Multi-sided Matching with Transfers: The Economics of Labor and Love

Pauline Corblet* Alejandro Robinson-Cortés†

December 22, 2025

Preliminary and incomplete
(Please do not distribute)

Abstract

This paper develops an empirical framework to study the joint equilibrium of interrelated matching markets. In our model, agents from an arbitrary number of sides (e.g., men, women, firms, schools, locations) form coalitions and split a transferable surplus. The central feature of the model is cross-market surplus complementarities: the surplus created in one match can depend on partners and matches in other markets. Our methodological contribution is to extend semiparametric identification results from two-sided markets with transfers to a multi-sided environment, providing conditions for identification of the match surplus from equilibrium match patterns and observed characteristics. We apply the model to study the coevolution of the marriage and labor markets in the United Kingdom during the first two decades of the 20th century. Preliminary results suggest that employment and partnership status are mutually reinforcing in match surplus, and counterfactual simulations quantify the role that the growth in women's education has played in the interplay between the two markets.

Keywords: matching markets, marriage, labor market, semiparametric identification

*Department of Economics, NYU Abu Dhabi, pauline.corblet@nyu.edu
†Department of Economics, University of Exeter, a.robinson-cortes@exeter.ac.uk

1 Introduction

The study of two-sided matching markets is prevalent in economics: men and women in the marriage market, workers and firms in the labor market, schools and students in school admission systems, houses and residents in the housing market, and so on. Anecdotal evidence suggests that matching markets may be interrelated. Do people consider their marriage prospects when choosing a career (and vice versa)? Do people decide where to live as a function of the school they wish their kids to attend to? Indeed, in the fields of labor and family economics, there is growing evidence that marriage and labor supply decisions are intimately interrelated, particularly for women (Goldin, 2021). Nonetheless, by and large, matching markets tend to be studied in isolation in the literature.

In this paper, we develop a general framework that allows to model the joint equilibrium of interrelated matching markets. We consider an economy with an arbitrary number of sides (men, women, jobs, schools, locations, etc). Agents from every side match to each other to form coalitions and split their joint surplus. Surplus complementarities across sides are the heart of our analysis. What role does your spouse, and their job, play on the surplus that you generate at your own job? Does making a good match on the labor market increase or decrease the surplus of a good match in the marriage market?

Our framework falls into the class of empirical matching models with transferable utility (TU) pioneered by Choo and Siow (2006). Our main methodological contribution to this literature is twofold: we show the existence of a stable matching in this larger class of models, and generalize the semiparametric identification results of two-sided models (see Galichon and Salanié (2022)). The framework we develop encompasses multiple environments studied in the matching literature, including one-sided matching models (roommate problems), two-sided matching models (marriage markets), and many-to-one matching models (college admissions).

In ongoing work, we use our model to study the joint evolution of the marriage and the labor markets in the United Kingdom in recent decades. The backdrop of our analysis is the steady increase of higher education in the population, particularly among women, as well as women's increased rates of participation in the labor market. In tandem, there has been a downward trend of marriage and cohabitation in recent decades, albeit one which has slowed down in recent years in the UK. The focus of the empirical exercise is on the measurement of the match-surplus complementarities underlying these trends. Preliminary findings point to the importance of employment in the surplus of marriage and vice versa. We conduct a counterfactual exercise to shed light on the role that an accelerated increase in women's education has played in the coevolution of both markets.

The paper is organized as follows. In what remains of the introduction, we review the related literature. We provide motivating evidence in Section 2, and spell out the model and our methodological contribution in Section 3. Section 4 presents the results of our empirical application. We conclude in Section 5 with next steps and future research. Proofs are omitted in this preliminary draft.

Related Literature. This paper contributes to a long tradition that studies matching markets as equilibrium objects and seeks to recover match surplus from observed sorting patterns. Classic analyses emphasize that who matches with whom reflects heterogeneity and complementarities in match surplus. Becker (1973) remains a foundational reference for studying assortative matching and the economic forces shaping marriage patterns. On the theory side, matching with transfers is naturally formalized through stable allocations and competitive equilibria. In the assignment game of Shapley and Shubik (1971), core allocations coincide with competitive outcomes in a one-to-one TU environment, and large-market limits such as the nonatomic assignment model of Gretsky, Ostroy, and Zame (1992) provide a useful benchmark when types represent masses of agents, as is typical in empirical work.

Empirically, marriage-market analysis has been shaped by the semiparametric TU framework pioneered by Choo and Siow (2006), which links equilibrium match frequencies to systematic components of match surplus while allowing for rich unobserved heterogeneity. This approach has supported applications that study which traits drive sorting and how they enter surplus, including personality and non-cognitive characteristics (Dupuy and Galichon, 2014), partner choice and parental investments in children's human capital (Chiappori, Salanié, and Weiss, 2017), comparisons of sorting in same-sex and different-sex households (Ciscato, Galichon, and Goussé, 2020), and recent work broadening the scope to gender and sexual orientation (Ciscato and Goussé, 2024). A key appeal of this literature is that equilibrium restrictions deliver discipline for counterfactuals without requiring a fully parametric structure.

Methodologically, our paper is closest to the semiparametric identification approach for TU matching models developed by Galichon and Salanié (2022). More broadly, it fits within a growing literature that uses convex analysis and optimal transport to study identification in matching, leveraging duality to characterize equilibrium allocations and welfare objects (see Villani, 2009; Galichon, 2016). For example, Chiappori, McCann, and Nesheim (2010) relate hedonic pricing equilibria to TU stable matching through optimal transport; Galichon, Kominers, and Weber (2019) extend the Choo–Siow framework to settings with imperfectly transferable utility; related connections to demand inversion and dynamic discrete choice

appear in Bonnet, Galichon, Hsieh, O'Hara, and Shum (2022), and Chiong, Galichon, and Shum (2016).

The substantive motivation for our application is the growing evidence that marriage-market outcomes and labor-market outcomes evolve jointly through household decisions, home production, and time allocation. Greenwood, Guner, Kocharkov, and Santos (2016) develop a unified quantitative equilibrium model that links marriage, divorce, education, and married women's labor-force participation, and use it to account for their joint evolution over time. Calvo, Lindenlaub, and Reynoso (2024) develop and estimate a model in which household labor supply links sorting in marriage and labor markets through home-production complementarities. Related evidence documents shifting assortative matching and inequality in the UK (Chiappori, Costa Dias, Crossman, and Meghir, 2020) and changing marital sorting in the United States (Chiappori, Costa Dias, Meghir, and Zhang, 2025), alongside broader demographic trends in economic homogamy (Gonalons-Pons and Schwartz, 2017). These developments coincide with long-run changes in women's educational attainment and career opportunities (Goldin, 2021), reflected in cohort evidence on gender gaps in high-powered careers (Bertrand, Goldin, and Katz, 2010) and in broader syntheses of evolving family roles (Lafortune, Salisbury, and Siow, 2024).

