Prediction and Congestion in Two-Sided Markets: Economist versus Machine Matchmakers☆

January 19, 2019

Kuan-Ming Chen[1], Yu-Wei Hsieh[2], Ming-Jen Lin[3]

Abstract

We study the recommender systems on the two-sided platforms. We find that the machine-learning algorithms create *congestion*: they often generate recommendations that are concentrated on a few users. We propose equilibrium machine-learning algorithms: they inherit the predicting power from machine-learning and solve the congestion problem by the market allocation mechanism in economics. We apply our recommenders to an online dating service that contains over 490,000 unique users. Our equilibrium recommenders outperform the baseline machine-learning algorithm in terms of the hit rate and accelerate the matching process by 200% in the counterfactual simulations.

Keywords: Online Dating, Two-Sided Matching, Recommender Systems, Matrix Factorization, Machine Learning, Reciprocal Recommender, Collaborative Filtering

☆**Acknowledgement**: We thank Stéphane Bonhomme, Whitney Newey, Matt Shum and attendees at the 2018 California Econometrics Conference, Taiwan R User Group Workshop (Dec. 2018) for comments. Ming-Jen Lin acknowledges the support from Taiwan Ministry of Science and Technology grant MOST 106-2410-H-002-041-MY2. Our special thank goes to Sheng-Wei Chen for his valuable inputs. We are grateful for Shao-Mou Cheng, Chu-En Ho, and Cheng-Kun Yang for their excellent research assistance.

Disclaimer: We use the confidential, anonymous data from SweetRing. The authors have no material financial relationships with entities related to this research. The algorithms and empirical results expressed herein are those of the authors and do not necessarily reflect SweetRing's algorithms and views.

[1] *Kenneth C. Griffin Department of Economics, University of Chicago*. Email: kmchen@uchicago.edu.

[2] *Department of Economics and Dornsife Institute of New Economic Thinking, University of Southern California*. Email: yuwei.hsieh@usc.edu.

[3] *Department of Economics, National Taiwan University*. Email: mjlin@ntu.edu.tw

1. Introduction

Online two-sided platforms have become increasingly popular over the past decade. Prominent examples include job matching platforms (e.g., LinkedIn), expertise finders, and online dating.[4] Modern e-Business often deploys a recommender to help consumers find products. In particular, websites allow their users to rank products, from which the algorithms can learn the consumer preferences.

The machine-learning algorithms often use the popularity of items or the nearest neighbor of similar users to make recommendations. While recommending the most popular movie to everyone might be harmless, recommending the most attractive partner to everyone could be a disaster—a form of *congestion* in the matching market.[5] Congestion has several welfare implications. First, congestion increases the matching time and decreases the number of matches made. Second, congestion discourages the participation of new users since the recommendation opportunity is skewed toward a few users. To the best of our knowledge, this is the first paper that studies the welfare consequences of recommender systems.

We propose integrating an economic model into a machine-learning algorithm. Our equilibrium recommender system inherits the predicting power from machine-learning, and it solves the congestion problem by the market allocation mechanism in economics. We use a rich dataset containing more than 490,000 unique users from an online dating service. On this platform, a user can browse other users' profiles, and he/she can "like" other users. The role of "like" is dual. From a user's perspective, it serves as a signaling device that is similar to that of Facebook. From the platform's perspective, it reveals information about mate preferences that is similar to that of Pandora. We exploit the information of likes and other

[4] According to a recent survey from the Pew Research Center, 15% of American adults have used online dating services. Source: http://www.pewinternet.org/2016/02/11/15-percent-of-american-adults-have-used-online-dating-sites-or-mobile-dating-apps/

[5] For example, several male users from a popular online dating service expressed their frustration when the algorithm matched more than 150 men with a single woman in New York City. Source: https://www.nytimes.com/2018/08/20/style/tinder-dating-scam-union-square.html

users' profiles to construct the recommenders.

Overview of the Algorithms

Our equilibrium matching algorithms consist of three modules: (1) a baseline machine-learning algorithm, (2) a pseudo-market, and (3) a separable matching model. For the baseline algorithm, we use both a flexible fixed-effect regression (OLS) and a matrix factorization (MF) to estimate the preference parameters. They represent two polar cases of machine-learning algorithms: *content filtering* and *collaborative filtering*, respectively.[6] Given these estimates, we apply the transferable utility (TU) matching model proposed by Choo and Siow (2006) (hereinafter CS) to the *pseudo-market*, an approximated matching market that includes potential partners and rivals constructed from the browsing history.

As the seminal marriage model studied in Becker (1973), the CS model adjusts the demand-and-supply through the "market price" or "matching cost" to dishearten users from seeking the relationship with the popular partners—those who have higher market prices. Specifically, the CS model computes the equilibrium utility, defined as the gross utility of the match net of the matching cost for any possible pair. For each user, we rank his/her potential partners within his/her pseudo-market according to the implied equilibrium utilities, and we recommend the N best women/men. Since agents may take the matching cost into account, deploying an economic model may help to improve the prediction power of the baseline machine-learning algorithm. On the other hand, the baseline machine-learning algorithms tend to produce similar recommendations due to the highly correlated gross-utility estimates. By introducing the matching cost, the equilibrium matching algorithms can achieve more diversified recommendations.

[6]Content filtering utilizes user-item attributes, whereas collaborative filtering extracts information only from the past user actions.

Main Findings

We evaluate our proposed algorithms from two aspects: the prediction power and diversification. These measures try to address two crucial concerns: First, do our recommender provide quality information? Second, are the recommendations diversified enough to promote the opportunity to be recommended and to avoid competition?

We use the *hit rate* to answer the first question. We define it as the probability that at least one member from the recommended list is liked. The baseline machine-learning algorithms have almost zero hit rates. When combined with our pseudo-market and equilibrium matching approach, the hit rate achieves about 35% when we recommend ten users. This suggests that our recommenders can predict users' behavior.

Next, we evaluate the diversity of our recommendations and its welfare consequences. First, we calculate the *coverage rate*: the number of distinct recommended users divided by the number of recommendations. The baseline algorithms have coverage rates lower than 30%, whereas the equilibrium matching algorithms achieve coverage rates above 60%. In particular, the equilibrium collaborative filtering algorithm has almost 100% coverage. Second, we simulate the number of rounds needed to clear the matching market in a Gale-Shapley style algorithm. Our equilibrium matching algorithms cut the number of rounds needed by 50% compared to the baseline machine-learning algorithms. These results from congestion suggest that our algorithms avoid recommending only a small number of people to most users, which in turn improves the welfare by accelerating the matching process and allocating the opportunity to be recommended more evenly.

Related Literature

Our paper is related to the board literature on both the recommender systems in computer science and the matching models in economics. In one-sided recommenders, Koren et al. (2009) adopt a novel collaborative filtering algorithm—Matrix Factorization—that outperforms Netflix's algorithm in predicting movie ratings. In a different vein, algorithms that are tailored for the two-sided setup, known as the *reciprocal* recommenders, are studied in

Malinowski et al. (2006), Brozovsky and Petricek (2007), Richards et al. (2008), Krzywicki et al. (2010), Akehurst et al. (2011), and Kunegis et al. (2012), among many others.

Our paper is also related to the matching literature in economics. An important class of matching models is the *separable* matching models, including CS, Chiappori et al. (2015), Galichon and Salanié (2015), and Mourifié and Siow (2017) for the TU cases.[7] We refer to Chiappori and Salanié (2016) for a more comprehensive survey. The primary focuses of this literature, however, are about parameter identification and comparative statics.[8] In contrast, we use the separable matching model for the prediction purpose.

