

Matching and Sorting in Online Dating

By GÜNTER J. HITSCH, ALİ HORTAÇSU, AND DAN ARIELY*

Using data on user attributes and interactions from an online dating site, we estimate mate preferences, and use the Gale-Shapley algorithm to predict stable matches. The predicted matches are similar to the actual matches achieved by the dating site, and the actual matches are approximately efficient. Out-of-sample predictions of offline matches, i.e., marriages, exhibit assortative mating patterns similar to those observed in actual marriages. Thus, mate preferences, without resort to search frictions, can generate sorting in marriages. However, we underpredict some of the correlation patterns; search frictions may play a role in explaining the discrepancy. (JEL C78, J12)

This paper studies the economics of match formation using a novel dataset obtained from a major online dating service. Online dating takes place in a new market environment that has become a common means to find a date or a marriage partner. According to comScore (2006), 17 percent of all North American and 18 percent of all European Internet users visited an online personals site in July 2006. In the United States, 37 percent of all single Internet users looking for a partner have visited a dating Web site (Mary Madden and Amanda Lenhart 2006). The Web site we study provides detailed information on the users' attributes and interactions, which we use to estimate a rich model of mate preferences. Based on the preference estimates, we then examine whether an economic matching model can explain the observed online matching patterns, and we evaluate the efficiency of the matches obtained on the Web site. Finally, we explore whether the estimated preferences and a matching model are helpful in understanding sorting patterns observed "offline," among dating and married couples.

Two distinct literatures motivate this study. The first is the market design literature, which focuses on designing and evaluating the performance of market institutions. A significant branch of this literature is devoted to matching markets (Alvin E. Roth and Marilda A. O. Sotomayor 1990), with the goal of understanding the allocation mechanism in a particular market, and assessing whether an alternative mechanism with better theoretical properties (typically in terms

* Hitsch: Booth School of Business, University of Chicago, 5807 S. Woodlawn Ave., Chicago, IL 60637 (e-mail: guenter.hitsch@chicagobooth.edu); Hortaçsu: Department of Economics, University of Chicago, 1126 E. 59th Street, Chicago, IL 60637 (e-mail: hortacsu@uchicago.edu); Ariely: Fuqua School of Business, Duke University, 1 Towerview Drive, Durham, NC 27708 (e-mail: dandan@duke.edu). We thank Babur De Los Santos, Chris Olivola, and Tim Miller for their excellent research assistance. We are grateful to Elizabeth Bruch, Jean-Pierre Dubé, Emir Kamenica, Derek Neal, Peter Rossi, Betsey Stevenson, and Utku Ünver for comments and suggestions. Seminar participants at the 2006 AEA meetings, Boston College, the Caltech 2008 Matching Conference, the Choice Symposium in Estes Park, the Conference on Marriage and Matching at New York University 2006, the ELSE Laboratory Experiments and the Field (LEaF) Conference, Northwestern University, the 2007 SESP Preconference in Chicago, SITE 2007, the University of Pennsylvania, the 2004 QME Conference, UC Berkeley, UCLA, the University of Chicago, UCL, the University of Naples Federico II, the University of Toronto, Stanford GSB, and Yale University provided valuable comments. This research was supported by the Kilts Center of Marketing (Hitsch), a John M. Olin Junior Faculty Fellowship, and the National Science Foundation, SES-0449625 (Hortaçsu).

TABLE 1—DATING SERVICE MEMBERS AND COUNTY PROFILE OF GENERAL DEMOGRAPHIC CHARACTERISTICS

Variable	San Diego			Boston		
	Dating service	General population	Internet users	Dating service	General population	Internet users
<i>General information</i>						
Population	11,024	2,026,020	1,180,020	10,721	2,555,874	1,581,711
Percentage of men	56.1	49.9	49.4	54.7	49.0	50.6
<i>Age composition</i>						
18 to 20 years	20.3	6.0	6.4	19.0	5.8	7.2
21 to 25 years	30.7	9.5	11.5	33.1	9.3	12.0
26 to 35 years	27.0	21.3	18.8	27.2	17.2	19.7
36 to 45 years	10.0	23.0	28.6	10.1	23.1	26.8
46 to 55 years	6.6	18.5	19.0	6.2	17.6	20.1
56 to 60 years	4.4	6.3	6.5	3.6	7.3	6.9
61 to 65 years	0.8	2.9	3.6	0.5	4.3	3.7
66 to 75 years	0.1	6.9	4.8	0.1	8.8	2.9
Over 76	0.2	5.7	0.8	0.2	6.8	0.7
<i>Race composition</i>						
Whites	72.4	61.9	71.3	82.2	84.2	89.1
Blacks	4.3	4.8	4.2	4.8	7.4	4.2
Hispanics	11.0	19.5	9.8	4.1	4.4	2.3
Asian	5.3	13.0	13.6	3.9	3.8	4.2
Other	7.1	0.9	1.1	5.0	0.3	0.2
<i>Marital status</i>						
Men						
Never married	65.6	31.8	28.5	67.2	35.3	36.8
Married and not separated	6.3	57.0	62.0	7.2	54.1	56.7
Separated	4.0	1.2	0.7	4.8	1.1	0.3
Widowed	1.8	2.3	1.5	1.4	3.6	1.0
Divorced	22.3	8.1	7.4	19.4	6.0	5.2
Women						
Never married	62.2	20.2	23.9	65.9	28.0	32.7
Married and not separated	2.6	57.0	62.5	2.0	49.0	55.9
Separated	3.7	3.9	1.9	4.3	2.4	0.9
Widowed	3.5	6.3	2.0	3.0	13.4	3.5
Divorced	28.1	12.3	9.7	24.7	7.2	7.0

