Search Frictions, Sorting, and Matching in Two-Sided Markets*

Hyesung Yoo† Song Yao‡ Ravi Bapna§ Jui Ramaprasad¶

This version: April 2023

Abstract

This paper studies the impact of friction and preferences on match formation in two-sided markets. Because every agent has private information about their preferences for potential matches' characteristics, forming a match based on mutual compatibility requires extensive costly search. We use field experiment data from an online dating platform to better understand the relative impact of search cost and preference on match outcomes. During the field experiment, randomly selected users are provided information about the preferences of their potential partners that can only be obtained through costly search otherwise. We find evidence suggesting that reducing frictions through this information provision leads to less sorting among matched couples in terms of their characteristics. This is because a user often assesses the match probability with a potential partner based on the similarity between their respective characteristics. The information provision allows a user to assess the match probability more accurately using the potential partner's preference rather than characteristics. Consequently, it encourages users to initiate contact with those who are more likely to match despite their characteristics differences, leading to less sorting among matched couples. To investigate the relative contribution of frictions and preferences on assortative matching, we develop and estimate a model that incorporates frictions and preference heterogeneity across users. Our estimation results reveal that frictions play a significant role in shaping matching outcomes. Using model estimates, we simulate matches under the frictionless Gale-Shapley protocol, and we find that removing frictions leads to significantly less sorting between couples. We also find that frictions in our platform lead to a significant reduction in efficiency. These results highlight the importance of platform designs that aim to reduce frictions. More importantly, with one-third of the marriages in the U.S. marriages originating from online encounters, this paper shows how the design of an online platform can contribute to diversity in the marriage market.

*We thank Rodrigo Belo, Tat Chan, Yulia Nevskaya, Raphael Thomadsen, and Maria Ana Vitorino for their comments. We also thank feedback by seminar participants at Arizona State University, Columbia University, INFORMS Marketing Science Conference 2020, KAIST, the Ohio State University, Santa Clara University, the University of Miami, the University of Minnesota, the University of Toronto, the University of Virginia, Washington University in St. Louis, and the WISE Conference 2019.

†University of Toronto: hyesung.yoo@utoronto.ca

‡Washington University in St. Louis: songyao@wustl.edu

§University of Minnesota: rbapna@umn.edu

¶University of Maryland: jramapra@umd.edu

1 Introduction

In many two-sided marketplaces, agents search potential partners (henceforth, candidates) to form a match based on a mutual agreement.[1] Since agents on both sides of the market have preferences regarding each other's characteristics, and because these preferences are often private, agents on each side do not know which counterparts are willing to transact.[2] Forming a match based on mutual compatibility in the presence of private preferences, therefore, generally requires extensive, costly search.

The question of who matches with whom has been a central question in the matching literature. In the presence of costly search, preferences and search frictions both shape the formation of a match. Understanding the relative effects of these two forces on match outcomes is not only theoretically important but also managerially relevant, because the design of online two-sided platforms should vary depending on whether match outcomes primarily result from preferences or search frictions.

In this paper, we try to obtain a better understanding of the relative impact of preferences and various frictions (including search frictions) on match outcomes using data from an online dating platform. We do this in two ways: First, we disentangle the relative impact of preference and frictions on assortative matching (i.e., positive sorting), a widely observed phenomenon where couples display resemblance across various characteristics such as age, education level, ethnicity, and income. Sorting in married couples is an important topic of study—among other things, it exerts long-term effects on economic development and inequality through its impact on the outcomes of children and accumulation of human capital (Raquel and Rogerson (2001); Raquel (2003)). Second, we quantify the gains in user welfare when frictions are removed from the platform. This quantification can offer insights on the gains that users can achieve when better platform designs reduce frictions.

There are two distinct explanations for assortative matching (Hitsch et al. (2010a)): (1) Positive sorting is an equilibrium outcome driven by agents' preferences. For example, if mate preferences are "horizontal," people may prefer to match with a similar partner, which in turn results in assortative matching. Or, if mate preferences are "vertical," in the sense that everyone evaluates candidates using the same criteria, then the ranks of matched partners will be positively correlated. In this case, couples will display positive sorting along attributes that are monotonically related with these rankings. (2) Search frictions influence how couples meet, regardless of their preferences. For

[1] Examples can be found across a wide range of industries, including marriage markets, college admissions, online labor markets (Taskrabbit and Upwork), and the hospitality industry (AirBnB).

[2] Fradkin (2015) shows that more than 40% of booking inquiries on AirBnB platform are rejected. Approximately 14% of those rejections are driven by hosts' preferences regarding the characteristics of the searcher or the trip.

example, in *offline* dating markets, the existence of couples who attended the same university might not be due to preferences, but instead might reflect the fact that it is more likely for the matched couples to meet and date because they are on the same campus. In reality, *both* preferences and frictions affect how couples meet.

In an *online* dating context, types of search costs vary depending on the platform design. Typically, on online dating platforms (and in the dating market in general) the decision to initiate a date request (or send a message) depends on the probability of the counterpart accepting that request. As a result, if a user has identical preferences towards two candidates, they should send the date request to the candidate with a higher likelihood of accepting the request. Generally, because people prefer others who are similar to themselves, candidates who are more similar to the focal user are ex ante more likely to accept an offer than those who are dissimilar. Therefore, unless a user searches for additional information that helps form a more accurate prediction about the match probability, the user will make the offer to a more similar candidate, which will in turn result in assortative matching. Hence, while assortative matching in online dating platforms may be due in part to people's preference for others who are similar to themselves, it can also be influenced by search frictions, because users with high search costs will not search for additional information; they will make decisions based only on the limited information provided in the default setting.

The unique feature of our data is that it was generated from a field experiment in an online dating platform. The treatment of the experiment was to provide randomly selected users a piece of information about the preference of their potential matches—an information as to whether a candidate in the profile had *ℓiked* the user.

To be specific, many online dating platforms give users the option to *ℓike* a candidate's profile by either clicking a *ℓike* button or swiping right (similarly, a user can not *ℓike* a candidate by simply not clicking on the *ℓike* button or swiping left). A *ℓike* from the opposite side serves as a positive signal about the likelihood of a successful match; When two users *ℓike* each other, they both get a notification about their mutual *ℓiking*, which can encourage users to start a conversation through messaging. In the default setting (control group), users did not know whether the candidate in the profile had *ℓiked* them. The only way for the focal user to find out was by *ℓiking* the candidate: If the focal user *ℓikes* a candidate, and if the candidate had also *ℓiked* the focal user, both get notified about the mutual *ℓiking*. If the focal user receives no notification upon *ℓiking* a candidate, this means that the candidate either did not *ℓike* the focal user or has yet to encounter/decide on the focal user's profile.[3]

The act of *ℓiking* a profile is costly. While the simple act of clicking a *ℓike* button seems costless, there is a psychological cost associated with the action of *ℓiking*. This cost can be incurred due to several reasons, including: (i) not finding a profile sufficiently appealing, (ii) prospect of not being *ℓiked* back (or not getting a response) by a candidate, which can hurt one's ego (Baumeister et al. (1993)), or (iii) the possibility of needing to reject a candidate after triggering a positive response from him/her, in which case the user faces the negative emotion associated with having to reject the candidate.[4]

Unlike for the control group, who had to incur the above mentioned psychological cost to find out if they were *ℓiked* by the candidate, the experiment allowed the treatment group to know *upfront* whether the candidate had *ℓiked* them, without having to *ℓike* a profile first. Because the majority of initiated messages receive no response, knowing whether someone had *ℓiked* them allows users to more precisely gauge the likelihood of getting a match. Since this information was revealed to the treatment group without them having to take any costly action, the treatment reduced search frictions.

In the data, we find descriptive evidence suggesting that the treatment changes sorting patterns between matched couples across various dimensions. Specifically, when users matched with those who had *ℓiked* them, couples in the treatment group displayed significant differences in attributes (age, education level, body type, race, popularity) compared to the control group.[5] Since the treatment reduces search frictions by revealing information about who *ℓiked* them without users having to *ℓike* first to find out, the descriptive patterns we see in the data suggest that reducing frictions can lower the degree of positive sorting between couples.

The mechanism behind this pattern, according to our data, is that the treatment encouraged users to initiate a conversation with candidates who had *ℓiked* them even though they are dissimilar from themselves. In other words, when *ℓiked* by dissimilar candidates, users in the treatment group knew this and were encouraged to initiate a conversation with them. Users in the control group, on the contrary, did not know this information (unless they had engaged in a costly search), and

[3]The focal user is not able to distinguish between these two causes.

[4]Research in psychology has shown that the object of unwanted affection experiences annoyance and frustration, and that rejecting the romantic overture may cause guilt, discomfort, and other distress (Baumeister et al. (1993)).

[5]We only consider matches where the users of our experiment initiated the conversation through messaging.

therefore were less likely to initiate a conversation with a dissimilar candidate because they perceived the match probability to be low.

To disentangle the impact of frictions and preferences on assortative matching, and to quantify the impact of frictions on user welfare, we need to compare the matches in a market with frictions to those in a frictionless environment where only preferences shape the matching outcomes. Although the experiment reduced treatment group's search frictions by providing the information about who *ℓiked* the users without users having to *ℓike* first to find out, some uncertainty nevertheless remains for the treatment group because they may still get no response after sending a message to a candidate despite being *ℓiked* (in this paper, we define a "match" as people exchanging at least four messages, following Bapna et al. (2022)). In addition, if the cost of composing a message is nonnegligible, the decision to start a conversation with a candidate would also depend on the perceived match probability. Therefore, we build a structural model of search for a partner and use model estimates to simulate matches under a counterfactual frictionless environment and compare those to matches formed in the presence of frictions.

More specifically, we model the decision process of a user who is considering whether to search for and to contact a candidate of the opposite gender. Our model incorporates preference heterogeneity across users and also allows for costly search for the information about *ℓikes*. For each candidate's profile that a user browses, he decides whether to take the costly action of *ℓiking* and/or messaging. This framework allows us to model users' decisions to search for and contact candidates, and to model how these decisions relate to their preferences and costs.

Based on the model estimates, we predict who matches with whom under various counterfactual protocols. Specifically, we simulate matches under the default (control) setting, treatment setting, as well as the frictionless environment. By comparing the matches in the presence of frictions to those in a frictionless environment, we disentangle the relative impact of frictions and preferences on sorting.

We find that frictions play a significant role is shaping assortative matching patterns and that removal of frictions leads to a significant reduction in positive sorting. Specifically, approximately 9% of positive sorting in age, 12% of positive sorting in years of schooling, 40% of positive sorting in popularity, and 21% of positive sorting in race is due to frictions.

We then turn to the question of efficiency. We first examine how much users become better off compared to the default control setting if the platform makes the information about *ℓikes* available to all users in both sides of the market, and then examine how much users gain when all the frictions are removed. To do this, we assign ordinal rankings to each matched partner based on estimated preference parameters and compare the average rankings achieved across different protocols. We find that reducing frictions by providing the information about $\ell ikes$ leads to a small but significant improvement (in terms of highest achievable ranking, 0.73 percentage points for men and 0.52 percentage points for women) in the average ranking of the partner compared to the default (control) setting. In terms of utility net of costs, providing the information about $\ell ikes$ leads to a gain in utility by 34% for men and 13% for women compared to the default setting. These results suggest that reducing search frictions significantly improves on the outcomes compared to a market with frictions. When all frictions on the platform are removed, the average ranking of the partner leads to a significant improvement of 9.9 percentage points for men and 3.5 percentage points for women in terms of the highest achievable ranking.

The rest of this paper is structured as follows. In the following section we review the related literature and this paper's contributions. Section 3 describes the institutional details of our dating platform. Section 4 details the experimental design. Section 5 summarizes the data, and Section 6 presents descriptive evidence suggesting that reducing frictions may lead to a reduction in positive sorting. In Section 7, we propose a model of search for a partner. Estimation details and estimation results are discussed in Sections 8 and Section 9, respectively. Section 10 presents counterfactual exercises. Section 11 concludes.

2 Related Literature

This paper is closely related to recent literature on frictions in online two-sided markets involving matching. Fradkin (2015) examines search frictions in Airbnb, an online market for short-term rental housing. The author shows that on Airbnb, even after a buyer expresses interest of renting an apartment, the transaction can fall through because the seller can reject the buyer, or because multiple buyers may contact the seller at the same time. The paper studies how ranking algorithms can increase the efficiency of the platform by increasing the number of matches. Horton (2014) shows that failed transactions due to information frictions are also common in online labor markets, where employers inefficiently pursue oversubscribed workers. The author explores how the platform can optimally allocate employers' attention to workers. While our research and these previous papers all study frictions in two-sided online markets, our focus is on the impact of frictions on assortative matching.[6]

Our paper is also related to empirical work that estimates mate preferences in romantic relationships (Wong (2003); Choo and Siow (2006); Flinn and Del Boca (2012); Chan et al. (2015); Richards Shubik (2015)). For example, Wong (2003) uses a two-sided search model to explain marriage outcomes using the Panel Study of Income Dynamics data (PSID). Choo and Siow (2006) applies a frictionless transferable utility matching model to the marriage market. Arcidiacono et al. (2016) estimate a two-sided search model of romantic-relationship formation; they show that individuals direct their search based on partners' characteristics, endogenously determined probabilities of matching, and the terms of a relationship (e.g., whether sex is included). While these papers use observed matches in equilibrium to estimate mate preferences, our data document each user's entire search process. This data enable us to estimate preferences based on users' actions at each stage of the decision process.

Finally, our paper is also related to the growing literature in speed-dating and online dating (Kurzban and Weeden (2005), Fisman et al. (2006), Fisman et al. (2008), Hitsch et al. (2010a), Hitsch et al. (2010b), Lee (2015), Lee and Niederle (2015), Bapna et al. (2016), Halaburda et al. (2017), Fong (2018)). Kurzban and Weeden (2005), Fisman et al. (2006), and Fisman et al. (2008) use speed-dating data to study mate preferences. Lee (2015) finds that online dating promotes marriages with weaker sorting along occupation and geographic proximity but stronger sorting along education and other demographic traits. Lee and Niederle (2015) study the effect of preference signaling (e.g., sending costly virtual roses with a dating request) on a major Korean online dating website; they find that the signaling increases the success probability of a dating request

More recently, Fong (2018) studies search and matching behavior in an online dating app, focusing on how users respond to market thickness; Bojd and Yoganarasimhan (2019) study the causal effect of popularity information in online dating; they find evidence of strategic shading due to fear of rejection. Bapna et al. (2022) study how revealing "who likes you" affects user behavior in online dating.