2 Empirical Motivation

Marriage and labor participation patterns have evolved significantly over the past three decades. We estimate the model using data from the UK Labour Force Survey (UKLFS), along job-vacancy data from the Office of National Statistics (ONS). We focus on the period 2001-2019. The UKLFS data is at the household level and representative of all UK households. The data include information on gender, age, education, marital status, and labor force status. Our model may be estimated with standard data, typically available in national labor surveys, such as the UKLFS. The main requirement is to observe joint characteristics and employment status of both partners in a couple. In this section, we present motivating evidence of the coevolution of the marriage and labor markets in the UK during the past two decades.

Figure 1 shows trends in educational attainment and labor market participation in the UK during the past two decades. The left panel reports the share of the population with a post-secondary education (what we dub "high skill" or "highly educated") by gender. The share of high-skill individuals in the population grew steadily between 2000 and 2020. Remarkably, the share of highly educated women has grown at a faster rate than that of men, with the share of highly-educated women overtaking that of men around 2007.

The right panel of Figure 1 reports the share of the population who is active in the labor market by gender, where being active means employed or actively seeking for a job in the past 4 weeks. The share of participation for men has increased if only slightly, staying at or above 90%. By contrast, the participation rate of women has steadily increased, from less than 75% in 2000 to over 80% in 2020.



Figure 1: Changes in Educational Attainment (left)
and Labor Market Participation (right) by gender

Figure 2 shows the changes in the rate at which people form couples, through marriage or cohabitation (what we dub "coupling rate"). Prior to 2010, the share of individuals in a couple steadily decreased (or remained flat) across all skills and ages. However, this trend stopped and, to some extent, reversed during the 2010s. This occurred among both low- and high-skill individuals (left panel of Figure 2). By contrast, the right panel of Figure 2 shows that the change has been primarily driven by middle-age individuals (35-44 years old), whose coupling rates steadily increased after 2010.

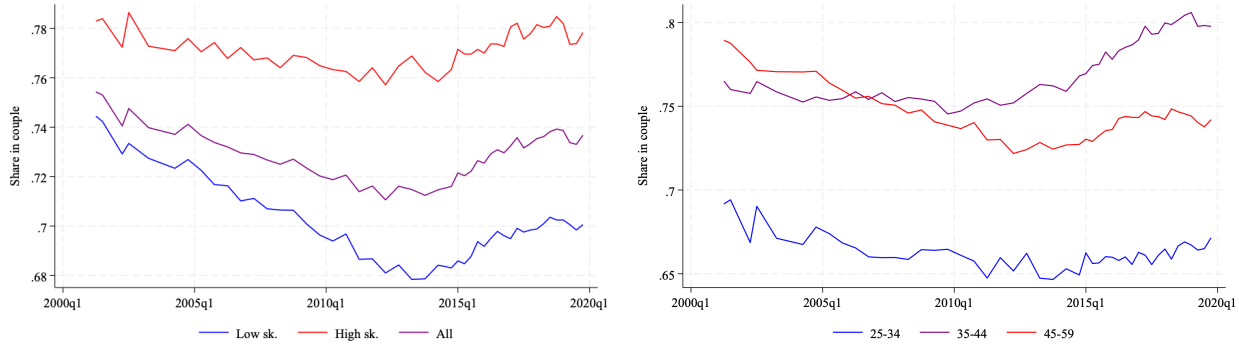


Figure 2: Share in a couple (marriage and/or cohabitation) by skill (left) and age (right)

Figure 3 reports the trends in sorting among couples by education. In terms of shares directly obtained from the data, the degree of assortativeness decreased during these two

decades. While the share of couples among high-skilled individuals steadily increased, the share of couples among low-skilled individuals decreased even further. As a result, the share of couples in which the woman and the man have different degrees of education has increased. Importantly, these trends reflect both changes in the underlying population of women and men, as well as any changes that may have occurred in terms of preferences during that period.

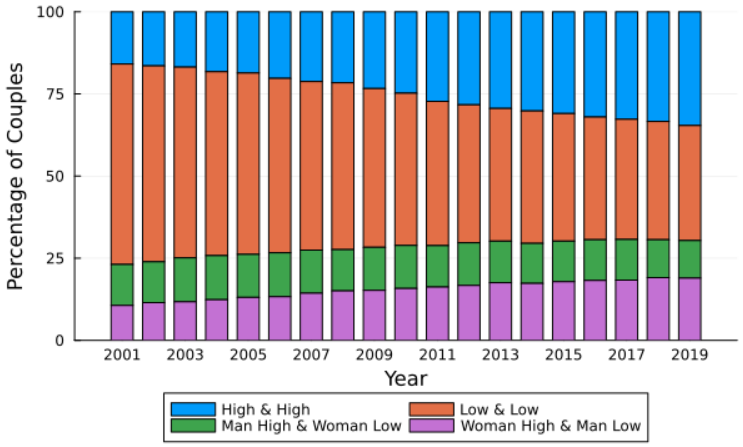


Figure 3: Couples by skills

3 Model

Our model is an extension of the two-sided empirical TU-matching model of Choo and Siow (2006) to multiple sides. We follow the generalization of Galichon and Salanié (2022).

Population model. Our empirical approach assumes that the data is sampled from a large population. Consider a population split into K sides, where $K \geqslant 1$ is finite. Each side contains a continuum of agents. Let A^k denote the set of agents in side $k \in [K]$.[1] Every agent belongs to one and only one side, so $\{A^k\}_{k\in[K]}$ is a partition of A, the set of all agents in the population. Agents from side k are referred to as k-agents. A match (or coalition) may include agents from any of the K sides, and the number of agents in a match may vary. Specifically, a match may include up to N^k agents from side k, where $N^k \geqslant 1$ for all k. For example, consider a market where $K = 3$, and the three sides are women, men and jobs. A match may be composed of a man, a woman, and two jobs; or two women and one job; or a single unemployed man, etc.

[1] For $K \in \mathbb{N}$, $[K] \equiv \{1, \ldots, K\}$.

