucial role of the family in the generation of HC, which lies at the core of some of the most important issues facing modern economies (growth and inequality being obvious examples). Second, family formation and family dissolution must be analyzed as endogenous phenomena that respond (at least in part) to the economic context individuals are facing. It follows that any assessment of the long-term consequences of a policy cannot possibly ignore its impact on these demographic determinants. To take only one example, the short-term effect of a tax reform may be a small change in individual savings or labor supply. In the long run, however, the indirect impact on incentives to marry and to invest in HC (whether one's own or one's children's) may be of major importance; neglecting these general equilibrium consequences may result in dramatic misconceptions.

A common message of most studies is the importance of assortative matching for long-term economic outcomes. Changes in assortativeness have a direct impact on (static) inequality. Much more important, however, are the long-term consequences in terms of investment in HC. In the inequality spiral described by Chiappori et al. (2017c), a sharp rise in the labor market premium to skill leads to more assortative matching on education and higher investment on children's HC by more educated families, generating more inequality and increased marital sorting for the next generation. It is interesting to note that a set of more macro-oriented contributions, based on diverse theoretical frameworks and relying on different empirical methodologies, reach similar conclusions.

Finally, it is important to keep in mind that different economic contexts may give rise to totally divergent evolutions. Most empirical work has concentrated on the United States (and, to a lesser extent, the United Kingdom); yet in other countries, even within the Western world, patterns and trends may be quite dissimilar. For instance, the rise in the college premium (and the consequent surge in assortative matching and inequality) seems to be specific to particular countries and to the recent decades. The approaches described in this survey could be used to analyze different data, potentially generating opposite conclusions. Here as elsewhere, more empirical work is needed. One can only hope that the advances in both theoretical and empirical approaches described in this review will foster these new developments.

DISCLOSURE STATEMENT

The author is not aware of any affiliations, memberships, funding, or financial holdings that might be perceived as affecting the objectivity of this review.

ACKNOWLEDGMENTS

I am grateful to Jean Marc Robin and to Bernard Salanié for many useful comments and suggestions. Errors are mine.

LITERATURE CITED

Abdulkadiroğlu A, Pathak P, Roth A. 2005. The New York City high school match. *Am. Econ. Rev.* 95:364–67

Abdulkadiroğlu A, Sonmez T. 2003. School choice: a mechanism design approach. *Am. Econ. Rev.* 93:729–47

Ahn S-Y. 2017. *Matching across markets: theory and evidence on cross-border marriage*. Work. Pap., Columbia Univ., New York

Aiyagari SR, Greenwood J, Guner N. 2000. On the state of the union. *J. Political Econ.* 109:213–44

Ashraf N, Bau N, Nunn N, Voena A. 2020. Bride price and female education. *J. Political Econ.* 128:591–641

Atakan A. 2006. Assortative matching with explicit search costs. *Econometrica* 74(3):667–80

Azevedo E, Hatfield J. 2015. *Existence of equilibrium in large matching markets with complementarities*. Work. Pap., Wharton Sch., Univ. Pa., Philadelphia

Banerjee A, Duflo E, Ghatak M, Lafortune J. 2013. Marry for what? Caste and mate selection in modern India. *Am. Econ. J. Microecon.* 5:33–72

Becker G. 1973. A theory of marriage: part I. *J. Political Econ.* 81:813–46

Becker G. 1974. A theory of marriage: part II. *J. Political Econ.* 82:S11–26

Bisin A, Tura G. 2019. *Marriage, fertility, and cultural integration of immigrants in Italy*. Work. Pap., New York Univ., New York

Bojilov R, Galichon A. 2015. *Matching in closed-form: equilibrium, identification, and comparative statics*. Work. Pap., Pontif. Cathol. Univ. Chile, Santiago

Bronson MA, Mazzocco M. 2018. *Cohort size and the marriage market: explaining nearly a century of changes in U.S. marriage rates*. Work. Pap., Univ. Calif., Los Angeles

Burtless G. 1999. Effects of growing wage disparities and changing family composition on the U.S. income distribution. *Eur. Econ. Rev.* 43(4–6):853–65

Chiappori PA. 2017. *Matching with Transfers*. Princeton, NJ: Princeton Univ. Press