The papers in these areas that are closest to ours are by Hitsch et al. (2010a) and Banerjee et al. (2013). Using data from an online dating website to study the efficiency of matches, Hitsch et al. (2010a) find that the matches predicted by the economic model (Gale-Shapley deferred-acceptance

[6]Related research on information frictions and market inefficiency is a paper by Fréchette et al. (2008). In many scenarios such as markets for new physicians/law graduates, professional athletes drafting, and college admissions, early matches can be inefficient if crucial information for determining the match quality evolves over time.

algorithm) are similar to the actual matches achieved on the dating website, suggesting that these matches are approximately efficient. They are also able to largely predict the assortative matching patterns observed in the matches, which suggests that assortative matching can arise in the absence of search frictions, primarily due to preferences and market mechanisms. Using a similar approach, Banerjee et al. (2013) study how preferences for caste can affect equilibrium patterns of matching. They find a very strong preference for within-caste marriages, and they also show that in equilibrium, ignoring caste-related preferences does not alter the matching patterns on non-caste attributes.

Compared to these papers, our experiment enables us to detect a source of search friction, which motivates us to disentangle the contribution of friction and preferences on the equilibrium sorting and user welfare.

3 Institutional Details

The online dating platform our data comes from is a typical "freemium" community: most of the users sign up for a free account that allows them to use the basic features (browsing profiles, *ℓiking*, and sending messages) to interact with other members of the platform, while some users pay a monthly subscription fee for a premium account that consists of a fixed bundle of free and premium features, including the ability to know whether the candidate in the displayed profile had *ℓiked* the focal user.[7] During the period of our study, there was no limit in the number of *ℓikes* and messages that a user could send.

The platform is accessible through both its website and mobile app. Although the experiment was conducted on randomly selected users who use either type of device, in this paper we focus solely on the users who used the mobile app. In addition, among a few different ways to search for a partner, we only focus on users who were searching for partners through the "rapid matching" process (the actual term of the process is disguised) due to its simplicity. This process will be described in detail below. Note that our approach is different from Bapna et al. (2022), who studied user behavior across all search processes, rather than solely focusing on the *rapid matching* process. To begin, we will first describe how the *rapid matching* process works for a mobile app user in the default setting (i.e., a user in the control group with a nonpremium account), then proceed to describe how the experiment changed the operation of the app for users in the treatment group in the following section.

[7] To avoid selection bias, we sample only from the users with a nonpremium account.

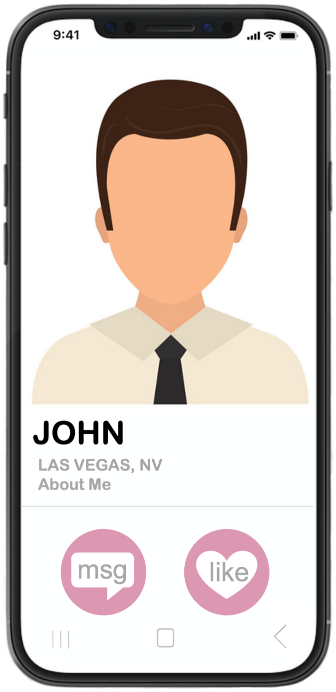

Note. This figure is for illustrative purpose only.

Figure 1: Profile of a Candidate Displayed to the Control Group

When a nonpremium user in the control group opens the mobile app, a random profile is displayed to him/her. Figure 1 illustrates what is displayed to a user. A user is able to see the candidate's profile picture and such characteristics as age, race, and education level. Upon browsing the profile, a user can choose to *like* and/or send a message. A user can choose to *like* by either clicking a *like* button (or by swiping right). Similarly, a user can choose to not *like* by simply not clicking the *like* button (or by swiping left). If the user chooses to *likes* a profile and the candidate had already *liked* the focal user, then both users receive a popup notification about their mutual *liking* (Figure 2b). In addition to the popup notification about the mutual *liking*, a heart icon appears next to the candidate's profile picture (in the upper right-hand corner) indicating that the candidate had *liked* the focal user.[8] Therefore, *liking* a candidate reveals whether or not the candidate had *liked* the focal user.

If neither a notification nor a heart icon appears upon choosing to *like*, it implies that the candidate had not *liked*, or had not yet seen and *liked*, the focal user (Figure 2a). The focal user is not able to distinguish between these two causes of "not *likes*."

[8]The heart icon is for illustrative purpose only. A different icon may appear in the actual app.

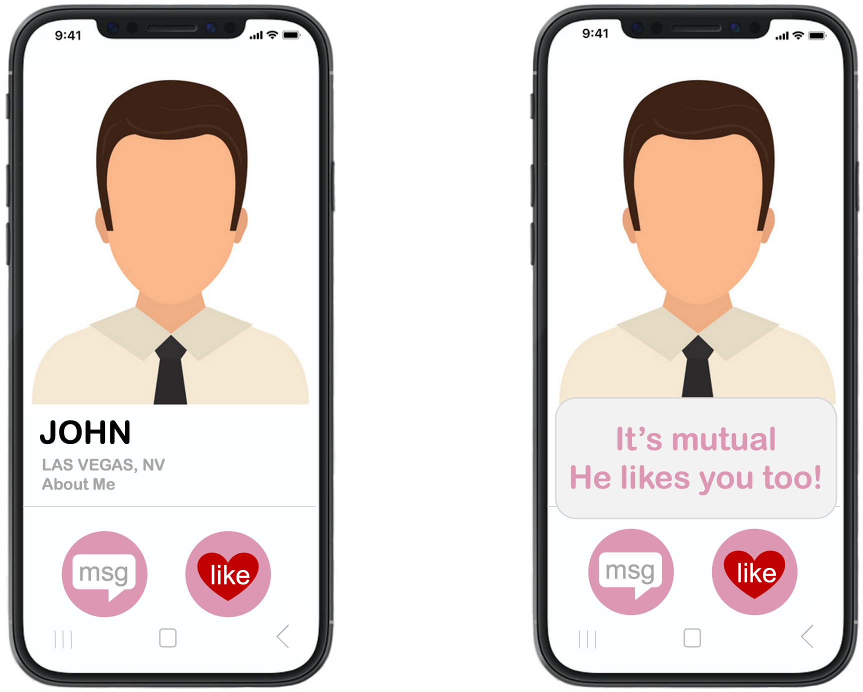

(a) The candidate did not *like* you

(b) The candidate *liked* you

Notes. (For illustrative purposes only. The actual icons and images may be different in the app). This figure illustrates how *liking* a profile reveals whether the candidate had *liked* the focal user or not. The *like* button turns red when the user chooses to *like*. If the candidate had *liked* the focal user, both users receive a notification about the mutual *liking*. In addition to a notification, a heart icon appears in the top right-hand corner. If the candidate had not *liked* the focal user, neither a notification nor a heart icon appears.

Figure 2: How *Liking* a Profile Reveals Information

A user will see a new candidate’s profile immediately after sending a message or clicking the “back” button. Also, if a user swipes in either direction as opposed to clicking the *like* button, a new candidate’s profile will be displayed—unless a mutual *liking* is reached, in which case Figure 2b is displayed and the user can decide whether to send a message. On the other hand, if a user clicks the *like* button instead of swiping right, he can continue browsing the current candidate’s profile (in this case, he/she sees Figure 2a if he/she hadn’t been *liked* and sees Figure 2b if he/she had been *liked*) and has to decide whether to send a message.

4 The Experiment

In this section, we describe the experiment design and how it changed the way the app operated for the treatment group. The experiment was conducted on 100,000 newly registered random users

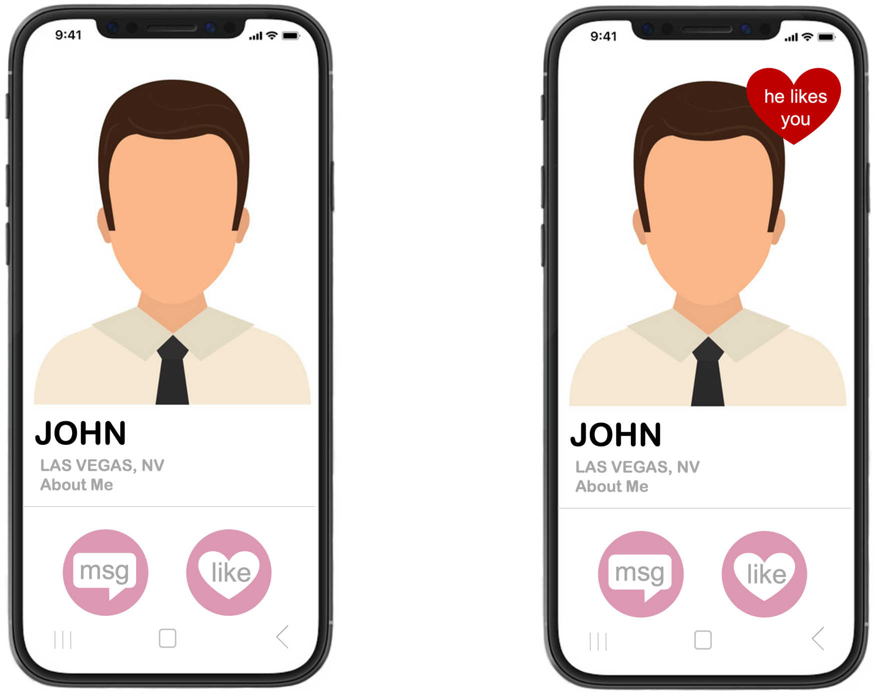

(a) The candidate did not *like* you

(b) The candidate *liked* you

Figure 3: Profile of a Candidate Displayed to the Treatment Group

of the platform (either website or the mobile app) over three consecutive months.[9] We refer to these months as the pretreatment period (month 1), the treatment period (month 2), and the posttreatment period (month 3).[10] On the first day of the treatment period, 50,000 randomly selected users received the following email:

> *Hey username,*
>
> *You have been randomly selected to receive a super power - for the next 30 days, we're giving you the ability to know whether someone had liked you! Normally this feature is restricted to paid premium users only. Enjoy!*[11]

The remaining 50,000 users who serve as our control group received the following email:

> *Hey username,*
>
> *It's a good time to visit our platform! Enjoy!*

[9] The target population was randomly selected from among newly registered users during a seven-day period in 2016. These users account for less than 1% of the entire population of the platform's users as of 2019.

[10] The gift of treatment expired after 30 days.

[11] To disguise the identity of the platform (and the terminologies specific to it), the messages presented in the paper are slightly modified from the actual messages that were sent to users.

Because the treatment was endowed on users by the platform without any required action on the users' part, users were unaware of being part of an experiment; therefore, observer bias wouldn't be present in this study.

Figure 3 illustrates what was displayed to a user in the treatment group when he/she opened the app. When a candidate's profile was displayed to a user, user could immediately see whether he/she was *ℓiked* by the candidate or not. If the user had been *ℓiked* by the candidate, a heart icon appeared in the upper right-hand corner (Figure 3b). If the candidate had not *ℓiked* the focal user, the heart icon was absent (Figure 3a). As mentioned earlier, a user was unable to distinguish whether the candidate had browsed his profile and decided not to *ℓike*, or whether his/her profile was not browsed yet by the candidate..

5 Data Description

For each of the 100,000 users in our experiment, we observe time-stamped actions (browsing, *ℓiking*, and messaging) over the three months. As explained earlier, the three months are pre-treatment period (1st month), treatment period (2nd month), and post-treatment period (3rd month). For each user, we have the following self-reported demographic variables: gender, sexual orientation, age, education level, race, and body type.[12] We also observe time-stamped actions and demographics for all correspondent users (i.e., candidates) who had interacted with the experimental users in any way. The data on correspondent users' *ℓiking* behavior allows us to observe who had *ℓiked* the focal user of our experiment. In addition, we observe whether a user was using a desktop or a mobile app, whether a user has a premium account, and whether the account is valid (whether or not the user is a spammer/bot).

Out of the initial sample of 100,000 experimental users, we limit our sample to mobile app users who were searching for a partner using the *rapid matching* process. We further limit our sample to heterosexual users who browsed at least one profile during the treatment period.[13] We drop users with a premium account to avoid selection bias, and we also drop users with a nonvalid account. The final sample consists of 8,142 treated and 7,663 control experimental users. In our data, less than 0.5% of the experimental users interacted with other members of the experiment, and less than 0.5% of the users in the treatment group interacted with other members in the treatment group. Hence, we do not worry about contamination bias.

[12] Only a few demographic variables were provided to us, due to privacy concerns.

[13] A large number of users became inactive just a few days after creating an account.

Table 1 summarizes characteristics of the experimental users in our sample, separately for men and women.[14] We have self-reported information on users' age, education level, body type, and ethnicity. Consistent with existing research that uses data from other online dating services, there are more men than women in our dataset (11,807 men, 3,998 women). Men are on average 31 years old and women are on average 35 years old. Among men and women who reported their education level, approximately 55% of men and 58% of women received their final degree from a university, and approximately 16% of users have a postgraduate degree. Most users in our dataset are White (64% men and 61% women), followed by Hispanic (11% men and 11% women), Asian (9% men and 13% women) and Black (8% men and 11% women). The test of randomization of the treatment is reported in Appendix A.

6 Descriptive Statistics

6.1 Impact of the Treatment on User Activities

We start by showing the effect of the treatment on user activities. From the perspective of the focal user, there are two types of candidates: (1) candidates who had *ℓiked* the focal user (henceforth "Likers"), and (2) candidates who did not *ℓike* the focal user (henceforth "NotLikers"). Since the treatment allows users to know whether the candidate had *ℓiked* him/her or not without having to take any further action, it is natural to think that treated users would behave differently depending on whether a candidate is a Liker or a NotLiker.

In Table 2, we present summary statistics of user activities toward Likers, separately for men and women. Column 1a (2a) summarizes activities of men (women) in the control group, and column 1b (2b) summarizes activities of men (women) in the treatment group. We also report the differences between the control and treatment group (columns 1c and 2c) and t-statistics (columns 1d and 2d) to show whether there are any significant differences between the two groups.

For men, treatment has no significant impact on the total number of profiles that a user browses, but it increases the number of *ℓikes* sent by 13% and increases number of initiated messages sent by 17%. We see a similar pattern for women: treatment has no significant impact on the total number of profiles that a user browses, but it increases the number of *ℓikes* sent by 32% and increases number of initiated messages sent by 34%.

[14]This table summarizes demographic characteristics of the experimental users only, not correspondent users.