An agent $a \in A$ is uniquely identified by their side $k_a \in [K]$, their observable type $x_a \in \mathcal{X}^k \subset \mathbb{R}^{\dim(x^k)}$, and their unobservable type $\varepsilon_a \in \mathbb{R}^{\dim(\varepsilon)}$. While ε_a is unobservable in the data, it is assumed to be observable to all the agents in the population. As agents are uniquely identified by their types, a match corresponds to a finite collection of types. Formally, it may be represented by a multi-set

$$G = [(k_1, x_1, \varepsilon_1), \ldots, (k_m, x_m, \varepsilon_m)], \tag{1}$$

where m denotes the number of agents in the match. The number of agents in a match m is determined in equilibrium and, in principle, can take any value from 1 to $\sum_k N^k$. Let $\mathcal{G}$ denote the set containing all the finite collections of types that contain at most N^k agents of every side k. That is, $\mathcal{G}$ is the set of all possible matches G.

A match $G \in \mathcal{G}$ generates surplus $\tilde{\Phi}(G) \in \mathbb{R}$, which is shared among the members of the match. A matching specifies a partition of the set of agents into matches. The final payoffs received by agents are described by a payoff function $u : A \to \mathbb{R}$. Given a matching, a payoff function u is feasible if the sum of the payoffs of the agents in every match is less than or equal than their match surplus. It satisfies no blocking if

$$\sum_{a \in G} u(a) \geqslant \tilde{\Phi}(G) \quad \text{for all } G \in \mathcal{G}.^{2} \tag{2}$$

The above no blocking condition has the standard interpretation: if there exists $G \in \mathcal{G}$ such that $\sum_{a \in G} u(a) < \tilde{\Phi}(G)$, then the members of G would block the allocation, as they would be better off matching to each other and splitting their match surplus. Under a continuity assumption on $\tilde{\Phi}(G)$, which is satisfied in our setting, it may be shown that a stable matching exists along with an associated payoff function that satisfies no blocking. Furthermore, feasibility implies that (2) is satisfied with equality for every match formed in equilibrium. In what follows, we let u denote the payoff function at equilibrium.

While an equilibrium is guaranteed to exist, it is not directly observable in the data, as it specifies who matches with whom on the basis of unobservables. In the data we only observe an aggregate matching that describes the distribution of observable types across the population matching. Alternatively, we may observe a final sample of matches, which approximates the aggregate matching up to sample variation. For the sake of the identification argument, we assume we have infinite data, i.e., that we observe the aggregate matching without sample variation.

[2] Abusing notation, we write $\sum_{a \in G} u(a)$ in (2) instead of an integral as G contains a finite number of elements (i.e., it is equal to an integral over a finite support).

Observable types. We assume that the space of observable characteristics is finite and may differ across sides.[3] Let $X^k = \{x_i^k\}_{i \in [I^k]}$ denote the space of observable characteristics for agents in side k. That is, the observable type of a k-agent can take any of $I^k \in \mathbb{N}$ values. We denote the set of agents in side k who are of type x_i^k by A_i^k, so that $\{A_i^k\}_i$ is a partition of A^k. From an empirical perspective, agents in A_i^k are observationally equivalent to each other as they only differ by their unobservable types. Since we index observable types by i, we refer to an agent's full type as (k, i, ε), and let i_a denote the index of the agent's observable type, i.e., $i_a = i$ iff $x_a = x_i^k$.

Let $M^k \in \mathbb{R}_{++}$ denote the total mass of k-agents; we normalize $M^1 = 1$. The mass of k-agents who are of type x_i^k is denoted by $n_i^k \in \mathbb{R}_{++}$. Let $n = (n_i^k)_{(i,k)}$ contain the distribution of types across all sides. By definition,

$$\sum_i n_i^k = M^k \quad \forall k. \tag{3}$$

We denote the distribution of observable types within a match by $\theta = (\theta_i^k)_{(i,k)}$, where $\theta_i^k \in \mathbb{N}_0$ denotes the number of k-agents in the match who have type x_i^k, and $\mathbb{N}_0 \equiv \mathbb{N} \bigcup \{0\}$. We refer to θ as a match-type. Let $\tilde{\Theta}$ denote the set of all possible match-types, given by:

$$\tilde{\Theta} = \left\{ \left(\theta_i^k\right)_{(i,k)} : \forall (k,i), \;\; \theta_i^k \in \mathbb{N}_0 \;\; \& \;\; \sum_i \theta_i^k \leqslant N^k \right\}. \tag{4}$$

Note that $\tilde{\Theta}$ is a finite set. Nevertheless, it may include match-types that we may wish to exclude from the analysis. For one, it includes an empty match-type, as there exists $\theta \in \tilde{\Theta}$ such that $\theta_i^k = 0$ for all (k, i). Denote this empty match-type by $\theta^\varnothing$. We may also wish to constrain the set of match-types to incorporate aspects of certain applications. For example, consider that side k corresponds to the location of a match, and that every match must be located in one and and only one location. Then, we may want to only include match-types θ such that $\sum_i \theta_i^k = 1$ for some k. As the next definition makes clear, we allow for general constraints to the set match-types as long as they satisfy minimum requirements.[4]

[3] Our results do not rely on having a finite observable type space. One can extend our setting to accommodate continuous covariates, along the lines of Dupuy and Galichon (2014).

[4] Consider an additional example. Let $\theta = (\theta^w, \theta^m, \theta^z)$, where θ^w indicates the type of the woman in a heterosexual household, θ^m the type of the man, and θ^z the type of the job(s) they hold (i.e., $N^w = N^m = 1 < 2 = N^z$). Then, we may want to allow a household to hold two jobs only if it consists of a married couple. That is, we may want to rule out matches with types θ in which a single man or woman holds two jobs, i.e., such that $\min\{\sum_{i \in [I^w]} \theta_i^w, \sum_{i \in [I^m]} \theta_i^m\} = 0$ and $\sum_{i \in [I^z]} \theta_i^z = 2$, while allowing types θ in which a single man or woman holds one job, i.e., with $\min\{\sum_{i \in [I^w]} \theta_i^w, \sum_{i \in [I^m]} \theta_i^m\} = 0$ and $\sum_{i \in [I^z]} \theta_i^z = 1$.

Definition 1 (Set of match-types)**.** A set of match-types Θ is a subset of the set of all possible match-types, $\Theta \subset \tilde{\Theta}$, such that:

1. $\theta^{\varnothing} \notin \Theta$;

2. for every (k,i), there exists $\theta \in \Theta$ such that $\theta_i^k > 0$.