The rest of the paper is organized as follows. Section 2 describes the mechanism of the online dating service and provides some summary statistics. Section 3 illustrates Equilibrium Content Filtering, and Section 4 presents Equilibrium Collaborative Filtering. We then show our results in prediction and congestion in Section 5 and Section 6, respectively. Section 7 concludes.

2. Data

2.1. Mechanism of Online Dating Service

The data we use comes from one of the major online dating services in Taiwan. Our user profile data includes 379,448 male accounts and 115,552 female accounts. Each user is required to sign up through his/her Facebook account, and the marital status must not be married. Furthermore, an account with more than 50 friends is required to deter robot or spam account. The users are then asked to create their profile, including age, height,

[7] Hitsch et al. (2010a) argue that the non-transferable utility (NTU) matching framework is more appropriate for online dating since it is unclear whether agents can exchange transfers without offline meetings. However, we find that the recommenders based on the NTU matching model of Galichon and Hsieh (2017) produce similar performances to the CS models. It is worth noting that the NTU matching models of Dagsvik (2000) and Menzel (2015) are not applicable since they produce only the equilibrium matching without the implied equilibrium utility.

[8] See, for example, Chiappori et al. (2010), Fox (2010), and Graham (2011)

physique, location, educational attainment, occupation, income, assets, and other characteristics. They could also upload their photos and a short paragraph of self-introduction. By entering search keywords on desired attributes, the users can *browse* the short profiles of other users including the profile photo, location, height, physique, and the short self-introduction. Three further actions can be made: "click", "like", and "sending message." First, the users can *click*[9] on the short profile and link to the detailed profile which includes all characteristics. Second, the users can *like* users of opposite sex. Third, the users can *send messages*.[10] We observe the detailed usage records from July 2016 to October 2017, such as who likes whom, along with several users' attributes.

An important feature on this website is the function of like. An unpaid male user can like at most 5 women in 12 hours, while the cap for a male user who pays the subscription fee (VIP) can like at most 15 women in 12 hours. On the other hand, the cap for a female user is 40 men in 12 hours regardless of the VIP status.[11]

2.2. Summary Statistics

We first present the like and click patterns from our data. On average, male users clicked 490.30 profiles and sent 115.76 likes, whereas female users clicked 331.38 profiles and sent 67.94 likes.

In the machine learning context, the difficulty of making a recommendation depends on the *sparsity* of the data. Namely, the number of the observed ranking data versus the whole sample. In our data, the total number of clicks divided by the number of all possible pairs is 0.04%, which is much smaller compared to the sparsity in the Netflix data (1%).

We compute several summary statistics of users' attributes from our data, and contrast with general population from the sample of Taiwan Social Change Survey (TSCS) in the year

[9]The previous study by Hitsch et al. (2010b) focus on estimating the mate preferences from the click data

[10]A female user could also send messages without additional charge, while a male user has to pay subscription fee to enable this function.

[11]One can also purchase additional quota for sending likes.

of 2015. TSCS is a academia-collected cross-sectional data surveyed in 2015[12], recording many attributes that overlap with our online dating dataset.

A first comparison of the two raw dataset is shown in Table 1. We find that people in the online dating dataset are quite different from the general population. First, there are much more men than women in the online dating website: the male to female ratio is almost 3 to 1. Second, the users of online dating service tend to be younger than the population. Third, more than half of our sample have Bachelor's degree, while only 35% have Bachelor's degree in the population.[13] Our sample is biased toward young, higher educated male population. However, since our objective is to predict matching in the online dating market, a bias sample is less of a concern.

3. Equilibrium Matching Algorithms

Our goal is to recommend partners that a user may want to send likes since this is the main source of the profit for the platform. We propose equilibrium matching recommenders based on the separable matching models à la CS. The model takes a collection of men and women, and each agent's utility over potential partners to form a discrete-choice demand system for both men and women. A pricing mechanism clears the market as in a typical competitive equilibrium analysis: to match with an over-demanded agent, one has to pay a "matching cost", which decreases the net utility of matching. The model essentially trades off between desirability and rivalry (measured by the endogenous matching cost). For each agent, our equilibrium matching algorithm ranks partners according to the corresponding net utilities and recommends the N best partners. Since the model explicitly takes into account competition, it is less likely that it recommends the same set of partners to observationally similar agents.

In what follows we discuss the implementation details about: (1) how to compute the

[12]For more details we refer to http://www.icpsr.umich.edu/icpsrweb/ICPSR/studies/23400

[13]This gap is driven by the significant expansion of colleges in Taiwan in 2000. As a result, the younger cohort in our online dating sample has much higher education attainment than the general population.

equilibrium matching, (2) how to estimate the mate preferences, and (3) how to choose the relevant set of men and women (the pseudo-market). Our algorithm is modular: different methods, such as different estimators for preferences or different matching models can be deployed to achieve a better performance.

3.1. Separable Matching Models

In this section we review the *separable* matching models as the framework to predict the matching outcome. At the core of these models is the combination of the classical matching theory (e.g., Gale and Shapley (1962) and Shapley and Shubik (1972)) that can predict matching outcomes and the random utility theory (e.g., McFadden (1976)) that can incorporate unobserved preference heterogeneity.

We first introduce a few notations. Let $m = \{1, 2, 3, \ldots, M\}$ be the index for men, and $w = \{1, 2, 3, \ldots, W\}$ be the index for women. Let x_m and z_w be the vectors of attributes of man m and woman w, respectively. The utility of man m who matches with woman w consists of a deterministic component α that depends on the observable attributes (x_m, z_w) and an *additively separable* random component $\varepsilon_{x_m z_w}$:

$$\alpha_{mw} = \alpha_{x_m z_w} + \varepsilon_{x_m z_w}. \tag{1}$$

Similarly, the utility of woman w who matches with man m is assumed to be

$$\gamma_{mw} = \gamma_{x_m z_w} + \eta_{x_m z_w}. \tag{2}$$

Agents can choose to be single: a single man obtains utility $\varepsilon_{x_m 0}$, and a single woman obtains utility $\eta_{0 z_w}$. We define Φ_{mw}, the social surplus, to be the sum of men's and women's deterministic utility matrices.

$$\Phi_{mw} = \alpha_{mw} + \gamma_{mw} \tag{3}$$

The separable models exclude interactions between unobserved characteristics of m and w conditional on observed types (x_m, z_w). It also implies that the preference ranking over partnership only depends on the observed characteristics. However, as the online dating contains rich set of regressors, we argue that the separability assumption is less problematic than in the traditional setup.