	Men				Women			
Variable	Mean	SD	Median	Obs.	Mean	SD	Median	Obs.
Age	31.4871	9.1936	29	11,807	34.6666	11.0415	32	3,998
Popularity	0.1046	0.5214	-0.0420	11,807	-0.0207	0.5543	-0.0934	3,922
Education								
HighSchool	0.1220	0.3274	0	2,057	0.0652	0.2470	0	859
TwoYear	0.1760	0.3809	0	2,057	0.1420	0.3493	0	859
University	0.5455	0.4981	1	2,057	0.5797	0.4939	1	859
PostGrad	0.1565	0.3635	0	2,057	0.1565	0.3635	0	859
Body type								
Thin	0.4706	0.4992	0	3,011	0.3129	0.4639	0	1,256
Average	0.4165	0.4931	0	3,011	0.5725	0.4949	1	1,256
LittleExtra	0.0917	0.2886	0	3,011	0.0892	0.2851	0	1,256
Overweight	0.0213	0.1443	0	3,011	0.0255	0.1576	0	1,256
Race								
Asian	0.0861	0.2805	0	5,100	0.1298	0.3362	0	2,026
White	0.6447	0.4786	1	5,100	0.6101	0.4879	1	2,026
Black	0.0843	0.2779	0	5,100	0.1091	0.3118	0	2,026
Hispanic	0.1082	0.3107	0	5,100	0.1076	0.3100	0	2,026
Other Race	0.0767	0.2661	0	5,100	0.0434	0.2039	0	2,026

Notes. Many users choose not to report some of their demographic information, which leads to different number of observations for each demographic variable. In the data (prior to selecting our sample), we also have users with the following education levels: LawSchool, MedSchool, and PhD. In our final sample, although we do have candidates with these education levels, we do not have experimental users with these education levels. This is because experimental users with these education levels were dropped during our sample selection process (browsed at least one profile, heterosexual, valid, nonpremium, mobile app users).

Table 1: Summary Statistics of Users Characteristics

We also test whether treatment has an impact on the number of successful matches achieved by our experimental users. While we do not observe whether users actually went on an offline date or the actual content of the messages exchanged between users, we do observe the number of messages that were exchanged by each pair. Prior research by Bapna et al. (2016) and anecdotal evidence from the online dating industry have pointed out that exchange of three messages between potential couples is a good predictor of an actual online match, in which phone numbers are exchanged or users ask the other out for an offline date. In fact, senior executives of the platform revealed that they strongly believe that this measure of a match is an accurate predictor of an offline date. Moreover, despite knowing the exact content of users' messages, the platform uses this metric as a measure of a successful match. Following Bapna et al. (2022), we take a more conservative stance and define a successful match as an exchange of at least four messages. Here we only consider "initiated" matches, where at least four messages were exchanged upon the experimental user starting a conversation.

These results are also reported in Table 2. We find that the treatment increases men's initiated matches with Likers by 12% and increases women's initiated matches with Likers by 28%.

Table 3 presents summary statistics of user activities towards NotLikers. Since treated users are unsure whether the candidate had browsed his/her profile and decided not to *ℓike*, or whether his/her profile was not browsed yet by the candidate, the effect of treatment on user behavior is ambiguous. Except for the reduction in the number of *ℓikes* sent by men in the treatment group, the treatment does not lead to a significant difference in user activities nor the number of successful matches.

6.2 Frictions and Sorting

Positive correlation in mate attributes has been widely documented and studied in previous research across multiple disciplines. Our experiment reduces the search friction present in our platform for the treatment group by revealing the information about *ℓikes* that users received from candidates without having to take a costly action. Therefore, by looking at how the treatment affects sorting patterns between matched couples, we can get insights into the impact of search frictions on assortative matching.

User attributes in our data are as follows: age, education level (high school = −4, two-year college = −2, university = 0, masters = 2, law school= 3, medical school = 3, PhD = 6), body type (skinny, average heavier, overweight), and race (Asian, Black, Hispanic, White, other). In addition, we create a synthetic variable that measures the popularity of a user, which is the total count of *ℓikes* received divided by the sum of *ℓikes* and "not *ℓikes*" received. Specifically, user i's popularity is calculated as:

$$popularity_i = \frac{\#\ell ikes\text{Received}_i}{\#\ell ikes\text{Received}_i + \#\text{Not}\ell ikes\text{Received}_i}$$

This measure is then standardized to have a mean 0 and standard deviation of 1, separately for men and women.[15]

To compare the sorting patterns between the two groups, we first construct a measure of attribute difference (henceforth "attribute difference") that can be used to test whether treatment leads to a significantly different sorting pattern between a matched man and woman. Specifically, the attribute difference between a man m and a woman w is calculated as $\Delta = |x_m - x_w|$, where x_m and x_w are m's and w's characteristics, respectively. We obtain the attribute difference between couples

[15]We use data of both experimental and correspondent users to calculate the mean and variance when standardizing the popularity measure.

	Men				Women			
	control	treated	diff	t-stat	control	treated	diff	t-stat
	(1a)	(1b)	(1c)	(1d)	(2a)	(2b)	(2c)	(2d)
Number of browses								
Mean	3.1772	3.3291	0.1520	0.9334	41.7138	41.8165	0.1027	0.0341
Median	1	1			10	10		
SD	10.3180	6.9631			89.4961	99.9036		
Likes sent								
Mean	1.7808	2.0140	0.2333	**2.5907**	4.8817	6.4207	1.5390	**3.3391**
Median	0	1			1	2		
SD	4.8778	4.9020			12.7792	16.0127		
Initiated messages								
Mean	0.7357	0.8912	0.1554	**4.1160**	0.9168	1.2372	0.3204	**3.2550**
Median	0	0			2.2510	3.7248		
SD	2.0060	2.0925			2.2510	3.7248		
Initiated messages that led to match								
Mean	0.2498	0.2827	0.0329	**2.0178**	0.3182	0.4078	0.0896	**2.0684**
Median	0	0			0	0		
SD	0.8683	0.9024			0.9166	1.6784		
Obs.	5,752	6,055			1,911	2,087		

Notes. If the difference between the treatment and control group is significant at the 5% level, the t-statistics are in bold.

Table 2: Summary Statistics of User Activities Toward Likers

	Men				Women			
	control	treated	diff	t-stat	control	treated	diff	t-stat
	(1a)	(1b)		(1c)	(2a)	(2b)		(2c)
Number of browses								
Mean	238.2098	221.492	-16.7179	-1.4673	239.0277	257.3297	18.3019	0.8087
Median	47	41			27	23		
SD	629.8658	608.125			655.5649	764.9958		
Likes sent								
Mean	104.9727	89.08423	-15.8885	**-2.7703**	20.33062	20.07483	0.2558	0.0915
Median	18	16			3	3		
SD	332.4176	290.2234			83.0087	93.7785		
Initiated messages								
Mean	9.5570	9.5538	-0.0033	-0.0048	1.298796	1.251078		
Median	1	1			0	0		
SD	35.6085	38.7018			6.5044	5.6285		
Initiated messages that led to match								
Mean	0.6777	0.7207	0.0430	0.8456	0.2166	0.2348	0.0181	0.5678
Median	0	0			0	0		
SD	2.5462	2.9579			0.9584	1.0539		
Obs.	5,752	6,055			1,911	2,087		

Notes. If the difference between the treatment and control group is significant at the 5% level, the t-statistics are in bold.

Table 3: Summary Statistics of User Activities Toward NotLikers

	Matches Between m and Liker w					
	Control (CT)		Treatment (TR)			
	Mean	SD	Mean	SD	TR−CT	t-stat
Age	4.1144	4.7604	4.5310	4.0655	0.4166	**3.1464**
	[2,045]		[2,563]			
Education	1.7241	2.0473	1.9925	1.9539	0.2684	**2.0715**
	[534]		[435]			
Race	0.5706	0.0164	0.6136	0.0139	0.0430	**1.9961**
	[906]		[1,219]			
Body type	0.5706	0.0186	0.6266	0.0171	0.0559	**2.2121**
	[708]		[798]			
Popularity	0.7493	0.7701	0.8073	0.8051	0.0580	**2.4770**
	[2,045]		[2,563]			

Notes. Number of observations in square brackets. For race and body type, we report standard error and z-statistic. If the difference between the treatment and control group is significant at the 5% level, the t-statistics are in bold.

Table 4: Attributes Differences With Initiated Matches

for age, education level, popularity, race, and body type. Since race and body type are categorical variables, we construct attribute difference as an indicator variable that takes value 1 when m and w are of different race/body type, and 0 otherwise.

Table 4 displays mean attribute differences and standard deviation (standard error for race and body type) of couples *who matched with Likers*, separately for the control and treatment group. We also report the difference between the two groups (TR−CT) and the t-statistic (z-statistic for race and body type). Interestingly, attribute differences of the treatment group are significantly larger compared to those of the control group, across all dimensions. The age difference between couples in the treatment group ($\Delta = 4.53$) is on average 0.4 years (or 10.2%) greater than the age difference between couples in the control group ($\Delta = 4.11$); The difference in years of education between couples in the treatment group ($\Delta = 1.99$) is on average 0.27 years (or 15.7%) greater than that of couples in the control group ($\Delta = 1.72$); Approximately 61% of the users in the treatment group matched with partners of different race, which is 4 percentage points greater than that of the control group (57 percent); Approximately 62% of the users in the treatment group matched with partners of different body type, which is 5 percentage points greater than that of the control group (57%); The popularity difference between couples in the treatment group ($\Delta = 0.81$) is roughly 8% greater than the popularity difference between couples in the control group ($\Delta = 0.75$).

Tables 5 shows differences in degree of sorting between the treatment and control group when they *matched with NotLikers*. Since the focal user is not able to distinguish whether the NotLiker

	Matches Between m and NotLiker w					
	Control		Treatment			
	Mean	SD	Mean	SD	TR-CT	t-stat
Age	5.0063	4.8298	5.1376	4.9419	0.1314	1.2837
	[4,311]		[4,854]			
Education	1.8301	1.9869	1.8775	1.9728	0.0474	0.4761
	[777]		[808]			
Ethnicity	0.5775	0.0116	0.5631	0.0106	-0.0144	-0.9145
	[1,813]		[2,179]			
Body type	0.6388	0.0134)	0.5868	0.0127	-0.0521	**-2.8065**
	[1,282]		[1,498]			
Attractiveness	0.7183	0.6886	0.6950	0.6294	-0.0234	-1.6931
	[4,285]		[4,837]			

Notes. Number of observations in square brackets. For race and body type, we report standard error and z-statistic. If the difference between the treatment and control group is significant at the 5% level, the t-statistics are in bold.

Table 5: Attributes Differences With Initiated Matches

choose to not *like* the him/her, or whether they have not yet seen the his/her profile, the impact of treatment on assortative matching is ambiguous.

Our results provide evidence suggesting that the treatment reduces the degree of positive sorting between couples. One possible explanation for this is that the treatment encourages users to initiate a match with candidates that seem ex ante unreachable in the absence of treatment. If, generally, attribute difference is negatively correlated with ex ante match probability, users in the default control condition are more likely to initiate a match with a candidate whose attributes are similar to their own. In other words, users in the default control condition are more likely to initiate a match with a similar candidate because similar candidates are more likely to accept the match offer. However, a user who was discouraged from reaching out to a candidate in the control condition due to the (ex ante) low match probability may be encouraged to start a conversation in the treatment condition once he/she sees that the candidate had *liked* him/her. Therefore, in the treatment condition, users may be encouraged to initiate a conversation with Likers whose attributes are different from their own, despite ex-ante low match probability.

To test this, we need to show that treated users are more likely, than control users, to send messages to candidates who are different from themselves, and that these candidates have ex ante low match probability. As a starting point, we first show that users have lower probability of matching (ex ante match probability) with dissimilar candidates compared to similar candidates. To obtain ex ante match probabilities, we use predicted values of a logistic regression where we regress a match indicator (exchanging at least four messages, with each sending at least two messages) on candidates' characteristics, as well as positive and negative differences between user and candidates' characteristics.[16] Table 6 reports the correlation between attribute difference and ex ante match probability. Except for the body type, attribute difference is negatively correlated with ex ante match probability, suggesting that indeed a match is generally more likely to occur if the pairs are similar in their characteristics.

Having shown that users have ex ante lower match probability with candidates whose attributes are different from their own compared to similar candidates, we now proceed to show that treated users are more likely, than control users, to match with candidates that have ex-ante low match probability. To do so, we need to compare the differences in ex-post match probability between the treatment and control group for different values of ex ante match probabilities. Ex post match probability is calculated as the average value of the match indicator for each percentile of ex ante match probability. To obtain the difference in ex post match probability, we subtract the control group's mean probability of a match with a Liker from that of the treatment group.

Figure 4 (a) plots the distribution of the difference (between the treatment and control group) in ex post match probability with Likers against its ex ante match probability. Specifically, the difference in ex post match probability is obtained by subtracting the control group's mean probability of a match with a Liker from that of the treatment group, separately for each percentile of ex ante match probability. We then plot the mean difference in ex post match probability for each decile of ex ante match probability. The positive value of this difference implies that the treatment group obtained more matches compared to the control group, and vice versa. We find that for sufficiently low values of ex ante match probabilities (e.g., 1st and 2nd decile), differences in match probabilities are positive. Positive (negative) differences for low (high) values of ex ante match probabilities suggest that users in the treatment group are more (less) likely, compared to the control group, to match with Likers of low (high) ex ante match probabilities. We repeat this analysis for matches formed with NotLikers. Results are plotted in Figure 4 (b). Differences in match probabilities are approximately centered around zero, for all values of ex-ante match probabilities.

We have shown that the treatment leads to less sorting when matches are formed with Likers, suggesting that reducing frictions can lead to less sorting between couples. Although the experiment reduced treatment group's search friction by providing the information about who *liked* the users

[16] In other words, we regress a binary match variable on right hand-side variables of equation 11 which specifies users' utility from matching with a candidate.