The second requirement guarantees that every type within a side is represented in the set of match-types. Let $\Theta_k^i \subset \Theta$ denote the subset of match-types in which type x_i^k is represented, i.e., for which $\theta_i^k > 0$. The second requirement guarantees that Θ_i^k is nonempty for every (k,i). Without it, the set of feasible matchings (defined below) could be empty. A simple way to satisfy this requirement is to allow every type within a side to remain single, as it is usually the case in standard two-sided matching models.[5]

An (aggregate) matching is a vector $\mu \in \mathbb{R}_+^{|\Theta|}$, where $\mu(\theta)$ stands for the mass of matches of type $\theta \in \Theta$ in the population. A matching μ is feasible if

$$\sum_{\theta \in \Theta} \theta_i^k \mu(\theta) = n_i^k \quad \forall (k,i). \tag{5}$$

Equation (5) is stated as an equality since all agents are accounted for in μ, even those who remain single (and have their own match-type).

Match surplus and utility maximization. Denote the composition of types within match $G \in \mathcal{G}$ by $\theta_G \in \Theta$. We specify the surplus function as follows:

$$\tilde{\Phi}(G) = \Phi(\theta_G) + \sum_{a \in G} \varepsilon_{a\theta_G}. \tag{6}$$

That is, $\Phi(\theta)$ captures the portion of the surplus explained by the observable types of the agents in a match, and $\varepsilon_{a\theta}$ represents the unobservable surplus contribution of agent a to a match with type composition θ. Hence, the unobservable type of agent a is given by the collection of that agent's unobservable contributions to all the match-types they belong to, i.e., $\varepsilon_a = (\varepsilon_{a\theta})_{\theta \in \Theta_{i_a}^{k_a}}$ for $k = k_a$ and $i = i_a$.

Assumption 1. *$\varepsilon_a = (\varepsilon_{a\theta})_{\theta \in \Theta_{i_a}^{k_a}}$ is a random vector with probability distribution $\mathbf{P}_{i_a}^{k_a}$. It satisfies the following:*

1. *The random variable $\max_{\theta \in \Theta} |\varepsilon_{a\theta}|$ has finite expectation under $\mathbf{P}_{i_a}^{k_a}$.*

2. *$\mathbf{P}_i^k$ has full support and is absolutely continuous with respect to the Lebesgue measure for every (k,i).*

[5] That is, for every (k,i) include a type θ in Θ such that $\theta_i^k = 1$ and $\theta_{i'}^{k'} = 0$ for all other (k',i').

Furthermore ε_a and $\varepsilon_{a'}$ are independent for every pair of distinct agents $a, a' \in A$.

The specification in (6) generalizes a standard assumption in empirical two-sided matching models with transfers, whereby an agent's unobservable surplus contribution depends on the observable attributes of the other agents in a match. As the following lemma shows, this specification is tightly connected with individual utility maximization.

Lemma 1. *Let u be an equilibrium payoff function. There exists a collection of real numbers $U_i^k(\theta)$, for all $\theta \in \Theta_i^k$, and (i,k), such that:*

(i) For all $\theta \in \Theta$, $\Phi(\theta) = \sum_{k,i} \theta_i^k U_i^k(\theta)$

(ii) For all $a \in A$, $u(a) = \max_{\theta \in \Theta_i^k} U_i^k(\theta) + \varepsilon_{a\theta}$

The proof of the lemma is omitted from this extended abstract. It shows that standard arguments from two-sided markets (e.g., Galichon and Salanié (2022)) carry over to the case with multiple sides. Intuitively, $U_i^k(\theta)$ is the systematic portion of the utility of an agent of type (k,i) in a match of type θ. It is an equilibrium object. While the individual equilibrium utility $u(a)$ may differ across agents who are observationally equivalent, it may only do so up to their unobservable types. Part (i) of the lemma shows that that the systematic component of the match surplus, $\Phi(\theta)$, may be decomposed into the sum of the individual systematic utilities. Part (ii) shows that all agents maximize at the equilibrium: the match-type composition of the match they choose to join, θ maximizes the sum of their systematic utility $U_i^k(\theta)$ and their unobservable contribution $\varepsilon_{a\theta}$.[6]

The unobservable surplus contribution $\varepsilon_{a\theta}$ is usually interpreted as an unobservable preference shock. Agents are assumed to have unobservable preferences over the observable attributes of the agents with whom they match. While the model rules out unobservable preferences over unobservable attributes, it allows for selection of agents into a match based on unobservables. Agents' unobservable preferences are over sets of agent types, imposing no structure on how they relate to preferences over matching with individuals. For instance, the model allows for an agent to prefer matching with an agent of type x to one with type y, $\varepsilon_{a(x)} > \varepsilon_{a(y)}$, while preferring to match with two agents of type y than with two agents of type x, $\varepsilon_{a(yy)} > \varepsilon_{a(xx)}$. Likewise, matching with one agent of each type, $\varepsilon_{a(xy)}$, may be superior or inferior to matching with a single agent and/or with two agents of the same type.

[6] In our exposition, we have set the surplus function $\Phi(\theta)$ as the primitive of the empirical model, and derived $U_i^k(\theta)$ as equilibrium objects, which are systematic individual utilities after transfers. An alternative but equivalent way of building the model is by specifying individual preferences (match amenities) as primitives and deriving equilibrium transfers that clear demand and supply, as in the original formulation of Choo and Siow (2006).

Remark 1. *In environments with (perfectly) transferable utility, our framework encompasses the following:*

1. *One-sided matching markets (roommate problems), in which $K = 1$, and $N^1 = 2$;*
2. *Two-sided one-to-one matching markets (marriage markets), in which $K = 2$, and $N^1 = N^2 = 1$;*
3. *Two-sided many-to-one matching markets (e.g., labor markets), in which $K = 2$, and $N^1 > 1 = N^2$.*

Definition 2. A competitive equilibrium is a tuple (μ, U) such that $\mu = (\mu(\theta))_{\theta\in\Theta}$ is a feasible aggregate matching and $U = \left(U_i^k(\theta)\right)_{\theta\in\Theta_i^k,(k,i)}$ is a vector of systematic utilities that satisfy:

$$\frac{\mu(\theta)\theta_i^k}{n_i^k} = \mathbb{P}\left[\theta \in \text{argmax}_{\Theta_i^k} U_i^k(\theta) + \varepsilon_\theta\right] \quad \forall i \in [I],\ k \in [K],\ \theta \in \Theta_i^k \tag{7}$$

The competitive equilibrium states that the share of (k, i) agents willing to enter match-type θ under systematic utility U has to be equal to the mass of (k, i) demanded by the match μ.