Following the celebrated TU framework of Becker (1973), CS postulate that a type-x man should pay τ_{xz} to match with a type-z women. The price system adjusts the demand and supply such that the number of type-x men who wish to match with a type-z woman equals the number of type-z women who wish to match with a type-x man:

$$n_x^M \text{Prob}\{z = \text{argmax}_{z\in\mathcal{Z}_0}(\alpha_{xz} - \tau_{xz} + \varepsilon_{x_x z})\} = n_z^W \text{Prob}\{x = \text{argmax}_{x\in\mathcal{X}_0}(\gamma_{xz} + \tau_{xz} + \eta_{xz})\}, \quad (4)$$

where n_x^M and n_z^W are the total number of type-x men and type-z women. Galichon and Salanié (2015) prove the existence and uniqueness of the TU matching equilibrium that satisfies Eq. (4) under fairly general conditions for the random utility shocks. In particular, if the random utility components $\varepsilon_{x_m z_w}$ and $\eta_{x_m z_w}$ follow an i.i.d. type-I extreme value distribution, CS show that the price system τ_{xz} in Eq. (4) can be concentrated out, and the equilibrium matching $(\mu_{xz}, \mu_{x0}, \mu_{0z})$ satisfies the following system of equations.

$$\begin{aligned} &\mu_{xz} = \mu_{x0}^{0.5}\mu_{0z}^{0.5}\exp\left(\frac{\alpha_{xz}+\gamma_{xz}}{2}\right) \\ &\mu_{x0} + \textstyle\sum_{z\in\mathcal{Z}} \mu_{xz} = n_x^M, \quad \forall x \in \mathcal{X} \\ &\mu_{0z} + \textstyle\sum_{x\in\mathcal{X}} \mu_{xz} = n_z^W, \quad \forall z \in \mathcal{Z}, \end{aligned} \quad (5)$$

where μ_{xz} is the number of type-x men who marry type-z women, and μ_{x0} (μ_{0z}) is the number of type-x men (type-y women) who remain single. It is also interesting to note that the equilibrium matching only depends on the social surplus $\Phi_{xz} = \alpha_{xz} + \gamma_{xz}$.

We rank partners according to the net utility $\alpha_{xz} - \tau_{xz}$. Therefore, our algorithm will recommend type-z women more often if the desirability of type-z women, as reflected in α_{xz}, is higher. On the other hand, our algorithm will recommend type-z women less often if the matching cost of type-z women, as reflected in τ_{xz}, is higher. The matching cost deserves further explanation. Clearly, if τ_{xz} is higher than τ_{zy}, it means that type-x men are less competitive than type-z men when facing type-z women. Such competition may be driven by the stronger preference for type-z men from the perspective of type-z women, or may be driven by more abundant type-x men in the market. The matching cost τ can summarize these two very different sources of competition into one single number. By taking into τ, the equilibrium matching algorithm avoids recommending partners with which there is only a dim chance to match. Moreover, by avoiding recommending the "superstars"—those who have higher market price τ —our algorithm can effectively solve the congestion issue.

While the separable matching models are well-known, the way we apply it to the recommender systems departs from the existing literature in three important ways. *First*, the separable matching models are often used to explain the aggregate marriage distribution μ, and the implied matching cost τ is often treated as the nuisance parameter in estimation. We, on the other hand, utilize τ to solve the congestion problem in recommender systems. *Second*, for the ease of estimation, researchers often group individuals into discrete categories, e.g., age category, when estimating the model. Since our focus is prediction, we do not face such a practical constraint in estimation. We solve the matching equilibrium by treating each individual as an unique type.[14] By doing so, the matching cost will be individual-specific, not group-specific as in CS. This further prevents recommending the same partners to everyone within the group defined by attributes. *Third*, online dating is a platform, where players can connect to many potential partners. The separable matching models, however, are designed for the one-to-one matching markets. The use of the one-to-one matching model is rather a simplified approximation to the complex truth.

To solve the matching equilibrium, one has to first estimate the utility parameters (α, γ) and specify the set of men and women, which we shall discuss next.

[14]This is a strategy advocated by Galichon et al. (2016) in maximum likelihood estimation of separable matching models.

3.2. Preference Estimation

Our baseline model is the following linear probability model

$$y_{mw} = \alpha_0 + \alpha_1 z_w + \alpha_2 d(x_m, z_w) + u_{mw}, \tag{6}$$

where y_{mw} is the binary indicator of whether man m likes woman w, and $d(x_m, z_w)$ is some distance function between the attributes of man and woman; e.g., the absolute value of difference in age. α_1 is intended to capture *vertical* preferences, in which users rank those attributes in the same way. α_2 is intended to capture *horizontal* preferences or homophily. For example, if $d(x_m, z_w)$ is the difference of years of schooling of man m and woman w, negative α_2 would mean the preference for similarity in education.[15] We use the fitted value $\hat{\alpha}_1 z_w + \hat{\alpha}_2 d(x_m, z_w)$ as the deterministic utility α_{mw} in the matching model. Analogously, for female users, we estimate $\gamma_{x_m z_w}$ by the following regression

$$y_{wm} = \gamma_0 + \gamma_1 x_m + \gamma_2 d(z_w, x_m) + v_{wm}. \tag{7}$$

The regressors we put include user's marriage status (never married or widowed), smoking and drinking status, occupation, religion, own evaluation of social ability, and whether one has kids. We also include whether one wants to have kids, ideal living arrangement, ideal house chores splitting, and one's attitude toward relationships. Distance measures include whether they live in the same city, their difference in heights and ages.

As users send likes after clicking the profiles, for each man m, we restrict to the set of women whose profile has been clicked by m. This leaves us with 31,301,828 observations when estimating Eq. (6). The same procedure also applies to each woman, leaving us with 16,333,134 observations when estimating Eq. (7). By the above sample construction, our data is a quasi-panel, in which N equals the number of men (women), and T equals the number of women (men) being clicked. Therefore, we further control for the individual fixed

[15] Similar specifications are also adopted in HHA, Logan et al. (2008) and Galichon et al. (2016).

effect when estimating Eq. (6) and (7).[16] For example, when estimating a man's preference, we add the man's fixed effect. The fixed effect can control for the fact that some users systematically send more likes than the others, and hence may bias the utility estimation. Lastly, we add rich interaction terms to account for preference heterogeneity. In particular, we divide income, asset, education, and physique into 3, 2, 4, and 4 categories, respectively. The interaction terms include all pairwise interaction of these 4 attributes. For instance, we interact man's income with woman's education, man's education with woman's income, etc. Since our main goal is prediction rather than explaining the mate preferences as in HHA, we do not report the regression results here.[17]

3.3. Equilibrium Content Filtering and Pseudo-Market

Below we formally introduce our first equilibrium matching algorithm—the Equilibrium Content Filtering based the OLS-estimated preferences.

The block 1 of Algorithm 1 constructs the pseudo-market. For a generic user m_i, we first find the set of women O_1^W he clicked before. O_1^W are the women he might be interested in as revealed by his browsing history. On the other hand, those women in O_1^W also clicked other men's profiles; therefore, men in O_1^M are m_i's competitors as revealed by the clicking histories of women in O_1^W. This procedure can continue until the desired number of men and women are reached. For example, women in O_2^W are competitors of those women in O_1^W. The pseudo matching market thus contains m_i's direct and indirect competitors and potential partners and these partners' competitors. The block 2 of our algorithm takes the estimated preference illustrated in section 3.2 as the inputs, and compute the equilibrium introduced in section 3.1 to make the recommendations.

The rationale for a pseudo-market is the following. While in principle all registered users constitute the matching market, it is computationally infeasible to deliver the recommended

[16]In principle, one can estimate a two-way fixed effect model to control for both men's and women's individual heterogeneity. However, our sample size goes beyond what canned software (e.g., Stata) can handle.

[17]The regression results are available upon request.