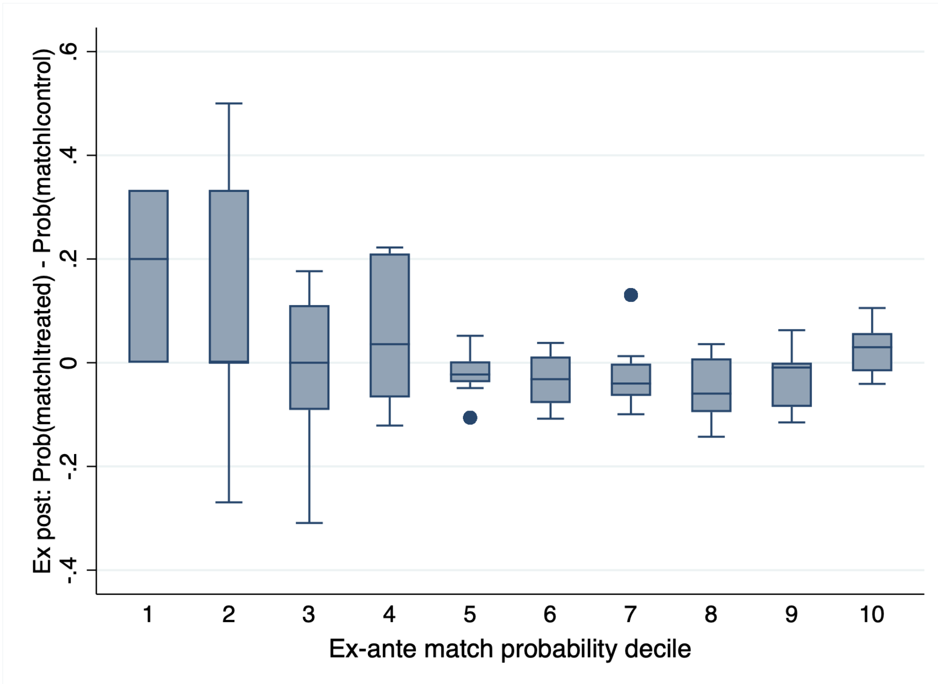

(a) Differences in Probability of Matching With a Liker

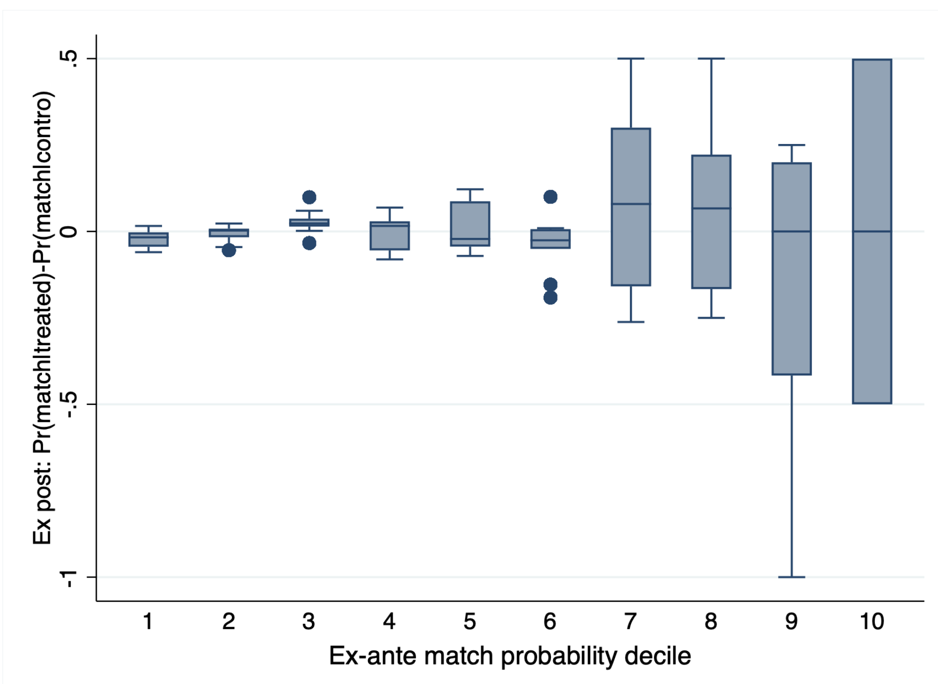

(b) Differences in Probability of Matching With a NotLiker

Notes. Predicted match probabilities are converted into percentiles.

Figure 4: Ex Post Probability of Matching With Respect to Ex Ante Match Probabilities

	Age	Education	Race	Body Type	Popularity
correlation	-0.1170	-0.1711	-0.1041	0.0851	-0.0551
stderr	(0.0024)	(0.0058)	(0.0041)	(0.0053)	(0.0040)

Table 6: Pearson Correlation of Attribute Gap and Ex Ante Match Probability

without users having to *ℓike* first to find out, some uncertainty about whether a match will happen nevertheless remains for the treatment group because they may still get no response after sending a message to a candidate despite being *ℓiked*. In addition, if the cost of composing a message is nonnegligible, the decision to start a conversation with a candidate would also depend on the perceived match probability. To disentangle the impact of frictions and preferences on sorting and user welfare, we need to compare the matches formed in a market with frictions to those formed in a frictionless environment where only preferences shape the matching outcomes. Therefore, in the following section, we develop a structural model of costly search and use the model estimates to simulate equilibrium matches in a frictionless environment.

7 The Search Model

We consider an online dating platform where in each period, N_M men and N_W women are searching for a partner. Time is discrete, and we assume that discounting across time is negligible, i.e., the time discount factor $\rho \approx 1$. This assumption is consistent with existing research that uses data from online dating platforms (Hitsch et al. (2010a); Fong (2018)). Each man is indexed by $m \in \mathcal{M} = \{1, 2, ..., N_M\}$ and has observed characteristics X_m. Similarly, each woman is indexed by $w \in \mathcal{W} = \{1, 2, ..., N_W\}$ and has observed characteristics X_w. In what follows, we describe the model from the perspective of a male user m. The model is symmetric for a female user w.

In each period, female candidate's profiles are drawn randomly from a distribution F_W and displayed to man m as long as he continues the search for a partner. A match occurs if both m and w agree to go on an offline date. If at least one does not agree, the search for a partner continues. Following the existing literature, we assume that the distribution of single users' profiles is exogenously given and is stationary over time. To guarantee stationarity, we assume that users who are matched exit the market and are immediately replaced by their "clones," as in McNamara and Collins (1990), Burdett and Coles (1997), Bloch and Ryder (2000), and Adachi (2003).[17]

[17] A clone of a man m has identical characteristics as man m.

For each candidate w that m browses, he has to pay a cost $c_{M,m}^{browse}$, which can be interpreted as time and cognitive effort spent on viewing the profile. Then, for each w's profile, m first needs to choose action $d_{mw} \in \{\text{like, nlike}\}$, which indexes his decision to ℓike or not ℓike w's profile ($d_{mw} =$ like if ℓike, and $d_{mw} =$ nlike if not ℓike). Then, he needs to choose action $\mu_{mw} \in \{\text{msg, nmsg}\}$, which indexes his decision whether or not to send a message to w ($\mu_{mw} =$ msg if message, and $\mu_{mw} =$ nmsg if not message). Let us denote these two decision stages as the *liking stage* (LS) and the *messaging stage* (MS), respectively.

There are two types of users, treated ($\tau_m = 1$) and control ($\tau_m = 0$).[18] At the liking stage, in contrast to treated users who know the true value of ℓ_{wm} which equals 1 if w had $\ell iked$ m and 0 otherwise, users in the control group do not observe the true value of ℓ_{wm} unless they choose $d_{mw} =$ like at the liking stage. Therefore, control users form an expectation about ℓ_{wm} conditional on their own and w's characteristics. The information structure of m depends on his type τ_m, the stage he is in $s \in \{LS, MS\}$, and his action $d_{mw} \in \{\text{like, nlike}\}$ at the liking stage.

Denote $\mathscr{L}_{wm}^{s,\tau_m}$ as the expected value of ℓ_{wm} that user of type τ_m has about ℓ_{wm} at stage $s \in \{LS, MS\}$. $\mathscr{L}_{wm}^{LS,\tau_m}$ is summarized as

$$\mathscr{L}_{wm}^{LS,\tau_m} = \begin{cases} \ell_{wm} & \text{if } \tau_m = 1 \\ \mathbb{E}[\ell_{wm}|X_m, X_w] & \text{if } \tau_m = 0. \end{cases} \tag{1}$$

The exact value of ℓ_{wm} will be revealed to control users if they choose $d_{mw} =$ like at the liking stage. Therefore, at the messaging stage, all users *except* those in the control group who have chosen to not ℓike a profile will observe the true value of ℓ_{wm}. $\mathscr{L}_{wm}^{MS,\tau_m}$ can then be summarized as

$$\mathscr{L}_{wm}^{MS,\tau_m} = \begin{cases} \ell_{wm} & \text{if } \tau_m = 1 \\ \ell_{wm} & \text{if } \tau_m = 0 \quad \& \quad d_{wm} = \text{like} \\ \mathbb{E}[\ell_{wm}|X_m, X_w] & \text{otherwise.} \end{cases} \tag{2}$$

Denote Ω_{mw}^{s,τ_m} as the information that m of type τ_m has about w at stage $s \in \{LS, MS\}$. The information set of m about w at the liking stage is given as:

$$\Omega_{mw}^{LS,\tau_m} = \{X_w, e_{mw}, \mathscr{L}_{wm}^{LS,\tau_m}\} \tag{3}$$

[18] The reason why we distinguish control and treated users in our model is because we need to estimate reservation values of both types of users for our counterfactual analysis.

and the information set at the messaging stage is given as follows:

$$\Omega_{mw}^{MS,\tau_m} = \{X_w, \mathscr{L}_{wm}^{MS,\tau_m}, e_{mw}\}. \tag{4}$$

In what follows, we first describe the details of the *messaging stage,* then we describe the details of the *liking stage* in a backwards induction manner.

7.1 Messaging Stage

Let $V_M(m)$ denote m's expected utility of remaining single and continuing the search for a partner. For the moment, let us suppose that $V_M(m)$ is given. Man m receives a utility of $u_{M,mw}$ by matching with a woman w. Denote $V_{M,mw}^{MS}(\mu_{mw}|d_{mw}, \Omega_{mw}^{MS,\tau_m})$ as m's expected utility from choosing action μ_{mw} conditional on his choice $d_{mw} \in \{\text{like, nlike}\}$ at the liking stage and his information set Ω_{mw}^{MS,τ_m}. Then we have

$$V_{M,mw}^{MS}(\mu_{mw}|d_{mw}, \Omega_{mw}^{MS,\tau_m}) = \begin{cases} -c_{M,mw}^{\text{msg}} + u_{M,mw} \cdot \mathbb{E}_m \left[P_{W,mw}(d_{mw}, \mu_{mw}|X_m, X_w, \ell_{wm}, \Omega_{mw}^{MS,\tau_m})\right] \\ \qquad +V_M(m) \cdot \Big(1 - \mathbb{E}_m \left[P_{M,mw}(d_{mw}, \mu_{mw}|X_m, X_w, \ell_{wm}, \Omega_{mw}^{MS,\tau_m})\right]\Big) & \text{if } \mu_{mw} = \text{msg} \\ \\ u_{M,mw} \cdot \mathbb{E}_m \left[P_{W,mw}(d_{mw}, \mu_{mw}|X_m, X_w, \ell_{wm}, \Omega_{mw}^{MS,\tau_m})\right] \\ \qquad +V_M(m) \cdot \Big(1 - \mathbb{E}_m \left[P_{M,mw}(d_{mw}, \mu_{mw}|X_m, X_w, \ell_{wm}, \Omega_{mw}^{MS,\tau_m})\right]\Big) & \text{otherwise} \end{cases} \tag{5}$$

If m decides to send a message, he has to pay a cost $c_{M,mw}^{\text{msg}}$, which can be interpreted as (1) time and effort to compose a message, or (2) aversion to experiencing negative emotions in case w does not respond or agree to go on a date (match). $P_{M,mw}(d_{mw}, \mu_{mw}|X_m, X_w, \ell_{wm}, \Omega_{mw}^{MS,\tau_m})$ is the probability that w will match with m, and is a function of d_{mw} and μ_{mw} conditional on X_m, X_w, ℓ_{wm}, and Ω_{mw}^{MS,τ_m}. When m does not know the true value of ℓ_{wm}, he forms an expectation about the match probability conditional on the information Ω_{mw}^{MS,τ_m} that he knows about w.

If m sends a message, he incurs a cost $c_{M,mw}^{\text{msg}}$ and gets utility $u_{M,mw}$ with expected probability $\mathbb{E}_m \Big[P_{M,mw}(d_{mw}, \text{msg}|X_m, X_w, \ell_{wm}, \Omega_{mw}^{MS,\tau_m})\Big]$, but with probability $1 - \mathbb{E}_m \Big[P_{M,mw}(d_{mw}, \text{msg}|X_m, X_w, \ell_{wm}, \Omega_{mw}^{MS,\tau_m})\Big]$ he does not get a match and remains single and continues the search. On the other hand, even if m does not send a message, he may nevertheless get a match with positive expected probability $\mathbb{E}_m \Big[P_{M,mw}(d_{mw}, \text{nmsg}|X_m, X_w, \ell_{wm}, \Omega_{mw}^{MS,\tau_m})\Big]$ if $d_{mw} = \text{like}$. This is because, if both m and w *like* each other, the information that m *ℓiked* w is revealed to w, which affects w's decision to initiate a message to m. We assume that if m neither *ℓikes* nor messages w, the match probability is zero, i.e., $\mathbb{E}_m\left[P_{M,mw}(\text{nlike}, \text{nmsg}|X_m, X_w, \ell_{wm}, \Omega_{mw}^{MS,\tau_m})\right] = 0$. Moreover, since sending a message is a stronger signal of interest than simply *ℓiking*, we assume that as long as m sends a message to w, his decision at the liking stage becomes irrelevant, i.e., $\mathbb{E}_m\left[P_{M,mw}(\text{like}, \text{msg}|X_m, X_w, \ell_{wm}, \Omega_{mw}^{MS,\tau_m})\right] = \mathbb{E}_m\left[P_{M,mw}(\text{nlike}, \text{msg}|X_m, X_w, \ell_{wm}, \Omega_{mw}^{MS,\tau_m})\right]$. Then, m will choose to message w if and only if

$$V_{M,mw}^{MS}(\mu_{mw} = \text{msg}|d_{mw}, \Omega_{mw}^{MS,\tau_m}) \geq V_{M,mw}^{MS}(\mu_{mw}{=}\text{nmsg}|d_{mw}, \Omega_{mw}^{MS,\tau_m}) \tag{6}$$

7.2 Liking Stage

At the liking stage, if m chooses $d_{mw} = \text{like}$, he receives utility $c_{M,mw}^{\text{like}} + \varepsilon_{mw}^{\text{like}}$, where $c_{M,mw}^{\text{like}}$ is the psychological cost of *ℓiking* and $\varepsilon_{mw}^{\text{like}}$ is an error term observed by m (but unobserved by the researcher) that affects m's decision to *ℓike* w.[19] Similarly, if m does not *ℓike* w, he receives $\varepsilon_{mw}^{\text{nlike}}$. Error terms $\varepsilon_{mw}^{\text{like}}$ and $\varepsilon_{mw}^{\text{nlike}}$ are distributed i.i.d Type I EV, denoted as $G(\varepsilon)$. We assume that if $c_{M,mw}^{\text{like}} = 0$, both error terms $\varepsilon_{mw}^{\text{like}}$ and $\varepsilon_{mw}^{\text{nlike}}$ equal zero and that m always chooses $d_{mw} = \text{like}$. This assumption will be useful later, when we show the connection between the outcomes of our model and those of the Gale-Shapley centralized model, because when $c_{M,mw}^{\text{msg}} = 0$, one can see with a bit of algebra that the condition in equation 6 reduces to $u_{M,mw} \geq V_M(m)$, in which case m's decision to *ℓike* does not affect the current-stage utility nor his decision at the messaging stage.