Duality. The following proposition shows that the equilibrium systematic utilities minimize the average surplus in the economy subject to no blocking.

Proposition 1. *The equilibrium vector of systematic utilities $U = \left(U_i^k(\theta)\right)_{\theta\in\Theta_i^k,(k,i)}$ is a solution to the following problem:*

$$\begin{aligned} \min_{U\in\mathbb{R}^{\dim(U)}} \quad & \sum_{k,i} n_i^k H_i^k(U) \\ \textit{subject to} \quad & \sum_{k,i} \theta_i^k U_i^k(\theta) \geqslant \Phi(\theta) \quad \forall \theta \in \Theta, \end{aligned} \tag{8}$$

where $H_i^k(U) = \mathbb{E}_{\varepsilon_a\sim\mathbf{P}_{ia}^{ka}}\left[\max_{\theta\in\Theta_i^k} U_i^k(\theta) + \varepsilon_{a\theta}\right]$.

Proposition 1 is a major step towards being able to compute the equilibrium in practice, as it uses our specific structure for types to rewrite the problem of finding equilibrium payoffs as a finite problem. In the next theorem we show that problem (8) is the dual to a social planner problem on aggregate matching.

Let $H^k(U,n) = \sum_{i\in[I^k]} n_i^k H_i^k(U)$ be the aggregated expected indirect payoff for side k. Define the Fenchel-Legendre transforms as

$$H_i^{k,*}(\mu) = \max_U \sum_{\theta\in\Theta_i^k} \mu(\theta)U_i(\theta) - H^k(U),$$
$$H^{k,*}(\mu,n) = \max_U \sum_{i\in[I^k]} \sum_{\theta\in\Theta_i^k} \mu(\theta)\theta_i^k U_i(\theta) - H^k(U,n),$$

Given Assumption 1, $H^{k,*}$ is differentiable and strictly convex.

Theorem 1. *The equilibrium matching μ maximizes social gain over all feasible matchings, i.e. it solves the social planner problem:*

$$\max_{\mu\in\mathbb{R}_+^{\dim\Theta}} \sum_{\theta\in\Theta} \mu(\theta)\Phi(\theta) - \sum_{k\in[K]} H^{k,*}(\mu,n), \tag{9}$$

which is the dual to problem (8). There exists a unique solution to problem (9).

The following theorem, which follows from the first-order condition of the social planner's problem, shows that the vector of payoffs $U = \left(U_i^k(\theta)\right)_{k,i,\theta}$ and the systematic portion of the match surplus $\Phi(\theta)$ is identified from observables.

Theorem 2. *Assume the surplus of singles is zero (outside option). The vectors of payoffs $U = \left(U_i^k(\theta)\right)_{k,i,\theta}$ and match surpluses $\Phi = (\Phi(\theta))_\theta$ are identified by μ, since:*

$$U_i^k(\theta) = \frac{\partial H_i^{k,*}(\mu)}{\partial\mu(\theta)} \quad \& \quad \Phi(\theta) = \sum_{k,i} \theta_i^k \left[\frac{\partial H_i^{k,*}(\mu)}{\partial\mu(\theta)}\right].$$

See that the function $H_i^{k,*}(\mu)$ only depends on quantities observable in the data. Furthermore, as the following corollary shows, our approach recovers the standard Choo-Siow formula under the logit assumption.

Corollary 1. *Assume ε_a is a vector of iid extreme value type-1 random variables. Then,*

$$\Phi(\theta) = \sum_{(k,i)\in\theta} \theta_i^k \log\left(\frac{\mu(\theta)\theta_i^k}{\mu_i^k}\right) \quad \& \quad \Phi_{xy} = \log\frac{\mu_{xy}}{\mu_{x\varnothing}} + \log\frac{\mu_{xy}}{\mu_{\varnothing y}} \quad \text{for } K=2,$$

where μ_i^k is the mass of (k,i) singles.

4 Empirical Application

In this section, we report the results of our empirical exercise using data from the UK (2001–2019). We divide the analysis into three sub-sections. First, we overview the model specification we follow in the empirical analysis. Second, we present the estimation results. Third, we report the results of a counterfactual exercise that investigates the role of the increased share of highly-educated women.

Model specification. We estimate a model with three sides: women, men, and jobs. A coalition may have at most one woman, one man, and two jobs, with the number of jobs being less than or equal to the number of agents (man and/or woman). That is, there are 8 match types in terms of who is a member of the coalition: (i) couples with two jobs, (ii) couples with one job, (iii) couples without a job, (iv) single woman with a job, (v) single man with a job, (vi) single woman without a job, (vii) single man without a job, and (viii) vacant job. The last three match types correspond to the "singles" in our matching model, equivalent to the outside option in a discrete-choice model. The surplus of remaining unmatched is normalized at zero.

Women and men have two observable characteristics, whether they are high-skilled or low-skilled, and whether they are active in the labor market. Only active agents may hold jobs. While participation in the labor market is exogenous, unemployment is an equilibrium outcome. All jobs are equal to each other in terms of observable characteristics. In total, once we account for observable characteristics, we have 49 match types.

Given the matching rates across all the match-types in the specification, $\{\hat{\mu}(\theta)\}_{\theta\in\Theta}$, we estimate the surplus of each match type under an extreme-value type 1 assumption (see Corollary 1):

$$\hat{\Phi}(\theta) = \sum_{(k,i)\in\theta} \theta_i^k \log\left(\frac{\hat{\mu}(\theta)\theta_i^k}{\hat{\mu}_i^k}\right).$$

That is, $\hat{\Phi}(\theta)$ is the estimated surplus for match type θ, where θ encodes the type of the match in terms of who is part of the coalition and their observable characteristics. We estimate the model for every point in time. For presentation purposes, we report trends of average match surpluses across distinct subsets of match types. Standard errors are computed using the bootstrap. Figures report confidence intervals at a 95% confidence level.