Algorithm 1: Equilibrium Content Filtering Algorithm

◇ **Block 1:**

input : history of clicks of all users

output : the pseudo matching market of man m_i: $\big(M(m_i), W(m_i)\big)$

begin

1. Find the set of women that man m_i clicked: O_1^W
2. For each woman $w \in O_1^W$, randomly select K_1 men that w clicked: $O^M(w)$.
   $O_1^M = \cup_{w \in O_1^W} O^M(w)$
3. For each man $m \in O_1^M$, randomly select K_2 women that m clicked: $O^W(m)$.
   $O_2^W = \cup_{m \in O_1^M} O^W(m)$

repeat

step 2 and step 3

until *desired number of men and women are reached*

4. **return** $M(m_i) = m_i \bigcup O_1^M$, $W(m_i) = O_1^W \bigcup O_2^W$

◇ **Block 2:**

input : $\big(M(m_i), W(m_i)\big)$ and the estimated preferences $\big(\alpha_{x_m z_w}, \gamma_{x_m z_w}\big)$,
$\forall m \in M(m_i), w \in W(m_i)$

output : N recommended women

begin

1. solve the matching equilibrium as defined in Eq. (4) and (5).
2. sort the resulting equilibrium net utility of m_i, $\alpha_{x_{m_i} z_w} - \tau_{x_{m_i} z_w}$, $\forall w \in W(m_i)$.
3. **return** *the women correspond to the first N highest utilities for m_i*

list in real time by solving such a large scale equilibrium problem. Further, since users send likes after clicking profiles, it suggests that one can restrict the attention to a subset of users from the clicking history.

4. Matrix Factorization and Equilibrium Collaborative Filtering Algorithm

Being the core algorithm deployed by the winning team of the Netflix Prize, Matrix Factorization attracts an increasing attention in the domain of machine learning and recommender systems. Here we briefly review Matrix Factorization, and we propose a novel Equilibrium Collaborative Filtering algorithm that combines the strength of Matrix Factorization and the CS model.

4.1. Matrix Factorization

Matrix Factorization (MF) assumes that y_{mw}, user m's rating on item w, can be approximated by the inner product of K-dimensional latent factors p_m and q_w, where p_m represents m's attributes, and q_w represents w's attributes. Since each user rates only a sparse fraction of the available items, most of the entries in y_{mw} are missing. MF uses these low-dimensional latent factors to impute the missing rating data.

In our case, y_{mw} is binary: $y_{mw} = 1$ if man m liked woman w, and $= -1$ otherwise. y_{mw} is missing if m did not click w's profile.[18] Formally, MF estimates the latent factors by solving the following minimization problem:[19]

$$\min_{p_m, q_w; m=1,\dots,M, w=1,\dots,W} \sum_{m=1}^{M} \sum_{w=1}^{W} (y_{mw} - p_m^{'} q_w)^2 + \frac{\lambda_p}{2} ||p_m||_2^2 + \frac{\lambda_q}{2} ||q_w||_2^2, \tag{8}$$

[18]In most of the applications, the percentage of the non-missing y_{mw} is rather small. In our case, only 0.2% of y_{mw} is non-missing, while in the Netflix case 1% is non-missing.

[19]It is possible to consider more general specifications as described in the influential paper by Koren et al. (2009). We only consider the baseline specification due to the limitation of the available canned software that will be discuss later in Appendix A.

where $||\cdot||_2$ is the l-2 norm. The penalization terms are added to avoid over-fitting as in the standard Lasso regression.

4.2. Equilibrium Collaborative Filtering

MF represents a polar case of recommender systems that do no rely on the observed user-item attributes. However, since the estimated latent factors essentially play the same role as the observed attributes, we argue that the fitted value of $\hat{y}_{mw} = p^{'}_{m} q_w$ in MF can be treated as the social surplus matrix Φ_{mw} in the CS matching model. MF implies that men's and women's j-th latent factor are *complement*: $\frac{\partial^2 y_{mw}}{\partial_{p_{mj}} \partial_{q_{mj}}} = 1$, which further implies that the fitted value of $\hat{y}_{mw}$ can be interpreted as a supermodular social surplus function.

We elaborate the idea with the following toy example. Suppose there are 2 latent factors, $K = 2$: the first latent factor indicates whether one smokes (1 for smokers and -1 for non-smokers), and the second latent factor indicates whether one drinks (1 for drinkers and -1 for non-drinkers). Suppose there are one man m_1 with latent factor $p_{m_1} = (1, -1)'$ and two women w_1 and w_2 with latent factors $q_{w_1} = (1, 1)'$ and $q_{w_2} = (1, -1)'$. In this example, both m_1 and w_2 smoke but do not drink, while w_1 smokes and drinks. The fitted value for the two matches would be: $r_{m_1 w_1} = 1 \times 1 + (-1) \times 1 = 0$ and $r_{m_1 w_2} = 1 \times 1 + (-1) \times (-1) = 2$. In this simple case, we can see that y_{mw} is higher when both sides have similar attributes. We summarize the Equilibrium Collaborative Filtering algorithm that combines the MF and the CS model below:

Algorithm 2: Equilibrium Collaborative Filtering

◇ **Block 1:**

input : history of clicks of all users

output : the pseudo matching market of man m_i: $\big(M(m_i), W(m_i)\big)$

begin

1. Find the set of women that man m_i clicked: O_1^W
2. For each woman $w \in O_1^W$, randomly select K_1 men that w clicked: $O^M(w)$.
   $O_1^M = \cup_{w \in O_1^W} O^M(w)$
3. For each man $m \in O_1^M$, randomly select K_2 women that m clicked: $O^W(m)$.
   $O_2^W = \cup_{m \in O_1^M} O^W(m)$

repeat
  step 2 and step 3
until *desired number of men and women reached*

4. **return** $M(m_i) = m_i \bigcup O_1^M$, $W(m_i) = O_1^W \bigcup O_2^W$

◇ **Block 2:**

input : $\big(M(m_i), W(m_i)\big)$ and the estimated social surplus from the MF as defined in Eq. (8): $\Phi_{mw} = p_m' q_w$, $\forall m \in M(m_i), w \in W(m_i)$

output : N recommended women

begin

1. solve the matching equilibrium as defined in Eq. (4) and (5).
2. sort the resulting equilibrium net utility of m_i, $\alpha_{x_{m_i} z_w} - \tau_{x_{m_i} z_w}$, $\forall w \in W(m_i)$.
3. **return** *the women correspond to the first N highest utilities for m_i*

5. Prediction

We compare the performances of two algorithms in terms of the *hit rate*, which is defined as the probability that at least one member from the recommended list is liked. Formally,

Definition 1: (Hit Rate) Consider I men, $\{m_1, m_2, \ldots, m_I\}$, and for each man we recommend N women. Suppose we recommend $\tilde{W}(m_i) = \{\tilde{w}_1^i, \tilde{w}_2^i, ..., \tilde{w}_N^i\}$ to m_i, and in the data, the set of women that m_i liked is $W^o(m_i) = \{w_1^i, w_2^i, ..., w_L^i\}$. Define $h_i = 1\{W^o(m_i) \cap \tilde{W}(m_i) \neq \emptyset\}$. The hit rate is defined as $\bar{H} = \frac{1}{I}\sum_{i=1}^{I} h_i$.

Intuitively, h_i measures that if the recommender system can suggest at least one woman that m_i will like. The hit rate is therefore the proportion of making a successful recommendation list.

We first discuss the out-of-sample hit rate of several recommenders out of 5,000 men. We only consider those who sent between 20 to 30 likes, which are around the average number of likes in the the data.[20] The results are summarized in Figure 1, and the implementation details are gathered in Appendix A.

Since our algorithms are modular, we can investigate the prediction power of each step. We consider the following recommenders:

---

1. OLS: for each m, choose women who correspond to the N highest fitted value $\hat{\alpha}_{mw}$ estimated by the procedure described in section 3.2.
2. MF: for each m, choose women who correspond to the N highest fitted value $p_m^{'} q_w$ estimated by the procedure described in section 4.
3. Pseudo-Market: randomly choose N women from m's pseudo-market defined in the block 1 of Algorithm 1.
4. OLS+Pseudo-Market: same as OLS, except here one chooses the top-N women from the pseudo-market instead of the whole sample.
5. MF+Pseudo-Market: same as MF, except here one chooses the top-N women from the pseudo-market instead of the whole sample.
6. OLS+Pseudo-Market+CS: Equilibrium Content Filtering defined in section 3

[20]The extreme cases in which men liked an exceedingly small or large number of women will lead to a hit rate that is close to zero or one, regardless of the algorithm.