The choice-specific expected utility at the liking stage from choosing d_{mw} when the information available about w is Ω_{mw}^{LS,τ_m} can be written as

$$V_{M,mw}^{LS}(d_{mw}|\Omega_{mw}^{LS,\tau_m}) = \begin{cases} -c_{M,mw}^{\text{like}} + \varepsilon_{mw}^{\text{like}} + E_m\left\{\max_{\mu_{mw}}\left[V_{M,mw}^{MS}(\mu_{mw}|d_{mw} = \text{like}, \Omega_{mw}^{LS,\tau_m})\right]\right\} & \text{if } d_{mw} = \text{like} \\ \varepsilon_{mw}^{\text{nlike}} + E_m\left\{\max_{\mu_{mw}}\left[V_{M,mw}^{MS}(\mu_{mw}|d_{mw} = \text{nlike}, \Omega_{mw}^{LS,\tau_m})\right]\right\} & \text{if } d_{mw} = \text{nlike}. \end{cases} \tag{7}$$

m will choose $d_{mw} = \text{like}$ if and only if

$$V_{M,mw}^{LS}(d_{mw}{=}like|\Omega_{mw}^{LS,\tau_m}) \geq V_{M,mw}^{LS}(d_{mw}{=}nlike|\Omega_{mw}^{LS,\tau_m}) \tag{8}$$

[19] As explained earlier, the psychological cost of *ℓiking* can be incurred due to several reasons, such as (i) not finding w sufficiently attractive, (ii) prospect of being rejected by w, or (iii) the possibility of having to reject w's overtures after triggering her response by *ℓiking* her.

Let $V^{LS}_{M,mw}(d_{mw}|\Omega^{LS,\tau_m}_{mw}) = \bar{V}^{LS}_{M,mw}(d_{mw}|\Omega^{LS,\tau_m}_{mw}) + e^{d_{mw}}_{mw}$. Since $\varepsilon^{\text{like}}_{mw}$ and $\varepsilon^{\text{nlike}}_{mw}$ follow i.i.d Type I EV distribution, the probability of choosing $d_{mw} = $ like is

$$\Pr(d_{mw} = \text{like}) = \frac{\exp\Big(\bar{V}^{LS}_{M,mw}(\text{like}) - \bar{V}^{LS}_{M,mw}(\text{nlike})\Big)}{1 + \exp\Big(\bar{V}^{LS}_{M,mw}(\text{like}) - \bar{V}^{LS}_{M,mw}(\text{nlike})\Big)} \tag{9}$$

7.3 Latent Utility

We assume that preferences for potential partners depend on observed own and candidate's attributes, and idiosyncratic preference shock e_{mw}, which follows i.i.d logistic distribution:

$$u_{M,mw} = u_M(X_m, X_w; \Theta_M) + e_{mw} \tag{10}$$

where $X_M = (x_m, x^d_m)$ and $X_W = (x_w, x^d_w)$ are $m's$ and w's observed characteristics which follow a distribution F_M and F_W, respectively. x_m and x_w are vectors that have continuous values, and x^d_m and x^d_w are sets of categorical variables. Θ_M is a vector of parameters that represent $m's$ preferences.

The latent utility that m gets if he matches with w is parameterized as

$$\begin{aligned} u_M(X_m, X_w; \Theta_M) = \quad & x'_w\beta_M + (|x_w - x_m|'_+)\beta^+_M + (|x_w - x_m|'_-)\beta^-_M \\ & + \textstyle\sum_{r,s=1}^{N} \mathbb{1}\{x^d_{mr} = 1 \text{ and } x^d_{ws} = 1\} \cdot \beta^d_{M,rs} \end{aligned} \tag{11}$$

where $|x_w - x_m|_+ = \max(x_w - x_m, 0)$ and $|x_w - x_m|_- = \max(x_m - x_w, 0)$. In other words,$|x_w - x_m|_+$ is the difference between m and w's attributes if this difference is positive, and $|x_w - x_m|_-$ is the absolute value of this difference if this difference is negative. More formally, $|x_w - x_m|_+ = \max(x_w - x_m, 0)$ and $|x_w - x_m|_- = \max(x_m - x_w, 0)$. Dummy variables x^d_{mr} and x^d_{wr} indicate whether m and w have certain categorial characteristics. For example, $x^d_{m,\text{Asian}} = 1$ if m is Asian, and 0 otherwise. The set of preference parameters to be estimated is $\Theta_M = (\beta_M, \beta^+_M, \beta^-_M, \beta^d_M)$ for men and $\Theta_W = (\beta_W, \beta^+_W, \beta^-_W, \beta^d_W)$ for women.

7.4 Negligible Costs and Equivalence With Gale-Shapley Stable Matches

In this subsection, we provide a link between the outcomes of a decentralized search model and stable matches of the Gale-Shapley marriage model. The Gale-Shapley marriage problem assumes the presence of a central matchmaker which recommends a matching to agents given individuals' preferences over potential partners, and hence does not describe the online matching process wherein agents have to incur costs to find a partner in the absence of a central matchmaker. Adachi (2003) shows, however, that as search costs become negligible, the set of equilibrium matches obtained in a two-sided search and matching model is identical to the set of stable equilibrium matches predicted by the Gale-Shapley algorithm. Moreover, repeated rounds of offer-making and corresponding rejections of the deferred-acceptance algorithm resemble the search and messaging behavior of the users on online dating platforms (Hitsch et al. (2010a)). Since the stable matching predicted by the Gale-Shapley algorithm is also Pareto-optimal, it also provides a theoretical efficiency benchmark that can be used to quantify the loss of efficiency caused by frictions in our platform.

Consider a hypothetical condition where all costs are removed, i.e. $c^{\text{like}}_{M,mw} = c^{\text{msg}}_{M,mw} = 0$, and consider m's decision problem of whether to send a message to w. If $c^{\text{msg}}_{M,mw} = 0$, the condition in equation 6 reduces to

$$u_{M,mw} \geq V_M(m) \tag{12}$$

The match probability no longer appears in equation 12, and m sends a message as long as the utility from a match is greater than the expected value of continuing the search. Since m's decision at the *ℓiking* stage does not affect his decision at the messaging stage, m's decision to *ℓike* will only depend on the error terms $\varepsilon^{\text{like}}_{mw}$ and $\varepsilon^{\text{nlike}}_{mw}$. To make a connection between the outcomes of our model to those of the Gale-Shapley centralized model, we will make a simplifying assumption that when $c^{\text{like}}_{M,mw} = c^{\text{msg}}_{M,mw} = 0$, both error terms $\varepsilon^{\text{like}}_{mw}$ and $\varepsilon^{\text{nlike}}_{mw}$ equal zero. Then since the decision to *ℓike* does not affect the utility, we will assume that m always chooses $d_{mw} = \text{nlike}$. In this case, m's expected utility of staying single and continuing the search can be defined by the following Bellman's equation:

$$\begin{aligned} V_M(m) &= V^{LS}_{M,mw}(d_{mw} = \text{like}|\Omega^{LS,\tau_m}_{mw}) \\ &= \rho \int u_{M,mw} \cdot A_M(m,w)A_W(m,w) + V_M(m)\left(1 - A_M(m,w)A_W(m,w)\right) dF_W(w \end{aligned} \tag{13}$$

where

$$A_W(m,w) = \mathbb{I}\{u_{W,mw} \geq V_W(m)\} \qquad \text{and} \qquad A_M(m,w) = \mathbb{I}\{u_{M,mw} \geq V_M(m)\} \tag{14}$$

Similarly, woman w's expected utility of staying single can be written as

$$V_W(w) = \rho \int u_{W,wm} \cdot A_W(m,w)A_M(m,w) + V_W(w)\left(1 - A_W(m,w)A_M(m,w)\right) dF_M(m) \tag{15}$$

The system of equations 13 and 15 define a monotone iterative mapping that converges to a profile of reservation utilities $\left(V_M^{GS}(m), V_W^{GS}(w)\right)$ solving the system, and hence characterize the stationary equilibrium in this market.

Adachi (2003) shows that as $\rho \to 1$, the set of equilibrium outcomes in a decentralized search model reduces to the set of stable matchings in a corresponding Gale-Shapley marriage problem. A stable match is defined, following Gale and Shapley (1962), as a pairing where there are no pairs (m, w) who are willing to abandon their partners and match with each other. Specifically, Adachi (2003) shows that the system of Bellman equations 13 and 15 coincide with the following system of equations characterizing the set of stable matchings in a Gale-Shapley marriage problem:

$$\begin{aligned} V_M^{GS}(m) &= \max_{\mathcal{W}\cup\{m\}} \left\{u_{M,mw} | u_{W,wm} \geq V_W^{GS}(w)\right\} \\ V_W^{GS}(w) &= \max_{\mathcal{M}\cup\{w\}} \left\{u_{W,wm} | u_{M,mw} \geq V_M^{GS}(m)\right\} \end{aligned} \tag{16}$$

where $V_M^{GS}(m)$ and $V_W^{GS}(w)$ is the expected utility of staying single in a frictionless environment for man m and woman w, respectively. If time is not discounted, and if there are no costs (costs of browsing, as well as the costs of $\ell iking$ and messaging) each man (woman) continues the search process until he (she) finds a woman (man) such that $u_{M,mw} \geq V_M^{GS}(m)$ ($u_{W,wm} \geq V_W^{GS}(w)$). Then a man will be matched with the best woman who is willing to match with him, and vice versa, which is how the set of stable matchings are characterized in Gale-Shapley problem.

8 Estimation

8.1 Match Probability

We obtain users' expectations about the match probability for each stage separately from the model. There are two match probabilities in our model, (i) $P_{W,mw}(d_{mw} = \ell ike, \mu_{mw} = \text{nmsg}|X_m, X_w, \ell_{wm})$ which is a match probability from sending only a ℓike, and (ii) $P_{W,mw}(d_{mw}, \mu_{mw} = \text{msg}|X_m, X_w, \ell_{wm})$ which is a match probability from sending a message.[20]

To obtain $\mathbb{E}_m\left[P_{W,mw}(d_{mw} = \ell ike, \mu_{mw} = \text{nmsg}|X_m, X_w, \ell_{wm})|\Omega_{mw}^{LS,\tau_m}\right]$, we use the predicted values from estimating the following linear probability model:

$$Match_{wm}^{\text{like}} = u_W(X_m, X_w; \Theta_W) + \psi^{\text{like}} \mathscr{L}_{wm}^{LS,\tau_m} + \xi_m + e_{wm} \tag{17}$$

[20] Recall that we assumed $\mathbb{E}_m\left[P_{M,mw}(\text{like}, \text{msg}|X_m, X_w, \ell_{wm})|\Omega_{mw}^{MS,\tau_m}\right] = \mathbb{E}_m\left[P_{M,mw}(\text{nlike}, \text{msg}|X_m, X_w, \ell_{wm})|\Omega_{mw}^{MS,\tau_m}\right]$

$Match^{\text{like}}_{wm}$ as a binary variable that equals 1 if m and w match (exchange at least 4 messages) when m *ℓikes* w but does not send a message to w, and 0 otherwise.[21] Here $u_W(\cdot)$ is defined similarly as in equation 11. The parameter ψ^{like} measures the impact of w's *ℓike* on the match probability.

The endogeneity of $\mathscr{L}^{LS,\tau_m}_{wm}$ is a potential concern. It is possible that unobserved high compatibility between m and w leads to a higher likelihood of w *ℓiking* m, leading to a biased estimate of ψ^{like}. Some users might be inherently more compatible with others, resulting in higher number of *ℓikes*. We address this endogeneity problem by including user fixed effects, ξ_m.

Similarly, to obtain $\mathbb{E}_m[P_{W,mw}(d_{mw}, \mu_{mw} = \text{msg}|X_m, X_w, \ell_{wm})|\Omega^{MS,\tau_m}_{mw}]$, let $Match^{\text{msg}}_{wm}$ be a binary variable that equals 1 if m and w match after m sends a message to w, and 0 otherwise. We estimate the following linear probability model and use predicted values from this regression as $\mathbb{E}_m[P_{W,mw}(d_{mw}, \mu_{mw} = \text{msg}|X_m, X_w, \ell_{wm})|\Omega^{MS,\tau_m}_{mw}]$:

$$Match^{\text{msg}}_{wm} = u_W(X_m, X_w; \Theta_W) + \psi^{\text{msg}}\mathscr{L}^{MS,\tau_m}_{wm} + \eta_m + e_{wm} \tag{18}$$

where η_m is user fixed effect. Parameter ψ^{msg} measures how much w is more likely to match with m in response to his message if she has *ℓiked* him. As before, we control for unobserved user compatibility using fixed effects.

In our field experiment, the experiment reduces the search friction for the users in the treatment group, but it does not reduce the search friction for the correspondent users who have interacted with the treatment group. Therefore, the experiment reduces the search friction for only one side of the market (the treatment group of our experimental users). To explore how a different platform design that reduces search frictions affects user welfare, we need to simulate equilibrium matches in a counterfactual scenario where users on *both* sides of the market are treated. When users on *both* sides of the market are treated, perceived match probabilities depend not only on the focal users' treatment, but also on the candidate's treatment. This is because a treated candidate can directly observe whether a focal user *ℓikes* him/her, and this will in turn affect the match probability.[22]

In our experiment, only 0.53 percent of the candidates are part of the experiment (0.26 percent control and 0.27 percent treated), making it is difficult to estimate match probabilities for the case where both sides of the market are treated. However, since we observe *ℓikes* and messages that treated users *received* from the other side, we are able to estimate equations 17 and 18 from

[21] Note that here we index the subscript as wm as opposed to mw. This is to reflect w's preferences and decision to accept m's offer

[22] If the focal user chooses to *ℓike* the candidate but decides not to send a message, the treated candidate will be able to observe that the focal user had *ℓiked* him/her without having to take a costly action.

the correspondent users' point of view: The probability of getting a match when $\ell iking$ and/or messaging a treated user. In other words, since we observe $\ell iking$ and message initiating activities of correspondent users, we are able to estimate equations 17 and 18 using the data in which the correspondent users (regardless of their treatment status) ℓike and initiate messages to treated experimental users.[23]

8.2 Reservation Value

In this subsection we describe how we estimate $V_M(m)$ and $V_W(w)$ in the presence of frictions. We will describe how to obtain reservation values in the absence of frictions, i.e., $V_M^{GS}(m)$ and $V_W^{GS}(w)$, in Section 10.1.

We assume that $V_M(m)$ and $V_W(w)$ remain constant across different profiles. Ideally, we want to estimate these reservation values using user-specific fixed effects, following Hitsch et al. (2010a) and Banerjee et al. (2013). But not only is this approach computationally burdensome due to the large number of users in our sample, it is also unsuitable in our setting due to the selection issue: to include user-specific fixed effects, we need to drop users who haven't $\ell iked$ or messaged any profiles. However, dropping these individuals may lead to biased estimates because higher costs of $\ell iking$ and messaging may prevent certain users from $\ell iking$ and messaging. Therefore, we group users with similar observed characteristics and estimate their reservation values using group fixed effects instead of individual user fixed effects, as users with similar characteristics are likely to have similar reservation values. Using this approach allows us to keep all users in the data even if they haven't $\ell iked$ or messaged any profiles. Specifically, we use K-means clustering to partition users into groups. K-means is an unsupervised learning approach that partitions the dataset into K predefined nonoverlapping clusters. Each user is assigned to a cluster such that the sum of the squared distance between the users' characteristics and the cluster's centroid (arithmetic mean of all users' characteristics that belong to that cluster) is at the minimum. We classify users into five clusters based on their observed characteristics X_m, separately for each type τ_m (total 10 clusters).