Estimation results. Figure 4 reports the average estimated surplus across distinct match types. Three observations stand out. First, the average match surpluses across match type are relatively stable over time. While there are a few variations within each match type,

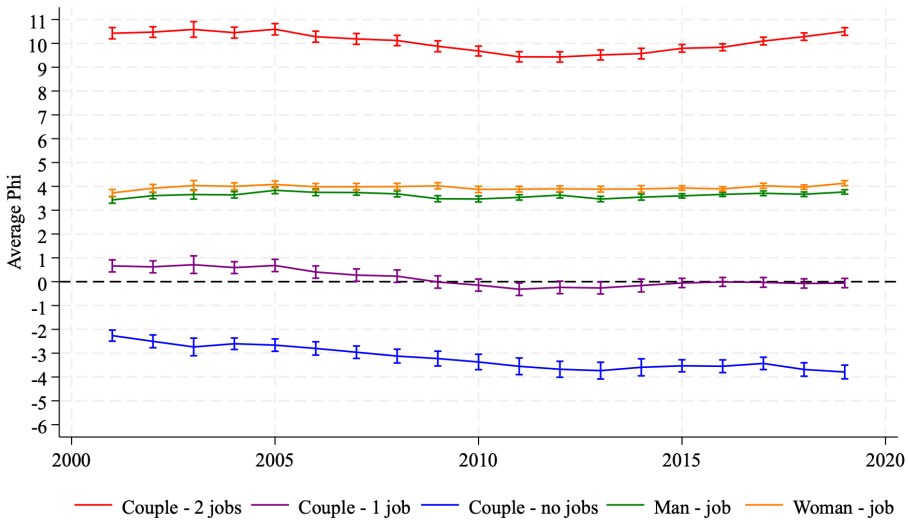


Figure 4: Average Estimated Surplus by match type

the ranking of match types according to the members of the coalition is very stable. The coalitions achieving the highest surplus are those consisting of couples in which both the woman and the man are employed (red series), followed by single women and men who are employed (yellow and green series, respectively). By contrast, couples with a single job (purple series) or no job (blue series) generate the lowest surplus, the latter generating less surplus than remaining single (dashed line at zero).

Second, the value of marriage is highly dependent on the employment status of its members. Notably, the surplus of couples with two jobs (red series) is higher than the sum of the surplus of the surplus of an employed woman (yellow series) and an employed man (green series). This reflects the value of marriage. However, the value of marriage is negative when one of the members in the couple does not have a job, as an employed single woman (yellow series) or man (green series) generates higher surplus than a couple with a single job (purple series).

Third, the value of work depends on the couple status of an individual. Obtaining a job while married to an employed individual (difference between purple and red series) increases the surplus by a greater amount than what it does when married to a spouse without a job (difference between purple and blue). Taken together, these estimates show that the choices agents take when participating in the marriage and the labor markets are interrelated: the gains in economic surplus from both are affected by the other.

Figure 5 disaggregates some of the averages in the previous figure by the education composition of the coalition members. The top panel in Figure 5 reports the average surplus of couples with two jobs differentiating between couples in which both the man and the

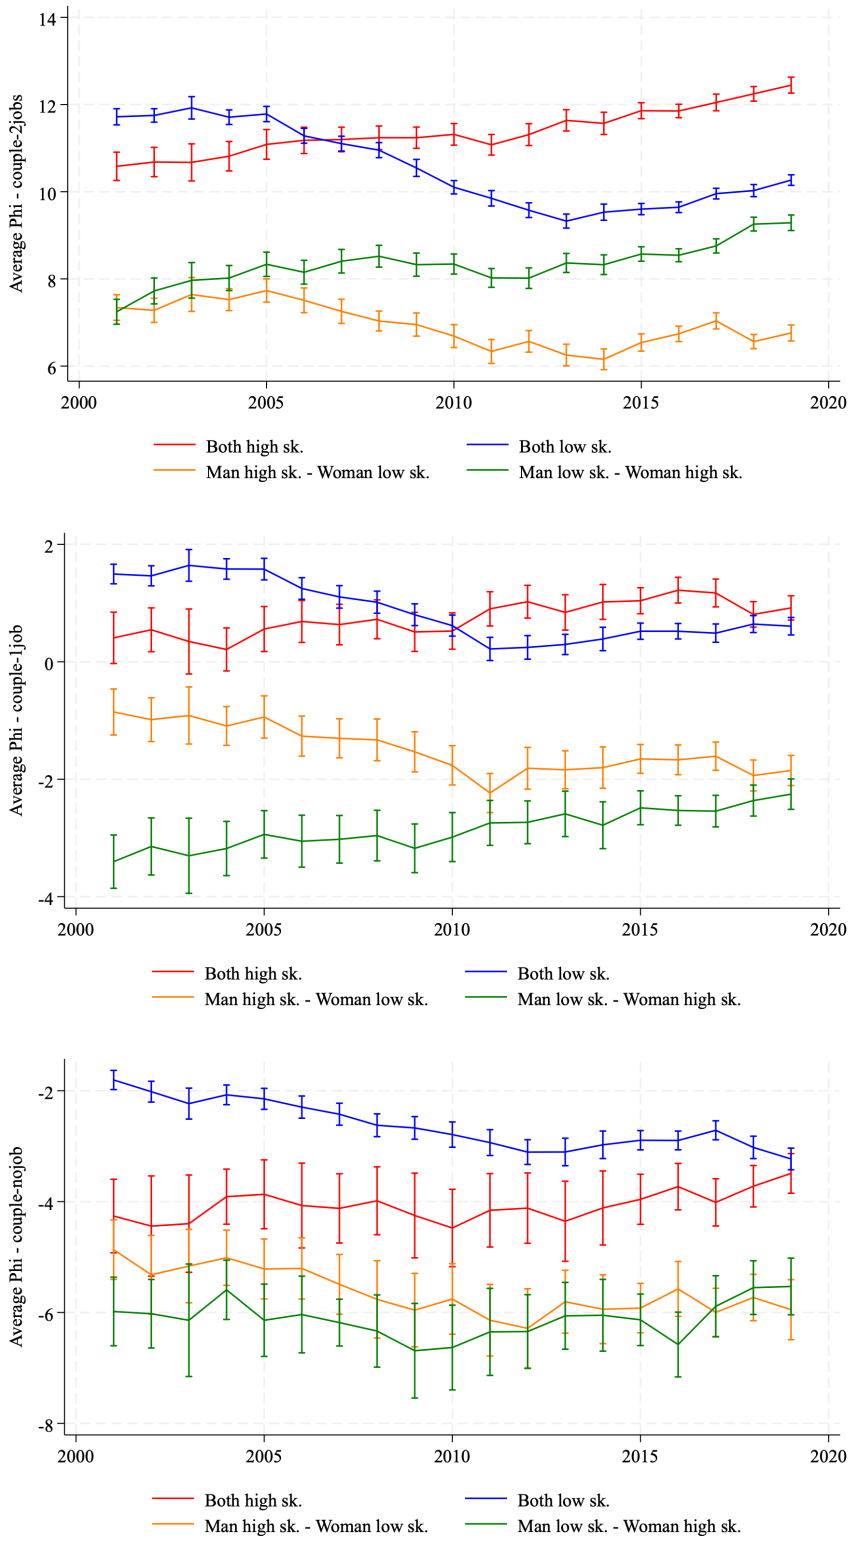


Figure 5: Average Estimated Surplus by education attainment of couples with two jobs (top), one job (middle), and without jobs (bottom)