7. MF+Pseudo-Market+CS: Equilibrium Collaborative Filtering defined in section 4

Figure 1: Hit Rate

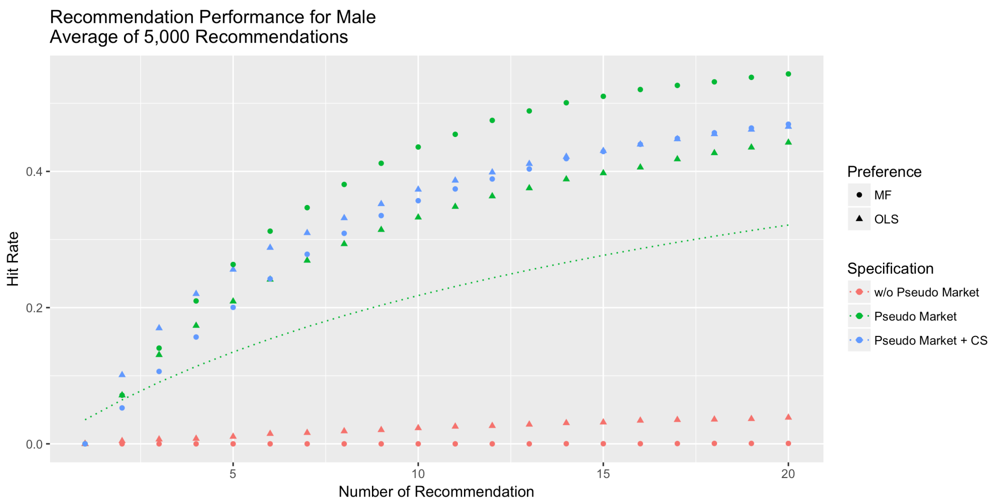

We begin with the OLS and the MF recommender as the baseline cases, and then add other modules such as the pseudo-market and the CS model. The average hit rate under different N, the length of the recommendation list, can be found in Figure 1. We use triangle to label algorithms that are based on OLS, and we use circle to label algorithms that are based on MF. We use the dashed line to present the hit rate of the pseudo-market. The algorithms are further color-coded according to the different modules deployed within the algorithm.

We find that the plain OLS recommender (a variant of content filtering) outperforms the plain MF recommender (a variant of collaborative filtering). In fact, the MF recommender has near zero prediction power for the like pattern even when $N = 20$. In contrast, the OLS recommender has about 2% hit rate when $N = 10$. The poor performances of both algorithms are mainly due to the extremely sparse observations of likes. By restricting to the pseudo-market, even random draws of recommendations performs much better than OLS or MF alone. When $N = 10$, the hit rate is slightly above 20%. If we combine pseudo-market and OLS approach (green triangle), when $N = 10$, the hit rate increases by a factor of 1.8 comparing with the random draw from the pseudo-market. The combination of pseudo-market and MF (green circle) doubles the hit rate under the random draw from the pseudo-market. Lastly, we add the TU matching model of CS to make recommendation. The Equilibrium Content Filtering (OLS+Pseudo-Market+CS, labeled by blue triangle) further improves the hit rate of the OLS+Pseudo-Market Algorithm, whereas the Equilibrium Collaborative Filtering (MF+Pseudo-Market+CS, labeled by blue circle) decreases the hit rate of the MF+Pseudo-Market Algorithm. Overall, the hit rate curves of these two equilibrium recommender algorithms are similar and outperform the randomized draws from the pseudo-market.

6. Congestion: Counterfactual Welfare Analysis

Both content and collaborative filtering can lead to *congestion*—the recommendation lists cluster around only few players. The content filtering utilizes user-item attributes to make predictions; therefore, men who share similar attributes will be recommended to similar women by construction. The collaborative filtering such as the nearest neighbour approach utilizes the choice patterns to make predictions; consequently, men who are within the same neighbourhood will be recommended to the similar women. On one hand, highly clustered recommendation list may increase dating competition. On the other hand, congestion also implies that the dating chance is not evenly allocated by the algorithm. If potential users realize that certain algorithms limit their opportunity to be recommended, it may discourage them from joining the platform, which can undermine the platform's long term revenue. In this section, we attempt to measure the degree of congestion and its welfare consequences. We find that deploying the CS model can alleviate the congestion problem. We present several Monte Carlo evidences in this section.

6.1. Coverage

First, we consider a measure for the *coverage rate of recommendation*, defined as the number of the distinct recommended women divided by the number of recommendations.

Definition 2: (Coverage Rate of Recommendation) Suppose we recommend $\tilde{W}(m_i) = \{\tilde{w}_1^i, \tilde{w}_2^i, ..., \tilde{w}_N^i\}$ to m_i, $i = 1, \ldots, I$. Define the coverage rate of the top-N recommendations as $c(N) = \frac{1}{I \times N} | \cup_{i \in I} \cup_{n=1}^{N} \tilde{w}_n^i |$, where $| \cdot |$ is the cardinality of a set.

Clearly, $c(N)$ is a decreasing function of congestion. In the case of top-1 recommendation, suppose every man is recommended to the same woman, i.e., $\tilde{w}_1^i = \tilde{w}_1^{i'} \quad \forall i, i'$, then $c(1) = \frac{1}{I} \times |1| = \frac{1}{I}$. On the other hand, if every m is recommended to a distinct w, then $| \cup_{i \in I} \tilde{w}_1^i | = I$ and hence $c(1) = 1$. Therefore, the higher the coverage rate, the more diversified the recommendation lists are.

For simplicity, for all simulations in the section, we will remove the module of the pseduo market.[21] We randomly draw 500 markets, with each market contains 1,000 men and 305 women (the gender ratio in the data). We use the OLS preference estimates and the MF latent factors obtained from the whole sample to compute the coverage rates of the top-1 recommenders produced by the OLS, MF, Equilibrium Content Filtering, and Equilibrium Collaborative Filtering.

Figure 2: Coverage Measure

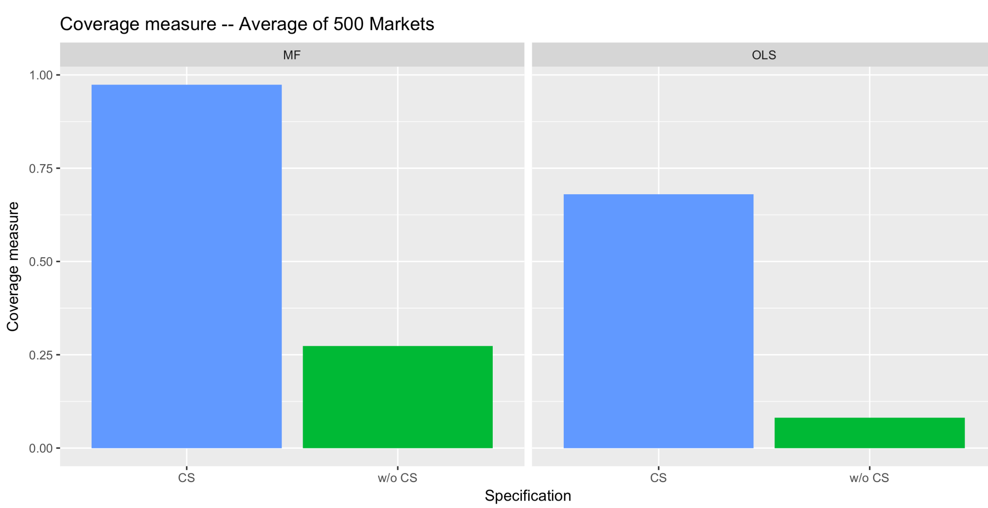

[21]The pseudo-market is *personalized*: every one faces a distinct matching market in the algorithm. This leads to some difficulty in counterfactual simulations. As will become clear later, the (simplified) Gale-Shapley algorithm requires a common matching market.