[23]We estimate equations 17 and 18 from the correspondent users' perspectives by including correspondent user fixed effects. To obtain predicted values of match probabilities for the experimental users, however, we need estimates of experimental user fixed effects. Since we do not have estimates of experimental user fixed effects for the case when both sides are treated, we add mean correspondent user fixed effects to the predicted match probabilities.

8.3 Likelihood

We maximize the joint likelihood of m's decision at the liking and messaging stages for each profile that he browses. The likelihood of our model is given as

$$L = \prod_{m=1}^{N_M} \prod_{w=1}^{J_m} \left(\Pr_{d,mw} \cdot \Pr_{\mu,\, mw|d,mw}^{\delta_{mw}} \left(1 - \Pr_{\mu,\, mw|d,mw}\right)^{1-\delta_{mw}} \right)^{\vartheta_{mw}} \times \left((1 - \Pr_{mwd}) \cdot \Pr_{\mu,\, mw|d,mw}^{\delta_{mw}} \left(1 - \Pr_{\mu,\, mw|d,mw}\right)^{1-\delta_{mw}} \right)^{1-\vartheta_{mw}} \tag{19}$$

where $\Pr_{d,mw}$ is the probability that m *ℓikes* w, $\Pr_{\mu,mw}$ is the probability that m messages w, and J_m is the total number of profiles that m browses. ϑ_{mw} indicates the decision made at the liking stage ($\vartheta_{mw} = 1$ if $d_{mw} =$ like, $\vartheta_{mw} = 0$ otherwise), and δ_{mw} indicates the decision chosen at the message stage ($\delta_{mw} = 1$ if $\mu_{mw} =$ msg, $\delta_{mw} = 0$ otherwise).

8.4 Identification

Identification is difficult in models where costs and preferences are interdependent. This is because *ℓiking* w could be due to high preferences for w's characteristics, or it could be due to the low cost of *ℓiking*. Similarly, messaging w could be due to high preference for w's characteristics or could be due to the lost cost of messaging. Thus, we rely on an exclusion restriction to separately identify preferences from costs. When we choose different sets of covariates to enter the utility and the cost function, covariates that enter the cost function (but not the utility function) serve as an exclusion restriction for identification.

In our model, ℓ_{wm} (or its expectation) enters the expected utility net of costs ($c^{\text{like}}_{M,mw}$ and $c^{\text{msg}}_{M,mw}$) only through its impact on match probability. Since the predicted match probability is estimated separately from the model as described in Subsection 8.1, ℓ_{wm} does not enter the expected utility function and hence can serve as an exclusion restriction.

As described earlier, the information set of a user depends on the type of the user, his actions, and the stage that he is in. We parameterize the cost of *ℓiking* as

$$c^{\text{like}}_{M,mw} = \exp\left(\lambda_0 + \lambda_1 \mathscr{L}^{LS,\tau_m}_{wm} + \lambda_2 \text{age}_m + \lambda_3 \text{popularity}_m\right).$$

We let the cost depends on the user's age and his overall popularity. Here $\mathscr{L}^{LS,\tau_m}_{wm}$ serves as an exclusion restriction.

Similarly, the expected cost of messaging is parametrized as

$$c^{\text{msg}}_{M,mw} = \exp\left(\gamma_0 + \gamma_1 \mathscr{L}^{s,\tau_m}_{wm} + \gamma_2 \text{age}_m + \gamma_3 \text{popularity}_m\right).$$

9 Estimation Results

Table 7 reports the maximum likelihood estimates of preference parameters, separately for men and women. Our estimation results are, generally, consistent with the findings of prior research (Kurzban and Weeden (2005); Fisman et al. (2006); Hitsch et al. (2010a); Hitsch et al. (2010b)). Both genders prefer partners who are young, and while men tend to have greater preference for women whose age is similar to their own, women prefer men who are older than them. With regards to education level, men's preferences are opposite of women's. Men prefer women with less years of education, while women prefer men with more years of education. Nevertheless, both genders tend to prefer candidates whose education level is greater than their own. We find that both genders prefer candidates with average body types to those who are heavy and overweight. Men show a strong preference for women with an average body type, while women show a strong distaste for men who are overweight. Moreover, men prefer being bigger than their female partners. We also find that popularity is an important determinant of preference for both genders. Both men and women prefer more popular candidates, but women tend to place greater emphasis on popularity than men. While men prefer women whose popularity is similar to their own, women prefer partners who are more popular than themselves. Finally, both men and women generally have a relative distaste for a candidate of a different race, but most of the coefficients on women's race preference are not statistically significant.

Table 8 Panel A reports parameter estimates of costs of *ℓiking*. Cost of *ℓiking* is significant and has important implications for user's search behavior. As expected, for both men and women, the coefficient on $\mathscr{L}^{LS,\tau_m}_{wm}$ is negative. This suggests that users are more likely to *ℓike* a candidate who is also more likely to *ℓike* (or has *ℓiked*) them. Age and popularity, on the other hand, has opposite effect for men and women. Older age and greater popularity increases the cost of *ℓiking* for men, but reduces the cost of *ℓiking* for women. Table 8 Panel B reports parameter estimates of costs of messaging. The coefficient on $\mathscr{L}^{s,\tau_m}_{wm}$, similarly to the cost of *ℓiking*, is negative for both genders. For both genders, the cost of messaging is significant, and the constant term is greater than that of the cost of *ℓiking*. For both men and women, greater age and popularity lowers the cost of messaging.

	Preferences of Men		Preferences of Women	
	Coefficients	SE	Coefficients	SE
Age	-0.1026***	0.0125	-0.6888***	0.0542
Age Difference (+)	-0.0607***	0.0237	-0.4974***	0.0426
Age Difference (−)	-0.0566	0.0396	0.7591***	0.0585
HighSchool	0.2151**	0.1096	0.1288	0.9534
TwoYear	0.2762**	0.1152	-0.5880	0.4472
University	-0.1532**	0.0732	-0.4773***	0.1834
Masters	-0.0695**	0.0333	-0.0122	0.8894
Law	-0.0451**	0.0213	0.6920***	0.2551
Medical	-0.0230*	0.0137	0.7965***	0.1924
PhD	-0.0130	0.0144	0.7666***	0.2897
Education Difference (+)	-0.5457***	0.0475	-0.4313**	0.1798
Education Difference (−)	0.2964***	0.0810	0.1927*	0.0998
Skinny	0.1691***	0.0585	-0.8016	0.4867
Average	0.3576***	0.1398	0.0294	0.3091
Heavier	0.0118	0.0221	-0.8375	0.5560
Overweight	0.0083	0.0245	-1.1515***	0.1654
BMI Difference (+)	0.1837**	0.0832	0.6543	0.8054
BMI Difference (−)	-0.2120***	0.0750	-0.1993	0.2103
Popularity	1.0819***	0.0505	2.9504***	0.2879
Popularity Difference (+)	-0.9354***	0.0626	-7.8175***	0.7120
Popularity Difference (−)	-0.1520***	0.0540	1.8292***	0.2577
Asian; mate White	-0.0563**	0.0239	1.5670	1.2275
Asian; mate Black	-0.0178*	0.0101	-0.0242	0.1287
Asian; mate Hispanic	-0.0174*	0.0101	0.0993	0.1426
Asian; mate other	-0.0073	0.0113	0.0437	0.1278
White; mate Asian	-0.0896**	0.0389	-0.5701	0.4524
White; mate Black	-0.1091**	0.0479	-1.0030	0.8148
White; mate Hispanic	-0.0495*	0.0296	-0.8049	0.6340
White; mate other	-0.0013	0.0111	-0.1335	0.2121
Black; mate Asian	0.0042	0.0117	-0.2764	0.2187
Black; mate White	0.0198	0.0165	-2.3525	1.7018
Black; mate Hispanic	0.0191	0.0143	-0.5329	0.3987
Black; mate other	0.0120	0.0091	-0.0912	0.1157
Hispanic; mate Asian	-0.0433**	0.0219	-0.1337	0.2044
Hispanic; mate White	-0.0110	0.0137	-0.1406	0.2328
Hispanic; mate Black	-0.0324*	0.0183	-0.4883	0.3924
Hispanic; mate other	-0.0025	0.0111	-0.2175	0.1796

*Notes.**Significant at the 10% level; **significant at the 5% level; ***significant at the 1% level.

Table 7: Preference Parameter Estimates

	Men		Women	
	Coefficients	SE	Coefficients	SE
A. Cost of Liking				
Constant (λ_0)	-4.2025***	0.0466	2.8612***	0.0807
$\mathscr{L}_{wm}^{LS,\tau_m}(\lambda_1)$	-3.0752***	0.0573	-0.3649***	0.0372
age (λ_2)	3.1156***	0.0606	-3.9779***	0.5474
popularity (λ_3)	4.4926***	0.0776	-0.6666***	0.1576
B. Cost of Messaging				
Constant (γ_0)	1.8186***	0.0182	4.4643***	0.1547
$\mathscr{L}_{wm}^{s,\tau_m}(\gamma_1)$	-0.0879***	0.0090	-0.1639***	0.0196
age (γ_2)	-0.5684***	0.0588	-5.3573***	0.6476
popularity (γ_3)	-0.3260***	0.0271	-1.1033***	0.1266
Log-Likelihood	-1,618,048.808 (1,757.053)		-131,463.146 (1,501.259)	

*Notes.**Significant at the 10% level; **significant at the 5% level; ***significant at the 1% level.

Table 8: Cost Parameter Estimates

10 Predicted Matching Patterns

We quantify the impact of frictions and preferences on assortative matching and user welfare. Specifically, we compare the equilibrium matches achieved across the following protocols: (1) equilibrium matches in a default setting where everyone is in the control group, (2) equilibrium matches when both sides of the market (men and women) are gifted with the treatment, and (3) equilibrium matches in a frictionless environment. By comparing matches in (1) and (2), we want to quantify the welfare gains that users will experience when the platform lowers users' search costs by making treatment a free feature. This can provide managerial insights as to whether more information about users' preferences should be revealed in the profile page. By comparing matches in (1) and (3), we can quantify the relative impact of frictions and preferences on sorting, and also the impact of frictions on user welfare.

For matches in (1) and (2), instead of using the actual matches that are observed in the data, we will simulate the matches using our structural model (see section 10.1). The costs of liking and messaging prevent users from reaching out to many candidates, which in turn leads to different sets of candidates ending up with a match across different protocols. To make a correct comparison across different protocols, the initial pool of available candidates in the market must be identical across different protocols. Therefore, we use all and only the users who are part of the experiment (11,807 men and 3,998 women) as our initial pool of men and women attempting to find a partner, and simulate matches among these users.

We simulate equilibrium matches in a frictionless environment using the Gale-Shapley deferred-acceptance algorithm. Predicted matches under the control and treatment settings are obtained by introducing frictions to the deferred-acceptance algorithm using the structural model. This is similar to Banerjee et al. (2013) where they introduce ad-hoc constraints to the deferred-acceptance algorithm to account for search frictions.

In what follows, we first describe how we compute predicted matches in a frictionless environment using the Gale-Shapley deferred-acceptance algorithm. Then, we describe how we compute equilibrium matches when users engage in costly *liking* and messaging. Predicted matches from these simulations will then be used to answer questions regarding the relative impact of frictions and preferences and the impact of frictions on user welfare.

10.1 Empirical Strategy

The man-optimal stable matching using the deferred-acceptance algorithm is executed as follows:

1. All men first propose to their most highly ranked woman, as long as $\hat{u}_{M,mw} \geq V_M^{GS}(m)$.
2. Among all the offers that each woman receives, she selects the most highly ranked man, as long as $\hat{u}_{W,wm} \geq V_W^{GS}(w)$.
3. All men who haven't been selected then propose to their second most highly ranked woman.
4. If a woman receives a new offer that is higher-ranked than the one she is currently holding, the woman releases the old offer and keeps the new offer. Released man then has to propose to the next woman in his ranking list.
5. This process continues until all men go through all women such that $\hat{u}_{M,mw} \geq V_M^{GS}(m)$.

Ties are broken randomly. The above process describes how we obtain a set of stable matches implied by the estimated preferences when frictions are negligible.

To compute the stable matches using the Gale-Shapley deferred-acceptance algorithm, we need to obtain reservation values in the absence of frictions, i.e., $V_M^{GS}(m)$ and $V_W^{GS}(w)$. We use the following procedures to obtain $V_M^{GS}(m)$ for each user m:

1. For each m, we calculate m's utility for each candidate that he browses using estimated preference parameters, then we calculate the average utility over all candidates, i.e., $E_w[u_{M,mw}]$

2. Next, we regress $V_M(m)$ on $E_w[u_{M,mw}]$, $E_w[c^{like}_{M,mw}]$, $E_w[c^{msg}_{M,mw}]$ and their respective squared terms

3. Then we obtain the fitted value of $\hat{V}_M(m)$ while setting $E_w[c^{like}_{M,mw}]$ and $E_w[c^{msg}_{M,mw}]$ to zeros.

4. We use this fitted value of $\hat{V}_M(m)$ as $V^{GS}_M(m)$.

Note that the cost of browsing, $c^{\text{browse}}_{M,m}$, is omitted from the regression. This is because we do not estimate $c^{\text{browse}}_{M,m}$ separately in our model, but instead only estimate $V_M(m)$, which can be written as $V_M(m) = -c^{\text{browse}}_{M,m} + \max\left\{V^{LS}_{M,mw}(d_{mw} = \text{like}|\Omega^{LS,\tau_m}_{mw}), V^{LS}_{M,mw}(d_{mw} = \text{nlike}|\Omega^{LS,\tau_m}_{mw})\right\}$. While $c^{like}_{M,mw}$ and $c^{msg}_{M,mw}$ may be correlated with $u_{M,mw}$ through user characteristics' association with ℓ_{wm}, there is no reason to believe that $c^{\text{browse}}_{M,m}$ is correlated with $u_{M,mw}$, because we assume that the cost of browsing is constant across candidate profiles. However, it is highly likely that $c^{\text{browse}}_{M,m}$ is positively correlated with $c^{like}_{M,m}$ and $c^{msg}_{M,m}$, which can in turn lead to biased estimates of effects of $c^{like}_{M,m}$ and $c^{msg}_{M,m}$. Since we are interested only in fitted values of $\hat{V}_M(m)$ and not regression coefficients per se, any bias in coefficients of $c^{like}_{M,m}$ and $c^{msg}_{M,m}$ will absorb the effect of $c^{\text{browse}}_{M,m}$. Therefore, we are not too concerned about the omission of $c^{\text{browse}}_{M,m}$.