Chiappori PA, Costa Dias M, Meghir C. 2018a. The marriage market, labor supply and education choice. *J. Political Econ.* 126(S1):S26–72

Chiappori PA, Costa Dias M, Meghir C. 2019a. *Measuring assortativeness*. Work. Pap., Columbia Univ., New York

Chiappori PA, Costa Dias M, Meghir C, Xiao PP. 2019b. *Education, marriage and child development*. Work. Pap., Yale Univ., New Haven, CT

Chiappori PA, Galichon A, Salanié B. 2019c. On human capital and team stability. *J. Hum. Cap.* 13(2):236–59

Chiappori PA, Gugl E. 2020. Transferable utility and demand functions. *Econ. Theory.* In press

Chiappori PA, Iyigun M, Lafortune J, Weiss Y. 2017a. Changing the rules midway: the impact of granting alimony rights on existing and newly formed partnerships. *Econ. J.* 127(604):1874–905

Chiappori PA, Iyigun M, Weiss Y. 2009. Investment in schooling and the marriage market. *Am. Econ. Rev.* 99(5):1689–717

Chiappori PA, McCann R, Nesheim L. 2010. Hedonic price equilibria, stable matching, and optimal transport: equivalence, topology, and uniqueness. *Econ. Theory* 42(2):317–54

Chiappori PA, McCann R, Pass B. 2017b. Multi- to one-dimensional transportation. *Commun. Pure Appl. Math.* 70(12):2405–44

Chiappori PA, Nguyen L, Salanié B. 2019d. *Matching with random components: simulations*. Work. Pap., Columbia Univ., New York

Chiappori PA, Oreffice S, Quintana Domeque C. 2012. Fatter attraction: anthropometric and socioeconomic matching on the marriage market. *J. Political Econ.* 120(4):659–95

Chiappori PA, Oreffice S, Quintana Domeque C. 2018b. Bidimensional matching with heterogeneous preferences in the marriage market. *J. Eur. Econ. Assoc.* 16(1):161–98

Chiappori PA, Oreffice S, Quintana Domeque C. 2020. Fatter attraction: anthropometric and socioeconomic matching on the marriage market: a corrigendum. *J. Political Econ*. In press

Chiappori PA, Salanié B. 2016. The econometrics of matching models. *J. Econ. Lit.* 54(3):832–61

Chiappori PA, Salanié B, Weiss Y. 2017c. Partner choice, investment in children, and the marital college premium. *Am. Econ. Rev.* 107(8):2109–67

Chiappori PA, Weiss Y. 2006. Divorce, remarriage and welfare: a general equilibrium approach. *J. Eur. Econ. Assoc.* 4(4):415–26

Chiappori PA, Weiss Y. 2007. Divorce, remarriage and child support. *J. Labor Econ.* 25(1):37–74

Choo E, Siow A. 2006. Who marries whom and why. *J. Political Econ.* 114:175–201

Ciscato E. 2019. *Matching models with and without frictions: applications to the economics of the family*. PhD Thesis, Paris Inst. Political Stud., Paris

Ciscato E, Galichon A, Goussé M. 2020. Like attract like? A structural comparison of homogamy across same-sex and different-sex households. *J. Political Econ.* 128(2):740–78

Ciscato E, Weber S. 2019. *The role of evolving marital preferences in growing income inequality*. Work. Pap., Inst. Étud. Politiques Paris, Paris

Cole HL, Mailath GJ, Postlewaite A. 2001. Efficient non-contractible investments in large economies. *J. Econ. Theory* 101(2):333–73

Dagsvik JK. 2000. Aggregation in matching markets. *Int. Econ. Rev.* 41:27–57

Diamond PA. 1982. Wage determination and efficiency in search equilibrium. *Rev. Econ. Stud.* 49:217–27

Diamond PA. 1984. *A Search-Equilibrium Approach to the Micro Foundations of Macroeconomics*. Cambridge, MA: MIT Press

Doepke M, Tertilt M. 2016. Families in macroeoconomics. In *Handbook of Macroeconomics*, Vol. 2, ed. JB Taylor, H Uhlig, pp. 1789–891. Amsterdam: Elsevier