The results are shown in Figure 2. The coverage rate of OLS (the green bar in the right panel) is close to 10%, whereas the coverage rate of MF (the green bar in the left panel) is about 25%. This is the evidence that the non-equilibrium based algorithms can produce highly clustered recommender list. Moreover, the MF has higher coverage rate than the OLS algorithm. This is because the MF approach implicitly estimate *individual-level* coefficients. Richer preference heterogeneity allowed in the MF thus translate into more diversified recommender lists.[22]

While it is possible to deploy more complex models to increase the coverage rate, it will generally lead to an over-fitting issue. However, we show that the matching model can improve the coverage rate without increasing the complexity of the model. From Figure 2, the coverage rate of both OLS and MF (the blue bars) increase several folds. In particular, MF almost achieves full coverage when combined with the CS model.

6.2. Congestion

We further conduct a simulation study to understand the effect on congestion from the highly clustered recommendation lists. As pointed out by Che and Koh (2016), economists have yet to develop a benchmark model to analyze the issue of congestion.[23] We therefore resort to the classical Gale-Shapley (GS) algorithm (Gale and Shapley (1962)), with some modifications to fit the online dating market.

The simplified Gale-Shapley works as follows:

---

1. Each man contacts the first woman in the recommendation list generated by the recommender algorithm.

[22]It is possible to consider random coefficient specifications in the OLS approach to achieve similar effects. We leave this as a future endeavor.

[23]Several attempts have been made including Che and Koh (2016) and Galichon and Hsieh (2017), and there references therein. However, it is unclear how to apply those models to simulate congestion.

2. Each woman replies to the best man she receives according to the utility estimated by OLS in Eq. (7) and ignores all others. Additionally, if the best man gives the woman negative utility, she does not reply to anyone.

3. The men who received response by some women drop out from the market. The men who are ignored continue to contact the second best women in the recommendation list.

4. Each woman replies to the best man from the new contact list and ignores all others.

5. repeat step 3-4 for several rounds

The simplified GS algorithm is only meant to simulate how fast that a man can find a woman who is willing to reply him. It has nothing to do with stability or strategy-proofness as studied in the classical matching theory (e.g. Roth and Sotomayor (1990)). Several remarks are in order. First, men are assumed to be *myopic*: they simply follow their recommendation lists as if it is their true preference ordering. Here we abstract from the fact that in the real world, men can both conduct their own search and consider the recommendations. On the other hand, women use the estimate preference parameters to determine to which men they reply. This asymmetric behavioral assumption is merely a way to isolate the congestion effect from the clustered recommendation lists of one side of the market. Second, there is no "deferred-acceptance" phase as in the classical GS algorithm. The platform we study serves as a platform for participants to send likes and exchange messages that does not enforce "one-to-one" matching. Therefore, as soon as a man receive a feedback from a woman, the purpose of the platform is achieved, and we can drop out that man to simplify the congestion calculation. Moreover, the dating platform is not a real one-to-one marriage market, hence at each round there is no need for women to compare the existing offers versus the new proposals as in the classical GS algorithm. We therefore assume women keep in touch with the men to whom they previously reply (those who are dropped from the simplified GS algorithm), and they continue to consider other men from the new contact lists.

We run the simplified GS algorithm. The boxplot of the number of matches made in the first round over the 500 simulated markets is shown in Figure 3. We focus on MF and OLS

Figure 3

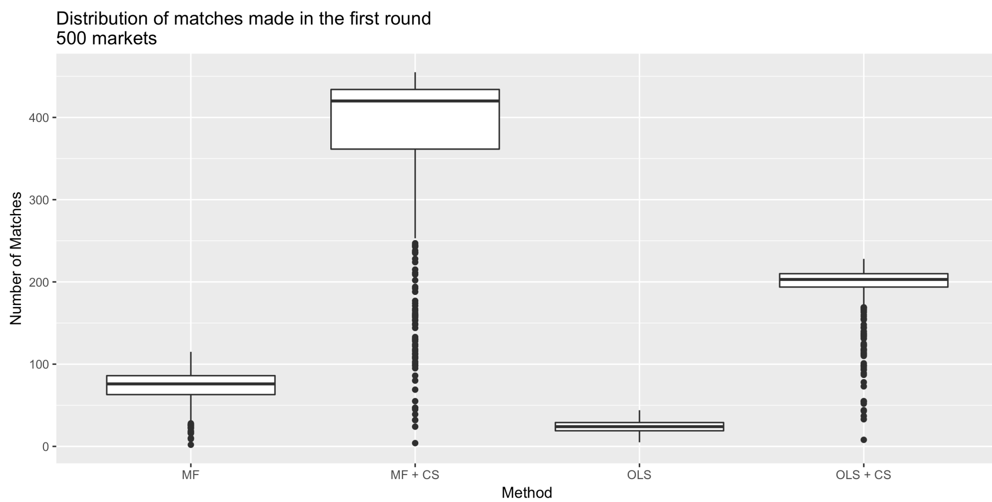

without CS first. The median number of matches of MF and OLS is 75 and 25 respectively. Since we have 1,000 men in the simulated market, it is about 7.5% and 2.5%. This patterns is consistent with the coverage rate reported in Figure 3. [24] The interquartile range of the boxplots of OLS and MF are tight, suggesting little Monte Carlo variations. When combining with CS algorithm, the number of matches made in the first round increase several times for both MF and OLS. For MF+CS, the median is more than 400, or 40% of the men in the market. For OLS+CS, the median can also be as high as 200, or 20% of the men in the market. In other words, deploying the CS model creates 5 times more matches for the MF and 8 times more matches for the OLS, respectively.

We further summarize the average number of matches made in for the first 10 rounds (those men who are dropped from the algorithm) in Figure 4.

The green lines correspond to the MF and OLS algorithm without CS, whereas the blue lines correspond to the results combining with CS. The market clears much faster in the

[24] Generally, the number in these two graphs will not be equal because in the simplified GS algorithm matches are made only when women accept the offers.

Figure 4

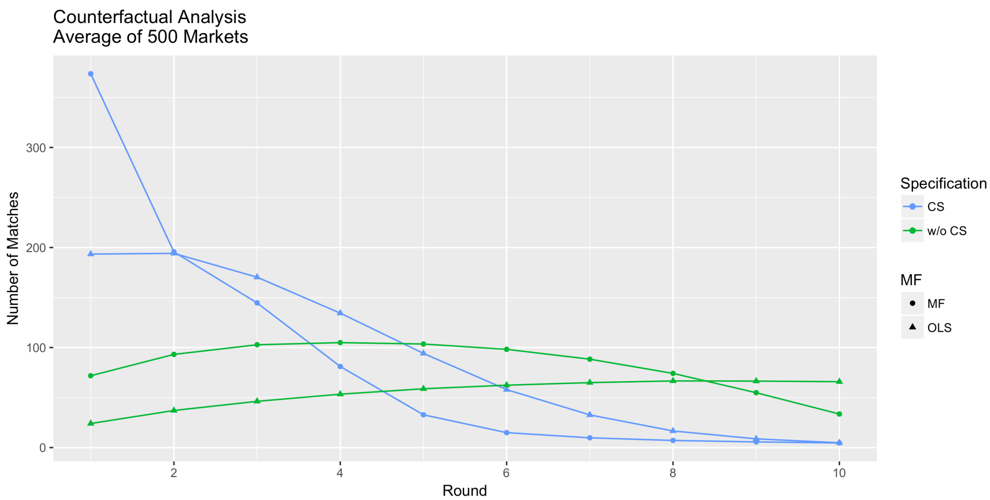

presence of the CS module; the blue lines exhibit an exponential decreasing pattern as they successfully match more pairs in the first few rounds. By contrast, the green lines are relatively flat; the lesser-diversified lists limit the number of matches at each round.