Next, we need to construct ordinal preferences (i.e., rankings) over the entire set of women (men) for each man (woman) using estimated preference parameters from the model. Specifically, estimated preference parameters from the model are used to construct the utility that each man would get from matching with each woman in the market (and vice versa for women) using the following equation:

$$\hat{u}_{M,mw} = u_M(X_m, X_w; \hat{\Theta}_M) \tag{20}$$

Predicted utility $\hat{u}_{M,mw}$ is then transformed into an ordinal ranking $R_m(w)$ of user m with respect to woman w as

$$R_m(w) = n \qquad \text{if} \qquad \begin{cases} \hat{u}_{M,mw'} > \hat{u}_{M,mw} > \hat{u}_{M,mw''} \\ \text{and } R_m(w') = n-1 \text{ and } R_m(w'') = n+1 \end{cases} \tag{21}$$

where n is an integer. We apply this methodology to all users in the sample to obtain a full set of ordinal preferences for each user with respect to all users of the opposite gender.

We next describe how we incorporate frictions to the Gale-Shapley algorithm (1) in the default control setting, and (2) the treatment setting where users on both sides (both men and women) are gifted with the treatment. Since we are only using the pool of experimental users (who were either assigned to the treatment or control group) to form the two sides of the market, we need to simulate the values of ℓ_{wm}. That is, we simulate the $\ell ikes$ that a user receives from the opposite gender before he engages in search, for each pair of m and w. Specifically, the simulation of ℓ_{wm} and matching in an environment with frictions is executed as follows:

1. Construct the $\ell ikes$ that women send to men for each pair of users (w, m) as follows:
   (a) Draw random utility terms, e_{wm}, for each pair of w and m.[24]
   (b) For each w, using the draw and estimated preference parameters, construct the expected utility from $\ell iking$ and not $\ell iking$ given her information set, i.e., $E_w \left\{ \max_{\mu_{wm}} \left[V^{MS}_{W,wm}(\mu_{wm} | d_{wm} = like, \Omega_{wm}) \right] \right\}$ and $E_w \left\{ \max_{\mu_{wm}} \left[V^{MS}_{W,wm}(\mu_{wm} | d_{wm} = nlike, \Omega_{wm}) \right] \right\}$
   (c) w decides to ℓike m if the expected utility from $\ell iking$ is greater than its cost:

$$\ell_{wm} = \begin{cases} 1 & \text{if } E_w \left\{ \max_{\mu_{wm}} \left[V^{MS}_{W,wm}(\mu_{wm} | d_{wm} = like, \Omega_{wm}) \right] \right\} \\ & \quad -E_w \left\{ \max_{\mu_{wm}} \left[V^{MS}_{W,wm}(\mu_{wm} | d_{wm} = nlike, \Omega_{wm}) \right] \right\} + \varepsilon^{\text{like}}_{wm} - \varepsilon^{\text{nlike}}_{wm} > c^{\text{like}}_{W,w} \\ 0 & \text{otherwise} \end{cases} \tag{22}$$

2. Taking ℓ_{wm} constructed in the previous step as given, create ordinal preferences of the "expected" utility for each user, for all users of the opposite gender
   (a) Draw the random utility term, e_{mw}, for each pair of m and w.
   (b) For each pair m and w, m first chooses $d_{mw} \in \{\text{like,nlike}\}$. m chooses $d_{mw} = \text{like}$ if and only if

$$E_w \left\{ \max_{\mu_{wm}} \left[V^{MS}_{W,wm}(\mu_{wm} | d_{wm} = like, \Omega_{wm}) \right] \right\} - E_w \left\{ \max_{\mu_{wm}} \left[V^{MS}_{W,wm}(\mu_{wm} | d_{wm} = nlike, \Omega_{wm}) \right] \right\} + \varepsilon^{\text{like}}_{mw} - \varepsilon^{\text{nlike}}_{mw} > c^{\text{like}}_{M,m} \tag{23}$$

   (c) Conditional on d_{mw}, m then chooses $\mu_{mw} \in \{\text{msg,nmsg}\}$. m chooses $\mu_{mw} = \text{msg}$ iff

$$V^{MS}_{M,mw}(\mu_{mw} = \text{msg} | d_{mw}, \Omega_{mw}) \geq V^{MS}_{M,mw}(\mu_{mw} = \text{nmsg} | d_{mw}, \Omega_{mw}) \tag{24}$$

[24] As opposed to Banerjee et al. (2013), who assume that the noise in the utility function comes from a measurement error, we follow Hitsch et al. (2010a), who assume that the error term is a "structural noise." This is because several important dimensions of the profile (such as picture or income) that affect choice and are observed by users are unobserved by the econometrician, whereas researchers in Banerjee et al. (2013) observe everything that is observed by the agent.

(d) For each m, compute the predicted expected utility $E\hat{U}_{mw}$:

$$
\begin{aligned}
E\hat{U}_{mw} = u_{M,mw} \cdot \mathbb{E}_m \left[P_{W,mw}(d_{mw}, \mu_{mw} | X_m, X_w, \ell_{mw}, \Omega_w)\right] \\
+ V_M(m) \cdot \left(1 - \mathbb{E}_m \left[P_{M,mw}(d_{mw}, \mu_{mw} | X_m, X_w, \ell_{mw}, \Omega_w)\right]\right) - c_m^{\text{like}} \mathbb{1}\{d_{mw} = \text{like}\} - c_m^{\text{msg}} \mathbb{1}\{\mu_{mw} = \text{msg}\}
\end{aligned}
\tag{25}
$$

(e) Transform $E\hat{U}_{mw}$ into ordinal ranking $R_m(w)$ such that

$$
R_m(w) = n \qquad \text{if} \qquad \begin{cases} E\hat{U}_{M,mw'} > E\hat{U}_{M,mw} > E\hat{U}_{M,mw''} \\ \text{and } R_m(w') = n-1 \text{ and } R_m(w'') = n+1 \end{cases} \tag{26}
$$

where n is an integer.

3. Compute equilibrium matches:

(a) Define I_m as the set of all profiles that m either *ℓiked* or messaged.

(b) All men first propose (either message or *ℓike*) to their most highly-ranked woman within the set I_m. Ties are broken randomly.

(c) Women consider all offers they receive. If a woman received a message from a man, the net utility of selecting this man becomes $u_{W,wm} - V_W(w)$. On the other hand, if a woman only receives a *ℓike* from a man but did not receive a message, the net utility from selecting this man becomes $u_{W,wm} - c^{msg}_{W,w} - V_W(w)$. This is because if a woman only receives a *ℓike*, she must initiate a conversation with a man, which is costly. Woman selects a man that gives the highest net utility as long as it is greater than zero.

(d) All men who haven't been chosen by women then propose to the next best woman within the set I_m.

(e) If a woman receives a new offer that is preferable to the one she is currently holding, she releases the previous offer. The released man then has to propose to the next woman on his list within the set I_m.

(f) Steps 3 (a) through (e) continue until each man m exhausts the list of women in his set I_m.

10.2 Frictions and Assortative Matching

	Control (CT)		Treated (TR)		
	Mean	SD	Mean	SD	TR-CT 95% CI
			Men-Optimal		
Age	4.7064	3.1564	5.6517	3.3799	0.6624, 1.2282
	[999]		[1,062]		
Education	2.3979	1.0450	2.6279	1.0287	0.0081, 0.4519
	[160]		[180]		
Popularity	0.4694	0.2688	0.4962	0.2436	0.0047, 0.0489
	[999]		[1,062]		
BodyType	0.4793	0.4287	0.4144	0.4369	-0.0203, 0.1501
	[197]		[203]		
Race	0.5857	0.4125	0.6748	0.3722	0.0155, 0.1627
	[207]		[233]		
			Women-Optimal		
Age	5.1265	5.2442	5.6856	4.2316	0.1596, 0.9586
	[844]		[1,358]		
Education	1.2741	1.0718	1.8702	1.1889	0.3446, 0.8476
	[126]		[224]		
Popularity	0.5466	0.3944	0.6948	0.4370	0.1120, 0.1844
	[844]		[1,358]		
BodyType	0.5407	0.4116	0.5389	0.4425	-0.0780, 0.0816
	[185]		[287]		
Race	0.5372	0.4839	0.5827	0.4456	-0.0352, 0.1262
	[208]		[320]		

Notes. This table reports the attribute difference for matches achieved with Likers. Average number of matches across 100 simulations is in square brackets.

Table 9: Attribute Difference in Predicted Matches with Likers

Since we have already compared the matches under the control and one-sided treatment in Section 6, we are mostly interested in comparing sorting patterns of matches achieved in control and frictionless settings. Nevertheless, we simulate the mean attribute difference for the treatment setting to check whether the simulation using our structural model produces consistent results with what we have seen in the data in Section 6. Table 9 reports mean attribute differences of couples *who matched with Likers*, under the assumption of control (CT) and (two-sided) treatment (TR) settings. We also report the 95 percent confidence interval of the difference between the treatment and control (TR-CT) setting. Consistent with the patterns we observe in our data, attribute differences of the treatment setting are generally significantly larger compared to those of the control setting (except for race in women-optimal matches and body type). These patterns provide validation to simulated matches generated from our structural model.

Table 10 reports the mean attribute difference between couples for predicted matches under the assumption of control (CT), (two-sided) treatment (TR) and frictionless (GS) environments. Note that here we report the mean attribute for all matches (both Likers and NotLikers), not just matches with Likers. We also report the 95 percent confidence intervals of the difference between the protocols (TR-CT is the difference between treatment and control, GS-CT is the difference between frictionless and control). While the attribute differences of the treatment setting are generally larger compared to those of the control setting, the overall direction is ambiguous due to the presence of matches with NotLikers. As we completely remove frictions, however, we observe a significant increase in attribute difference between couples, compared to the control setting (except for body type in male-optimal matches). For example, when we look at male-optimal matches, the age difference between couples is 5.03 years in the control setting and 5.63 years in the frictionless setting. This suggests that approximately 11% of the positive sorting in age is due to frictions and 89% is due to preferences.[25] The difference in years of education between couples is 2.10 years in the control setting and 2.36 years in the frictionless setting. Similarly, approximately 11% of the positive sorting in years of education is due to friction and 89% is due to preferences. For popularity, 9%of the positive sorting in popularity is due to friction and 91% is due to preferences. Finally, 11% of the positive sorting in race is due to friction and 89% is due to preferences.

We find similar patterns in female-optimal matches. The age difference between couples is 5.3 years in the control setting and 6.34 years in the frictionless setting, attributing approximately 16% of sorting in age to frictions. The difference in years of schooling between couples is 1.47 years in the control setting and 1.94 years in the frictionless setting, attributing 24% of the sorting in years of schooling to frictions. For popularity, 13%of the positive sorting in popularity is due to frictions, and 19% of the positive sorting in race is due to frictions.

10.3 Welfare Analysis

In this section, we study (1) whether reducing frictions through provision of information makes users better off by comparing matches formed in control and treatment setting, and (2) we quantify the departure from efficiency caused by frictions on the platform by comparing matches formed in control and frictionless settings. Following Hitsch et al. (2010b), we implement this as follows: For

[25] 5.03/5.63=0.89

	Control (CT)		Treated (TR)		Gale-Shapley (GS)		TR− CT	GS − CT
	Mean	SD	Mean	SD	Mean	SD	95% CI	95% CI
				Men-Optimal				
Age	5.0363	3.3434	5.2346	3.1730	5.6365	4.9920	[-0.0087, 0.4053]	[0.3340, 0.8664]
	[1,895]		[1,916]		[2,115]			
Education	2.0970	0.8417	2.3601	1.0694	2.3599	1.5265	[0.1063, 0.4199]	[0.0660, 0.4598]
	[293]		[287]		[338]			
Popularity	0.5120	0.2663	0.5202	0.2411	0.5605	0.3330	[-0.0079, 0.0243]	[0.0298, 0.0672]
	[1,895]		[1,916]		[2,115]			
BodyType	0.5146	0.2003	0.4732	0.1656	0.5414	0.3339	[0.0153, 0.0675]	[-0.0120, 0.0656]
	[383]		[380]		[406]			
Race	0.5523	0.3375	0.6224	0.3525	0.6219	0.3897	[0.0232, 0.1170]	[0.0218, 0.1174]
	[405]		[434]		[521]			
				Women-Optimal				
Age	5.3044	3.9277	6.3199	4.0022	6.3414	7.1154	[0.3763, 0.8549]	[0.7756, 1.2554]
	[2,030]		[2,178]		[2,528]			
Education	1.4653	1.0733	1.7117	1.2545	1.9422	1.5652	[0.0610, 0.4318]	[0.2692, 0.6846]
	[293]		[326]		[408]			
Popularity	0.6091	0.3812	0.7356	0.4333	0.7041	0.5381	[0.1018, 0.1512]	[0.0673, 0.1227]
	[2,030]		[2,178]		[2,528]			
BodyType	0.5306	0.3533	0.5450	0.3466	0.6094	0.4532	[-0.0302, 0.0590]	[0.0285, 0.1291]
	[464]		[483]		[580]			
Race	0.4602	0.3858	0.5300	0.3829	0.5728	0.4596	[0.0229, 0.1167]	[0.0634, 0.1618]
	[526]		[509]		[657]			

Notes. Average number of matches across 100 simulations is in square brackets.

Table 10: Attribute Correlations in Predicted Matches

	Men					Women				
	Mean	Median	SD	Obs.	95% CI	Mean	Median	SD	Obs.	95% CI
					Men-Optimal					
Panel A. Rank Differences										
$\Delta\bar{R}^{TR-CT}$	29.160	9.742	241.802	10,389	[24.510, 33.810]	61.511	80.313	650.833	3,878	[41.027, 81.995]
$\Delta\bar{R}^{GS-TR}$	107.869	126.905	269.536	9,464	[102.439, 113.299]	183.884	221.820	1,193.369	3,878	[146.325, 221.443]
$\Delta\bar{R}^{GS-CT}$	132.567	137.237	328.267	9,129	[125.833, 139.301]	245.395	302.765	1,258.920	3,878	[205.772, 285.018]
$\%\Delta\bar{U}^{TR-CT}$	34.327	44.765	221.825	11,807	[30.326, 38.328]	13.364	11.258	89.741	3,922	[10.555, 16.173]
					Women-Optimal					
Panel B. Rank Differences										
$\Delta\bar{R}^{TR-CT}$	94.073	96.111	875.291	10,709	[77.495, 110.651]	41.110	3.0	252.392	3,849	[33.137, 49.084]
$\Delta\bar{R}^{GS-TR}$	266.903	291.0	1,103.264	8,809	[243.864, 289.942]	378.520	384.927	48.511	3,878	[376.993, 380.047]
$\Delta\bar{R}^{GS-CT}$	387.102	450.744	1,243.929	8,598	[360.809, 413.395]	418.918	387.803	252.726	3,918	[411.005, 426.831]
$\%\Delta\bar{U}^{TR-CT}$	107.519	-0.0623	1,243.415	11,807	[85.091, 129.947]	1,019.449	63.160	14,016.441	3,922	[580.799, 1,458.099]

Table 11: Changes in Welfare

each user, we assign an ordinal ranking to all candidates based on the predicted utility. If there are N candidates, the most desirable candidate will be ranked N^{th}, the second desirable candidate will be ranked $N-1^{\text{th}}$, and so forth. For each user $i \in \{m, w\}$, let R_i^{CT} be the rank of i's matched partner predicted under the control setting, R_i^{TR} be the rank of i's matched partner predicted under the treatment setting, and R_i^{GS} be the rank of i's matched partner predicted under the frictionless setting. Denote $\Delta R_i^{s1-s2} = R_i^{s1} - R_i^{s2}$ as the difference between the ranks achieved under the different scenarios, where $s1$ and $s2$ are indexes denoting either control, treatment or frictionless protocols. Denote $\Delta \bar{R}^{s1-s2}$ as the mean difference in ranks across users. Then, if $\Delta \bar{R}^{s1-s2}$ is positive, $s1$ setting could have improved, on average, on the allocation from setting $s2$.