Dokko J, Li G, Hayes J. 2015. *Credit scores and committed relationships*. Work. Pap., Finance Econ. Discuss. Ser. 2015-081, Board Gov. Fed. Reserv. Syst., Washington, DC. **http://dx.doi.org/10.17016/FEDS.2015.081**

Dupuy A, Galichon A. 2014. Personality traits and the marriage market. *J. Political Econ.* 122(6):1271–319

Eika L, Mogstad M, Zafar B. 2019. Educational assortative mating and household income inequality. *J. Political Econ.* 127(6):2795–835

Fernández R. 2002. Education, segregation and marital sorting: theory and an application to the UK. *Eur. Econ. Rev*. 46:993–1022

Fernández R, Guner N, Knowles J. 2005. Love and money: a theoretical and empirical analysis of household sorting and inequality. *Q. J. Econ.* 120:273–344

Fernández R, Rogerson R. 2001. Sorting and long-run inequality. *Q. J. Econ*. 116:1305–41

Fiorio C, Verzillo S. 2018. *Watching in your partner's pocket before saying "Yes!": income assortative mating and inequality*. Work. Pap., Eur. Comm. Joint Res. Cent., Brussels, Belg.

Fox J. 2008. *Estimating matching games with transfers*. NBER Work. Pap. 14382

Fox J. 2010. Identification in matching games. *Quant. Econ.* 1:203–54

Fox J, Bajari P. 2013. Measuring the efficiency of an FCC spectrum auction. *Am. Econ. J. Microecon.* 5:100–46

Frémeaux N, Lefranc A. 2017. *Assortative mating and earnings inequality in France*. IZA Discuss. Pap. 11084, Inst. Labor Econ., Bonn, Ger.

Gale D, Shapley LS. 1962. College admissions and the stability of marriage. *Am. Math. Mon.* 69(1):9–15

Galichon A. 2015. *Optimal Transport Methods in Economics*. Princeton, NJ: Princeton Univ. Press

Galichon A, Kominers SD, Weber S. 2019. Costly concessions: an empirical framework for matching with imperfectly transferable utility. *J. Political Econ.* 127(6):2875–925

Galichon A, Salanié B. 2015. *Cupid's invisible hand: social surplus and identification in matching models*. Work. Pap., New York Univ., New York

Galichon A, Salanié B. 2017. The econometrics and some properties of separable matching models. *Am. Econ. Rev.* 107(5):251–55

Gayle GL, Shephard A. 2019. Optimal taxation, marriage, home production, and family labor supply. *Econometrica* 87(1):291–326

Gersbach H, Haller H. 2017. *Groups and Markets: General Equilibrium with Multi-Member Households*. New York: Springer

Goussé M, Jacquemet N, Robin J-M. 2017. Marriage, labor supply, and home production. *Econometrica* 85(6):1873–919

Greenwood J, Guner N. 2004. *Marriage and divorce since World War II: analyzing the role of technological progress on the formation of households*. NBER Work. Pap. w10772

Greenwood J, Guner N, Knowles J. 2003. More on marriage, fertility, and the distribution of income. *Int. Econ. Rev.* 44:827–62

Greenwood J, Guner N, Kocharkov G, Santos C. 2014. Marry your like: assortative mating and income inequality. *Am. Econ. Rev.* 104(5):348–53

Greenwood J, Guner N, Kocharkov G, Santos C. 2016. Technology and the changing family: a unified model of marriage, divorce, educational attainment, and married female labor-force participation. *Am. Econ. J. Macroecon.* 8(1):1–41

Greinecker M, Kah C. 2019. *Pairwise stable matching in large economies*. Work. Pap., Univ. Gratz, Gratz, Austria

Grossbard S. 1993. *On the Economics of Marriage: A Theory of Marriage, Labor, and Divorce*. Boulder, CO: Westview Press

Harris M, Cronin C. 2014. *The effects of prospective mate quality on investments in healthy body weight among single women*. Work. Pap., Univ. Tenn., Knoxville

Hitsch G, Hortacsu A, Ariely D. 2010. Matching and sorting in online dating. *Am. Econ. Rev.* 100:130–63

Iyigun MF, Lafortune J. 2016. *Why wait? A century of education, marriage timing and gender roles*. IZA Discuss. Pap. 9671, Inst. Labor Econ., Bonn, Ger.