Lastly, we present the boxplot of rounds needed to clear 95% of the men in the market in Figure 5.[25] The medians suggest that it takes about 10 rounds for MF and 15 rounds for OLS to clear 95% of the market.[26] When combining with CS, both MF and OLS clear the market much faster. The CS module effectively reduces the number of rounds by a factor of two: It takes less than 5 rounds for MF + CS and about 6 rounds for OLS + CS.

In summary, we find that deploying the CS model can substantially accelerate the matching process, approximated by the simplified Gale-Shapley algorithm. The recommendations are more diversified and the number of rounds required is only half of that without the CS module. These results suggest that one can mitigate the congestion problem by deploying the CS model.

[25]We truncate it at the 95% due to the tail behavior of the simplified Gale-Shapley algorithm: there exist few unattractive men that require excessive rounds to fully clear the market.

[26]While the interquartile ranges are tight, there are several outliers. We find that these outliers correspond to the case when women are more "selective." Namely, the utilities of the recommended men are all negative.

Figure 5

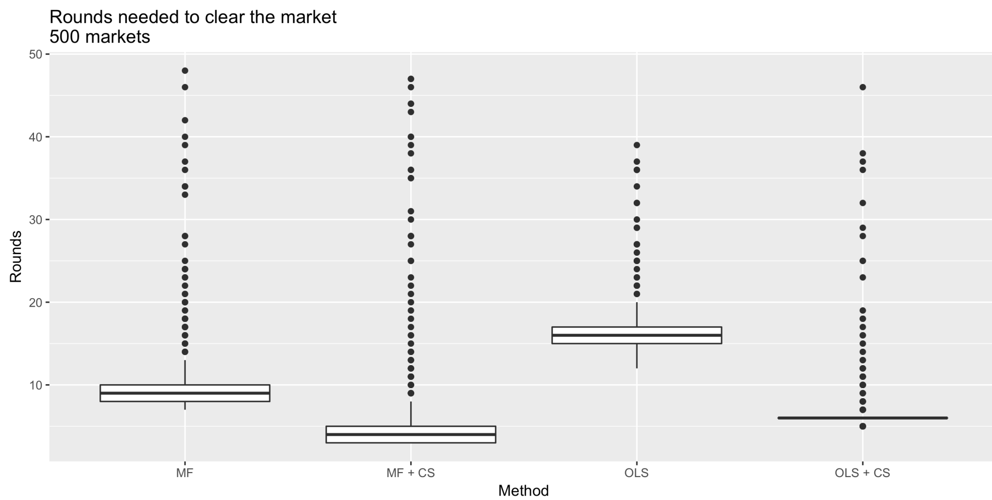

7. Conclusion and Discussion

We suggest using the separable matching model in conjunction with the content or collaborative filtering, which are originally designed for one-sided recommenders. The separable matching models are easy to deploy in practice. It is related to the optimal transport theory in which there exist several high performance algorithms. Moreover, it can seamlessly blend into either content or collaborative filtering, without increasing the complexity of the underlying machine-learning algorithm.[27] The resulting equilibrium recommenders can accommodate two-sided preferences and rivalry in online dating. In particular, we find that it can improve the hit rate of the baseline OLS recommender.

On the other hand, the equilibrium recommenders help solve the issue of congestion. Interestingly, despite the efforts to make matching more efficient, the platform's incentives might not be aligned with the users. Currently, many platforms charge a monthly subscription fee for usage; therefore, a platform may not have the incentive to reduce congestion or

[27] One can construct a recommender with a combination of (1) any machine-learning algorithm that produces an estimate or an approximation to the gross utility, and (2) any matching model that produces a rankable equilibrium utility indices. This combination of machine-learning algorithms and matching models is therefore a "tuning parameter" that can be optimized in practice.

minimize the time an user spent on the platform. The Two-sided platforms face a trade-off between the length of an user's subscription and the platform's reputation in making matches—two conflicting goals in terms of congestion. Platforms may choose a less congested algorithm under alternative pricing model, such as a "matching guarantee program": the platform charges a one-time fee until a match is found. We leave this for future research.

Reference

Akehurst, J., I. Koprinska, K. Yacef, L. A. S. Pizzato, J. Kay, and T. Rej (2011): "CCR-A Content-Collaborative Reciprocal Recommender for Online Dating." in *IJCAI*, 2199–2204.

Becker, G. S. (1973): "A theory of marriage: part I," *Journal of Political Economy*, 81, 813–846.

Brozovsky, L. and V. Petricek (2007): "Recommender System for Online Dating Service," *arXiv:cs/0703042*.

Che, Y.-K. and Y. Koh (2016): "Decentralized College Admissions," *Journal of Political Economy*, 124, 1295–1337.

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

Chiappori, P.-A. and B. Salanié (2016): "The Econometrics of Matching Models," *Journal of Economic Literature*, 54, 832–861.

Chiappori, P.-A., B. Salanié, and Y. Weiss (2015): "Partner Choice and the Marital College Premium: Analyzing Marital Patterns Over Several Decades," *working paper*.

Choo, E. and A. Siow (2006): "Who Marries Whom and Why," *Journal of Political Economy*, 114, 175–201.

Dagsvik, J. K. (2000): "Aggregation in Matching Markets," *International Economic Review*, 41, 27–57.

Fox, J. (2010): "Identification in Matching Games," *Quantitative Economics*, 1, 203–254.

Gale, D. and L. S. Shapley (1962): "College Admissions and the Stability of Marriage," *The American Mathematical Monthly*, 69, 9–15.

Galichon, A. and Y.-W. Hsieh (2017): "A Model of Decentralized Matching without Transfers," *working paper*.

Galichon, A., S. Kominers, and S. Weber (2016): "Costly Concessions: An Empirical Framework for Matching with Imperfectly Transferable Utility," *Forthcoming, Journal of Political Economy*.

Galichon, A. and K. O'Hara (2016–2018+): "TraME: Transportation Methods for Econometrics," an open-source software available at https://github.com/trame-project.

Galichon, A. and B. Salanié (2015): "Cupid's Invisible Hand: Social Surplus and Identification in Matching Models," *working paper*.

Graham, B. S. (2011): *Handbook of Social Economics*, North-Holland, vol. 1, chap. Econometric Methods for the Analysis of Assignment Problems in the Presence of Complementarity and Social Spillovers, 965–1052.

Hitsch, G. J., A. Hortacsu, and D. Ariely (2010a): "Matching and Sorting in Online Dating," *American Economic Review*, 100, 130–163.

——— (2010b): "What Makes You Click? Mate Preferences In Online Dating," *Quantitative Marketing and Economics*, 8, 393–427.

Juan, Y.-C., W.-S. Chin, Y. Zhuang, B.-W. Yuan, M.-Y. Yang, and C.-J. Lin (2016): *LIBMF: A Matrix-factorization Library for Recommender Systems*.