In Table 11 Panels A and B, we report means, medians and standard deviations of predicted average rank differences across protocols for men-optimal and women-optimal equilibrium, respectively. In men-optimal matches, we find that the average rank difference between matches formed in treated settings and control settings is 29.2 for men and 61.5 for women. To interpret the magnitude of this difference, we can express the rank differences as a percent of the highest achievable rank ($100 \times \Delta \bar{R}/N$), which gives us rank improvement of 0.73 percentage points for men and 0.52 percentage points for women. Despite small magnitude, this suggests that if the platform makes treatment a free feature for everyone, users will be better off in terms of the ranking of their matched partner. When we completely remove all frictions, men experience an improvement in their partner's rankings by 132.6 and women experience an improvement in partner ranking by 245.4 compared to the control setting. In terms of percent of the highest achievable rank, man experience an improvement of 3.4 percentage points and women experience an improvement of 2.1 percentage points. We find similar patterns in women-optimal matches. The average difference between the treatment matches and the control matches is 94.1 for men and 41.1 for women (improvement of 2.4 percentage points for men and 0.34 percentage points for women), and the average difference between the Gale-Shapley matches and the control matches is 387.1 for men and 418.9 for women (improvement of 9.9 percentage points for men and 3.5 percentage points for women).

Unlike in a frictionless environment where only preferences shape the matching outcomes, users in control and treatment settings have to pay the costs of $\ell iking$ and messaging if they decide to ℓike or send a message. To see whether treatment leads to better outcomes for users in terms of overall welfare (not just ranking of the partner), we compare the utility that users receive net of costs of $\ell iking$ and messaging. Specifically, we subtract the total sum of costs of $\ell iking$ and messaging that a user incurs during the entire process of the deferred-acceptance algorithm from the utility that a user obtains from a match, and we compare the percent change of this value between the treatment and control environments (denoted as $\%\Delta\bar{U}^{TR-CT}$). If this value is positive, the treatment setting could have improved user welfare compared to the control setting. The average percentage change in net utility between the treatment and control setting in men-optimal matches is 34% for men and 13% for women. We find similar patterns in female-optimal matches. These numbers suggest that treatment leads to a higher level of welfare for users compared to the default control setting, even after considering the costs of $\ell iking$ and messaging.

11 Conclusion

This paper investigates the impact of frictions on the formation of a match in two-sided markets. With agents on both sides having private preferences regarding each others' characteristics which are often private, finding a match based on mutual agreement requires extensive costly search. Using data from an online dating platform, we estimate a model of costly search for information that incorporates preference heterogeneity across users. Our estimation results reveal that frictions play a significant role in shaping matching outcomes. Our counterfactuals reveal that couples matched in a frictionless environment display greater difference in attributes (in terms of age, education level, race, and popularity), compared to couples who matched in a market with frictions. We also find that reducing frictions lead to significant gains in terms of partner rankings and utility.

Our findings provide important managerial implications for the pricing of premium features, in how much users are willing to pay for an additional piece of information about the preferences of the other side. In addition, our findings shed light on what type of information should be displayed on users' profiles. Information that is helpful in gauging the preferences of other users can greatly improve consumer experience. Our findings are not limited to online dating contexts; they are broadly applicable to other two-sided market contexts where matches may not form due to imperfect information about each other's preferences.

Our paper also provides insights into how the design of online platforms can contribute to diversity, given that one-third of U.S. marriages now originate from online encounters. However, due to data limitations, we are unable to quantify the long-term effects of reducing frictions in online dating, such as the impact on marriages and the accumulation of human capital through children's education. Future research should aim to address this gap, as quantifying the long-term effects of reducing frictions in online dating platforms can provide solutions to alleviating persistent social inequality.

References

Adachi, H. (2003): "A search model of two-sided matching under nontransferable utility," *Journal of Economic Theory*, 113, 182–198.

Arcidiacono, P., A. Beauchamp, and M. McElroy (2016): "Terms of endearment: An equilibrium model of sex and matching," *Quantitative Economics*, 7, 117–156.

Banerjee, A., E. Duflo, M. Ghatak, and J. Lafortune (2013): "Marry for What? Caste and Mate Selection in Modern India," *American Economic Journal: Microeconomics*, 5, 33–72.

Bapna, R., E. I. McFowland, P. Mojumder, J. Ramaprasad, and A. Umyarov (2022): "So, Who Likes You? Evidence form a Randomized Field Experiment," *Management Science*, 0, 1–19.

Bapna, R., J. Ramaprasad, G. Shmueli, and A. Umyarov (2016): "One-Way Mirrors in Online Dating: A Randomized Field Experiment," *Management Science*, 62, 3085–3391.

Baumeister, R. F., S. R. Wotman, and A. M. Stillwell (1993): "Unrequited Love: On Heartbreak, Anger, Guilt, Scriptlessness, and Humiliation," *Journal of Personality and Social Psychology*, 64, 377–394.

Bloch, F. and H. Ryder (2000): "Two-Sided Search, Marriages, and Matchmakers," *International Economic Review*, 41, 93–115.

Bojd, B. and H. Yoganarasimhan (2019): "Star-Cursed Lovers: Role of Popularity Information in Online Dating," *Working Paper*.

Burdett, K. and M. G. Coles (1997): "Marriage and Class," *The Quarterly Journal of Economics*, 112, 141–168.

Chan, T. Y., B. H. Hamilton, and N. W. Papageroge (2015): "Health, Risky Behavior and the Value of Medical Innovation for Infectious Disease," *Review of Economic Studies*, 83, 1456–1510.

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

Fisman, R., S. S. Iyengar, E. Kamenica, and I. Simonson (2006): "Gender Differences in Mate Selection: Evidence From a Speed Dating Experiment," *The Quarterly Journal of Economics*, 121, 673–697.

——— (2008): "Racial Preferences in Dating," *Review of Economic Studies*, 75, 117–132.

Flinn, C. J. and D. Del Boca (2012): "Endogenous household interaction," *Journal of Econometrics*, 166, 49–65.

Fong, J. (2018): "Search, Selectivity, and Market Thickness in Two-Sided Markets," *working paper*.

Fradkin, A. (2015): "Search, Matching, and the Role of Digital Marketplace Design in Enabling Trade: Evidence from Airbnb," *Working Paper*.

Fréchette, G. R., A. E. Roth, and M. U. Ünver (2008): "Unraveling yields inefficient matchings: evidence from post-season college football bowls," *The RAND Journal of Economic*, 38, 967–982.

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

Halaburda, H., M. Jan Piskorski, and P. Yildirim (2017): "Competing by Restricting Choice: The Case of Matching Platforms," *Management Science*, 64, 3574 – 3594.

Hitsch, G. J., A. Hortaçsu, and D. Ariely (2010a): "Matching and Sorting in Online Dating," *American Economic Review*, 100, 130–63.

——— (2010b): "What makes you click?-Mate preferences in online dating," *Quantitative Marketing and Economics*, 8, 393–427.

Horton, J. J. (2014): "Misdirected Search Effort in a Matching Market: Causes, Consequences and a Partial Solution," *Proceedings of the Fifteenth ACM Conference on Economics and Computation*, 357–357.

Kurzban, R. and J. Weeden (2005): "HurryDate: Mate preferences in action," *Evolution and Human Behavior*, 26, 227–244.

Lee, S. (2015): "Effect of Online Dating on Assortative Mating: Evidence from South Korea," *Journal of Applied Econometrics*, 31, 1120–1139.

Lee, S. and M. Niederle (2015): "Propose with a rose? Signaling in internet dating markets," *Experimental Economics*, 18, 731–755.

McNamara, J. M. and E. J. Collins (1990): "The job search problem as an employercandidate game," *Journal of Applied Probability*, 27, 815–827.

Raquel, F. (2003): "Household formation, inequality, and the macroeconomy," *Journal of the European Economic Association*, 1, 683–697.

Raquel, F. and R. Rogerson (2001): "Sorting and Long-Run Inequality," *Quarterly Journal of Economics*, 116, 1305–1341.

Richards Shubik, S. (2015): "Peer effects in sexual initiation: Separating demand and supply mechanisms," *Quantitative Economics*, 14, 663–702.

Wong, L. Y. (2003): "Structural Estimation of Marriage Models," *Journal of Labor Economics*, 21, 699–727.

Appendix

A Randomization Check

	Men							
	Control			Treatment				
Variable	Mean	SD	Obs.	Mean	SD	Obs.	t-stat	*P* value
Age	30.5935	10.6457	30,869	30.6530	10.5916	30,653	0.0822	0.9345
Popularity	5.2541	2.8646	30,869	5.2409	2.8600	30,653	0.5726	0.5669
HighSchool	0.1461	0.3533	6,261	0.1445	0.3516	6,166	0.2595	0.7953
TwoYear	0.1760	0.3809	6,261	0.1790	0.3834	6,166	-0.4428	0.6579
University	0.5351	0.4988	6,261	0.5292	0.4992	6,166	0.6552	0.5123
PostGrad	0.1428	0.3499	6,261	0.1473	0.3544	6,166	-0.7076	0.4792
Thin	0.1654	0.3716	5,634	0.1605	0.3671	5,582	0.7036	0.4817
Average	0.6560	0.4751	5,634	0.6550	0.4754	5,582	0.1175	0.9065
LittleExtra	0.1409	0.3480	5,634	0.1467	0.3539	5,582	-0.8739	0.3822
Overweight	0.0376	0.1903	5,634	0.0378	0.1907	5,582	-0.0476	0.9620
Asian	0.1031	0.3041	8,146	0.1018	0.3024	8,085	0.2782	0.7808
White	0.6701	0.4702	8,146	0.6639	0.4724	8,085	0.8379	0.4021
Black	0.1062	0.3081	8,146	0.1126	0.3161	8,085	-1.2996	0.1938
Hispanic	0.1053	0.3070	8,146	0.1072	0.3094	8,085	-0.3943	0.6934
Other Race	0.0849	0.2788	8,146	0.0826	0.2753	8,085	0.5351	0.5926

Table A.1: Randomization Check—Men

	Women							
	Control			Treatment				
Variable	Mean	SD	Obs.	Mean	SD	Obs.	t-stat	*P value*
Age	30.9072	10.7242	19,131	30.7887	10.7629	19,347	1.0901	0.2757
Popularity	4.2480	3.2444	19,131	4.2436	3.2393	19,347	0.1338	0.8936
HighSchool	0.1352	0.3420	3,329	0.1343	0.3410	3,380	0.1027	0.9182
TwoYear	0.1787	0.3832	3,329	0.1719	0.3773	3,380	0.7365	0.4614
University	0.5314	0.4991	3,329	0.5414	0.4984	3,380	-0.8236	0.4102
PostGrad	0.1547	0.3617	3,329	0.1524	0.3594	3,380	0.2651	0.7909
Thin	0.2742	0.4462	2,498	0.2701	0.4441	2,514	0.3286	0.7425
Average	0.5320	0.4991	2,498	0.5155	0.4999	2,514	1.1703	0.2420
LittleExtra	0.1405	0.3476	2,498	0.1591	0.3659	2,514	-1.8446	0.0652
Overweight	0.0532	0.2246	2,498	0.0553	0.2286	2,514	-0.3199	0.7491
Asian	0.1271	0.3331	4,548	0.1314	0.3379	4,764	-0.6200	0.5352
White	0.6546	0.4756	4,548	0.6474	0.4778	4,764	0.7304	0.4652
Black	0.1321	0.3387	4,548	0.1299	0.3363	4,764	0.3164	0.7517
Hispanic	0.1062	0.3081	4,548	0.1075	0.3097	4,764	-0.1986	0.8426
Other Race	0.0556	0.2292	4,548	0.0542	0.2263	4,764	0.3119	0.7551

Table A.2: Randomization Check—Women

B Attribute Difference: Men and Women Separately

	Men: Matches Between m and Liker w					
	Control (CT)		Treatment (TR)			
	Mean	SD	Mean	SD	TR−CT	t-stat
Age	4.021	3.9858	4.4287	4.8440	0.4079	2.5487
	[1,437]		[1,712]			
Education	1.6818	1.9328	2.0	2.0085	0.3182	2.0703
	[308]		[355]			
Race	0.5540	0.0178	0.5892	0.0207	0.0352	1.2892
	[574]		[762]			
Body type	0.5613	0.0230	0.5913	0.0214	0.0300	0.9529
	[465]		[526]			
Popularity	0.5387	0.4368	0.5808	0.4655	0.0421	2.6006
	[1,437]		[1,712]			

Notes. Number of observations is in square brackets. When exchanging at least four messages, pairs sent at least two messages each.

*Significant at the 10% level; **significant at the 5% level; ***significant at the 1% level.

Table B.1: Men Only: Attributes Differences With Initiated Matches

	Women: Matches Between m and Liker w					
	Control		Treatment			
	Mean	SD	Mean	SD	TR−CT	t-stat
Age	4.3355	4.2429	4.7368	4.5835	0.4013	1.7000
	[608]		[851]			
Education	1.8268	2.0080	1.9777	2.1278	0.1509	0.6255
	[127]		[179]			
Race	0.5994	0.0269	0.6543	0.0222	0.0549	1.5770
	[332]		[457]			
Body Type	0.5885	0.0316	0.6949	0.0280	0.1064	2.5179
	[243]		[272]			
Popularity	1.2470	1.0919	1.2629	1.0983	0.0159	0.2733
	(1.0919)		(1.0983)			
	[608]		[851]			

Notes. Number of observations is in square brackets. When exchanging at least four messages, pairs sent at least two messages each.

*Significant at the 10% level; **significant at the 5% level; ***significant at the 1% level.

Table B.2: Women Only: Attributes Differences With Initiated Matches