woman are high-skilled (red series), neither of them is (blue series), only the woman is (green series), and only the man is (yellow series). The middle and bottom panels reproduce the same for couples with one job and for couples without a job, respectively. Taken together, the plots reflect how the value of skills differs depending on the job composition of a couple.[7] Nevertheless, we observe that the surplus function displays positive assortativeness across education: the surplus generated by same-skill couples is higher that that of different-skill couples, even after accounting for job status. This confirms the assortativeness in matching observed on Figure 2 is driven by preferences on education, as is generally osberved in the literature, even after controlling for employment.

Counterfactual analysis. Our counterfactual exercise is focused on the role of women's education in the coevolution of the marriage and labor markets. What would coupling and employment rates have been if the share of high-skilled women had risen more slowly, at the same pace as the share of high-skilled men?

Figure 6 shows the evolution of the coupling rate. As observed, the model is able to replicate the U-shaped pattern observed in the data, whereby the decreasing trend of coupling rate reversed around 2010. In the counterfactual, the coupling rate would be slightly lower. This means that the accelerated pace at which women's education increased during these two decades *increased* the coupling rate. The rate of singlehood would be higher if women's education had not increased so rapidly.

Figure 7 shows that the change in coupling rate is driven mostly by low-skilled women and high-skilled men. First, one can note that that the coupling rate of high-skill women is higher in the counterfactual. This means that the accelerated pace at which women's education increased in these two decades *decreased* the coupling rate of high-skill women. However, this is not the case for low-skill women, who would have had a lower coupling rate if women's education had not increased so rapidly. Because women's education has increased at a higher rate than that of men's, more low-skilled women and high-skill men are in a couple, and this increase overcompensates for the lower coupling rates of high-skill women. Figure 8 shows the effect on the employment rates by gender and skill level.

[7] As skill is an exogenous variable in the model, the estimated surplus does not capture the preferences underlying the decision to acquire education. Rather, it captures the preferences/productivity that drive the decision to couple with or employ a skilled individual.

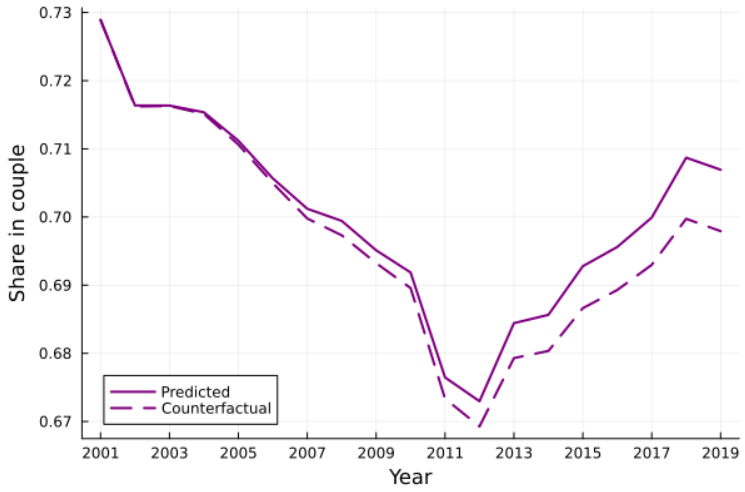


Figure 6: Counterfactual: Coupling Rate

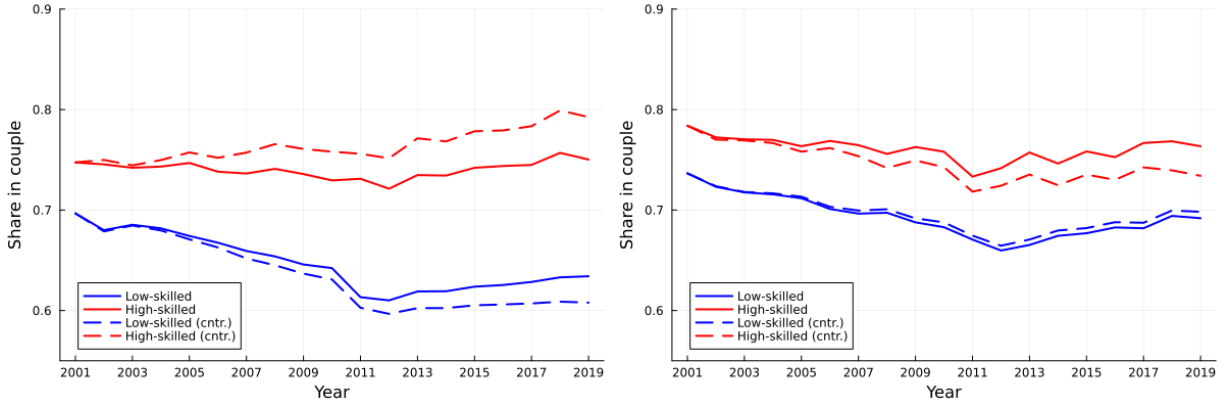


Figure 7: Counterfactual: Coupling Rate by Skill and Gender
(left panel: women; right panel: men)