Koren, Y., R. Bell, and C. Volinsky (2009): "Matrix factorization techniques for recommender systems," *Computer*, 30–37.

Krzywicki, A., W. Wobcke, X. Cai, A. Mahidadia, M. Bain, P. Compton, and Y. S. Kim (2010): "Interaction-based collaborative filtering methods for recommendation in online dating," in *International Conference on Web Information Systems Engineering*, Springer, 342–356.

Kunegis, J., G. Gröner, and T. Gottron (2012): "Online dating recommender systems: The split-complex number approach," in *Proceedings of the 4th ACM RecSys workshop on Recommender systems and the social web*, ACM, 37–44.

Logan, J. A., P. D. Hoff, and M. A. Newton (2008): "Two-Sided Estimation of Mate Preferences for Similarities in Age, Education, and Religion," *Journal of the American Statistical Association*, 103, 559–569.

Malinowski, J., T. Keim, O. Wendt, and T. Weitzel (2006): "Matching People and Jobs: A Bilateral Recommendation Approach," in *Proceedings of the 39th Annual Hawaii International Conference on System Sciences-Volume 06*, IEEE Computer Society, 137–3.

McFadden, D. (1976): *Behavioral Travel-Demand Models*, Heath and Co., chap. The Mathematical Theory of Demand Models, 305–314.

Menzel, K. (2015): "Matching Markets as Two-Sided Demand Systems," *Econometrica*, 83, 897–941.

Mourifié, I. and A. Siow (2017): "The Cobb Douglas Marriage Matching Function: Marriage Matching with Peer and Scale Effects," *working paper*.

Qiu, Y., C.-J. Lin, Y.-C. Juan, W.-S. Chin, Y. Zhuang, B.-W. Yuan, and M.-Y. Yang (2017): *recosystem: Recommender System using Matrix Factorization*, r package version 0.4.2.

Richards, D., M. Taylor, and P. Busch (2008): "Expertise recommendation: A two-way knowledge communication channel," in *Autonomic and Autonomous Systems, 2008. ICAS 2008. Fourth International Conference on*, IEEE, 35–40.

Roth, A. E. and M. A. O. Sotomayor (1990): *Two-Sided Matching: A Study in Game-Theoretic Modeling and Analysis*, Cambridge University Press.

Shapley, L. S. and M. Shubik (1972): "The Assignment Game, I: The Core," *International Journal of Game Theory*, 1, 111–130.

Appendix A. Software and Numerical Details

To construct the pseudo-market used in both Algorithm 1 and 2, we set $K_1 = 5$[28] and $K_2 = 3$. We further perform an additional random sampling to ensure the gender ratio in the pseudo-market coincides with that of the whole sample. Under these tuning parameters, roughly there are 264 men and 81 women in the pseudo-market.

To solve the problem of MF, we use **recosystem**[29], an R wrapper of the **LIBMF** library developed by Juan et al. (2016). **LIBMF** uses a novel parallel stochastic gradient descent algorithm to efficiently solve the high dimensional optimization problem as defined in Eq. (8). We use the tuning function provided by **recosystem** to select the user-chosen parameters. We choose $K = 15$ latent factors for both men and women, and we choose the regularization parameters $\lambda_p = \lambda_q = 0.0001$. We choose the following parameter values that are associated with the numerical optimization. The learning rate is 0.01, the number of thread for parallel computation is 8, the number of bin is 100, and the maximum number of iteration is 200,000.

We use **TraME**, an R library developed by Galichon and O'Hara (2016–2018+), to solve the matching equilibrium. We treat each individual as an unique type, and hence the number of each type defined in Eq. (4) is set to one: $n_x^M = n_z^W = 1$. Notice that the CS model only depends on the social surplus matrix Φ_{mw} as can be seen in Eq. (5). For Algorithm 1, $\Phi_{mw} = \big(\alpha_1 z_w + \alpha_2 d(x_m, z_w)\big) + \big(\gamma_1 x_m + \gamma_2 d(z_w, x_m)\big)$, where (α, γ) are estimates by the fixed effect regression as described in section 3.2. For Algorithm 2, $\Phi_{mw} = p_m^{'} q_w$. The sovler returns the men's equilibrium net utility $U_{mw} \equiv \alpha_{mw} - \tau_{mw}$ and women's equilibrium net utility $V_{mw} \equiv \gamma_{mw} + \tau_{mw}$, where $\Phi_{mw} = U_{mw} + V_{mw}$. We use U_{mw} to rank potential partners for both algorithms.

We use the following Bootstrap sampling procedure to construct the training and testing sample when comparing the out-of-sample performance of the hit rate of the Equilibrium Content Filtering algorithm.

1. Randomly draw m_i
2. Construct his pseudo-market $(M(m_i), W(m_i))$
3. The pseudo-market of m_i is the testing sample. The rest of the observations are the training sample for estimating (α, γ).

[28] If w clicked less than K_1 men then we include all these men.
[29] See Qiu et al. (2017)

4. Use the estimated regression coefficient to impute Φ_{mw} for $m \in M(m_i)$ and $w \in W(m_i)$.

5. Solve the matching equilibrium and compute h_i as defined in Definition 1

6. Repeat I times to calculate the empirical hit rate $\bar{H}$ as defined in Definition 1

Regarding the MF method, it is not possible to split the sample by completely excluding people in the pseudo-market. Unlike the regression-based approach, in which the utility parameters in the testing sample can be extrapolated from the training sample, the MF requires that some non-missing observations in any row and any column of y_{mw} to be supplied to the algorithm. While repeated OLS fitting and model-solving in the Boostrap sampling is computationally cheap, this is not the case for MF. To avoid repeated training, we use all available observations to train the MF model. When we make the recommendation from the pseudo-market as in the Equilibrium Collaborative Filtering, it is inevitable that all women in the pseudo-market have been used in the training phrase. As a result, the hit rate of the MF-related methods reported in Figure 1 is rather an in-sample prediction. To investigate this potential issue, We further use only 50% and 20% of the available y_{mw} to train the MF model. Under these cases, only part of the women in the pseudo-market are used in the training phrase. However, the results remain similar.

Appendix B. Tables

Table 1: Raw Comparison

		Dating Dataset	TSCS
Gender			
	Male	76.66%	48.94%
	Female	23.34%	51.06%
Education			
	<Secondary	1.07%	19.74%
	Secondary	4.28%	11.82%
	High School	34.70%	25.75%
	College	52.83%	35.84%
	Master	6.65%	5.96%
	Doctor	0.48%	0.83%
Physique			
	Slim	11.70%	
	Medium	67.26%	
	Chubby	16.49%	
	Fat	4.54%	
Income			
(Annual NTD)	No Report	24.53%	
	< 300K	24.56%	
	300K to 450K	18.45%	
	450K to 600K	14.31%	
	600K to 800K	8.86%	
	800K to 1M	5.06%	
	1M to 2M	2.97%	
	2M to 5M	0.65%	
	> 5M	0.62%	
Assets			
(NTD)	No Report	63.18%	
	< 3M	30.71%	
	3M to 9M	3.53%	
	9M to 15M	1.17%	
	15M to 30M	0.68%	
	30M to 60M	0.26%	
	60M to 90M	0.07%	
	> 90M	0.41%	
Marriage			
	Married	0%	57.80%
	Single	88.00%	29.54%
	Divorced	10.89%	4.19%
	Widowed	1.11%	8.37%
Among Non-married			
	Single	88.00%	70.17%
	Divorced	10.89%	9.95%
	Widowed	1.11%	19.88%
Mean-Age			
		31.59	56.53