Iyigun M, Walsh R. 2007. Building the family nest: a collective household model with competing pre-marital investments and spousal matching. *Rev. Econ. Stud.* 74:507–35

Lauermann S, Nöldeke G. 2015. Existence of steady-state equilibria in matching models with search frictions. *Econ. Lett.* 131:1–4

Legros P, Newman AF. 2007. Beauty is a beast, frog is a prince: assortative matching with nontransferabilities. *Econometrica* 75(4):1073–102

Liu Q. 2018. *Rational expectations, stable beliefs, and stable matching*. Work. Pap., Columbia Univ., New York

Liu Q, Mailath GJ, Postlewaite A, Samuelson L. 2014. Stable matching with incomplete information. *Econometrica* 82:541–87

Low C. 2014. *Essays in gender economics*. PhD Thesis, Columbia Univ., New York

Low C. 2019. *A "reproductive capital" model of marriage market matching*. Work. Pap., Wharton Sch., Univ. Pa., Philadelphia

Lundberg S, Pollak R. 2009. *Marriage market equilibrium and bargaining*. Presidential address delivered at the Society of Labor Economists 14th Annual Meeting, Boston, MA

Mare RD. 1991. Five decades of educational assortative mating. *Am. Sociol. Rev.* 56(1):15–32

Mazzocco M. 2004. Saving, risk sharing, and preferences for risk. *Am. Econ. Rev.* 94:1169–82

Mortensen D. 1988. Finding a partner for life or otherwise. *Am. J. Sociol.* 94:S215–40

Nocera N. 2019. *Marital segregation and the labor market outcomes of immigrants*. Work. Pap., Univ. Chicago, Chicago

Nöldeke G, Samuelson L. 2015. Investment and competitive matching. *Econometrica* 83(3):835–96

Persson P. 2020. Social insurance and the marriage market. *J. Political Econ.* 128(1):252–300

Pissarides CA. 1990. *Equilibrium Unemployment Theory*. Oxford, UK: Blackwell

Pollak RA. 2003. Gary Becker's contributions to family and household economics. *Rev. Econ. Househ.* 1(1–2):111–41

Reynoso A. 2019. *Polygamy, co-wives' complementarities, and intra-household inequality*. Work. Pap., Univ. Mich., Ann Arbor

Roth AE. 1984. The evolution of the labor market for medical interns and residents: a case study in game theory. *J. Political Econ.* 92:991–1016

Roth AE, Sonmez T, Utku Unver M. 2004. Kidney exchange. *Q. J. Econ.* 119(2):457–88

Roth AE, Sonmez T, Utku Unver M. 2005. Pairwise kidney exchange. *J. Econ. Theory* 125(2):151–88

Roth AE, Sotomayor M. 1992. *Two-Sided Matching: A Study in Game-Theoretic Modeling and Analysis*. Cambridge, UK: Cambridge Univ. Press

Schulhofer-Wohl S. 2006. Negative assortative matching of risk-averse agents with transferable expected utility. *Econ. Lett.* 92(3):383–88

Schwartz CR. 2010. Earnings inequality and the changing association between spouses' earnings. *Am. J. Sociol.* 115(5):1524–57

Schwartz CR, Mare RD. 2005. Trends in educational assortative marriage from 1940 to 2003. *Demography* 42(4):621–46

Shapley LS, Shubik M. 1971. The assignment game I: the core. *Int. J. Game Theory* 1(1):111–30

Shephard A. 2019. *Marriage market dynamics, gender, and the age gap*. Work. Pap., Univ. Pa., Philadelphia

Shimer R, Smith L. 2000. Assortative matching and search. *Econometrica* 68:343–69

Smith L. 2011. Frictional matching models. *Annu. Rev. Econ.* 3:319–38

Tinbergen J. 1956. On the theory of income distribution. *Weltwirtschaftliches Arch*. 77(2):155–73

Weiss Y, Willis RJ. 1993. Transfers among divorced couples: evidence and interpretation. *J. Labor Econ.* 11:629–79

Weiss Y, Willis RJ. 1997. Match quality, new information, and marital dissolution. *J. Labor Econ.* 15:S293–329