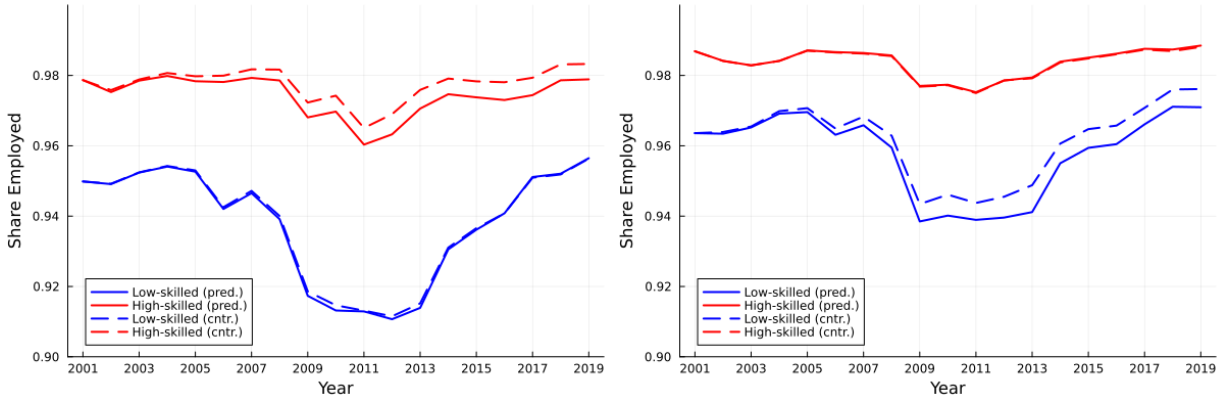


Figure 8: Counterfactual: Employment Rate by Skill and Gender
(left panel: women; right panel: men)

5 Conclusion and Next Steps

This paper develops an empirical framework to model the joint equilibrium of multiple matching markets. We apply our framework to study the coevolution of the marriage and the labor market in the UK during the first two decades of the 20th century. Our main methodological contribution of the paper is to extend identification results from two-sided matching markets to multi-sided ones. Our framework encompasses additional environments studied in the literature, such as one-sided matching markets, and two-sided many-to-one matching markets. Our application highlights the interdependence between the marriage and the labor markets. In ongoing work, we endogenize the labor supply decisions of individuals through a model that incorporates intra-match equilibrium allocations.

References

Becker, G. S. (1973): "A Theory of Marriage: Part I," *Journal of Political Economy*, 81, 813–846. (Cited on page 2)

Bertrand, M., C. Goldin, and L. F. Katz (2010): "Dynamics of the Gender Gap for Young Professionals in the Financial and Corporate Sectors," *American Economic Journal: Applied Economics*, 2, 228–55. (Cited on page 3)

Bonnet, O., A. Galichon, Y.-W. Hsieh, K. O'Hara, and M. Shum (2022): "Yogurts Choose Consumers? Estimation of Random-Utility Models via Two-Sided Matching," *The Review of Economic Studies*, 89, 3085–3114. (Cited on page 3)

Calvo, P., I. Lindenlaub, and A. Reynoso (2024): "Marriage Market and Labour Market Sorting," *The Review of Economic Studies*, 91, 3316–3361. (Cited on page 3)

Chiappori, P.-A., M. Costa Dias, S. Crossman, and C. Meghir (2020): "Changes in Assortative Matching and Inequality in Income: Evidence for the UK," *Fiscal Studies*, 41, 39–63, special Issue: 50th Anniversary of IFS: Part 2. (Cited on page 3)

Chiappori, P.-A., M. Costa Dias, C. Meghir, and H. Zhang (2025): "Changes in Marital Sorting: Theory and Evidence from the United States," *Journal of Political Economy*, 133, 3045–3077. (Cited on page 3)

Chiappori, P.-A., R. J. McCann, and L. P. Nesheim (2010): "Hedonic Price Equilibria, Stable Matching, and Optimal Transport: Equivalence, Topology, and Uniqueness," *Economic Theory*, 42, 317–354. (Cited on page 2)

Chiappori, P.-A., B. Salanié, and Y. Weiss (2017): "Partner Choice, Investment in Children, and the Marital College Premium," *American Economic Review*, 107, 2109–67. (Cited on page 2)

Chiong, K. X., A. Galichon, and M. Shum (2016): "Duality in Dynamic Discrete-Choice Models," *Quantitative Economics*, 7, 83–115. (Cited on page 3)

Choo, E. and A. Siow (2006): "Who Marries Whom and Why," *Journal of Political Economy*, 114, 175–201. (Cited on pages 1, 2, 5, and 9)

Ciscato, E., A. Galichon, and M. Goussé (2020): "Like Attract Like? A Structural Comparison of Homogamy across Same-Sex and Different-Sex Households," . (Cited on page 2)

Ciscato, E. and M. Goussé (2024): "Matching on Gender and Sexual Orientation," IZA Discussion Paper 17420, IZA – Institute of Labor Economics. (Cited on page 2)

Dupuy, A. and A. Galichon (2014): "Personality Traits and the Marriage Market," *Journal of Political Economy*, 122, 1271–1319. (Cited on pages 2 and 7)

Galichon, A. (2016): *Optimal Transport Methods in Economics*, Princeton University Press. (Cited on page 2)

Galichon, A., S. D. Kominers, and S. Weber (2019): "Costly Concessions: An Empirical Framework for Matching with Imperfectly Transferable Utility," *Journal of Political Economy*, 127, 2875–2925. (Cited on page 2)

Galichon, A. and B. Salanié (2022): "Cupid's Invisible Hand: Social Surplus and Identification in Matching Models," *The Review of Economic Studies*, 89, 2600–2629. (Cited on pages 1, 2, 5, and 9)

Goldin, C. (2021): *Career and Family: Women's Century-Long Journey toward Equity*, Princeton University Press. (Cited on pages 1 and 3)

Gonalons-Pons, P. and C. R. Schwartz (2017): "Trends in Economic Homogamy: Changes in Assortative Mating or the Division of Labor in Marriage?" *Demography*, 54, 985–1005. (Cited on page 3)

Greenwood, J., N. Guner, G. Kocharkov, and C. Santos (2016): "Technology and the Changing Family: A Unified Model of Marriage, Divorce, Educational Attainment, and Married Female Labor-Force Participation," *American Economic Journal: Macroeconomics*, 8, 1–41. (Cited on page 3)

Gretsky, N. E., J. M. Ostroy, and W. R. Zame (1992): "The nonatomic assignment model," *Economic Theory*, 2, 1432–0479. (Cited on page 2)

Lafortune, J., L. Salisbury, and A. Siow (2024): "The Changing Identities of American Wives and Mothers," *Journal of Economic Literature*, 62, 1538–88. (Cited on page 3)

Shapley, L. S. and M. Shubik (1971): "The assignment game I: The core," *International Journal of Game Theory*, 1, 111–130. (Cited on page 2)

Villani, C. (2009): *Optimal Transport: Old and New*, vol. 338 of *Grundlehren der mathematischen Wissenschaften*, Berlin, Heidelberg: Springer. (Cited on page 2)