of efficiency and/or stability) can improve on the allocation obtained by the current mechanism.¹ Online dating sites are similar to previously analyzed matching markets in that they are used to allocate indivisible “goods” without an explicit price or transfer mechanism. Unlike many of the markets studied by the matching literature, however, economists have rarely been consulted to design or analyze online dating sites. Furthermore, it is not clear whether any existing market “design” or mechanism in the online dating industry—including the specific site that we study—corresponds to a mechanism studied in the theoretical matching literature. Thus, the first objective of this paper is to explore whether an economic matching model can predict the matching outcomes obtained on the Web site and to assess how efficient these outcomes are.

The second literature that motivates this paper is the extensive body of work across economics (starting with Gary S. Becker 1973), sociology, social psychology, and anthropology that studies matching and sorting patterns in marriage markets. This literature has documented that

¹ For example, Roth (1984) studies the matching of medical interns to hospitals, Muriel Niederle and Roth (2003) study the matching of gastroenterologists to hospitals, Atila Abdulkadiroğlu and Tayfun Sönmez (2003) study the matching of children to schools, and Roth, Sönmez, and Utku Ünver (2004) study kidney transfers from donors to recipients.

TABLE 1—DATING SERVICE MEMBERS AND COUNTY PROFILE OF GENERAL DEMOGRAPHIC CHARACTERISTICS
(Continued)

Variable	San Diego			Boston		
	Dating service	General population	Internet users	Dating service	General population	Internet users
<i>Educational attainment</i>						
Have not finished high school	1.4	12.1	3.0	1.6	9.2	3.2
High school graduate	9.4	23.0	17.8	10.2	30.1	20.4
Technical training (two-year degree)	31.9	5.2	5.4	23.6	7.3	7.6
Some college	6.8	27.9	28.5	4.6	14.1	15.0
Bachelor's degree	28.6	22.7	31.5	33.9	22.2	29.7
Master's degree	11.2	6.0	9.0	16.4	11.7	16.3
Doctoral degree	3.5	1.5	2.6	3.8	3.3	5.2
Professional degree	7.3	1.7	2.3	5.8	2.0	2.6
<i>Income</i>						
Individuals with income data	6,549	283,442	224,339	6,349	396,065	281,619
Less than \$12,000	7.9	12.5	12.4	8.6	7.6	4.6
\$12,000 to \$15,000	5.1	3.0	1.9	3.9	5.0	6.0
\$15,001 to \$25,000	8.7	13.8	10.1	6.1	21.4	16.2
\$25,001 to \$35,000	14.1	23.3	22.3	12.5	19.9	21.4
\$35,001 to \$50,000	20.4	12.4	10.6	21.9	16.5	18.5
\$50,001 to \$75,000	20.0	17.3	20.2	22.8	21.7	24.6
\$75,001 to \$100,000	10.5	7.2	9.1	12.0	4.8	4.5
\$100,001 to \$150,000	6.7	7.5	9.5	7.1	1.9	2.7
\$150,001 to \$200,000	2.7	3.2	4.0	2.0	1.1	1.6
\$200,001 or more	3.9	0.0	0.0	3.0	0.0	0.0

Notes: The geographic information regards two Metropolitan Statistical Areas (MSAs). The Boston Primary MSA includes a New Hampshire portion. San Diego geographic information corresponds to the San Diego MSA. The site member information is from 2003. The figures for whites, blacks, Asians, and "other" ethnicities for the Current Population Survey (CPS) data correspond to those with non-Hispanic ethnicity. The income figures from the CPS data were adjusted to 2003 dollars. All the CPS estimates are weighted. We consider only individuals who are at least 18 years old. The percentages for the column "Internet users" are from the CPS sample, restricted to individuals who declare they use the Internet.

Source: Estimates from CPS Internet and Computer use Supplement, September 2001.

marriages are not random, but exhibit strong sorting patterns along many attributes (for recent surveys, see e.g., Matthijs Kalmijn 1998; Martin Browning, Pierre-André Chiappori, and Yoram Weiss 2008). For example, marriage partners are similar in age, education levels, and physical traits such as looks, height, and weight. Such sorting patterns can arise due to several distinct reasons (Kalmijn 1998). First, sorting can arise due to search frictions, independent of preferences. For example, sorting along educational attainment might not reflect a preference for a partner with a certain education level, but rather the fact that many people spend much of their time in the company of others with a similar level of education in school, college, or at work. Alternatively, sorting can arise in the absence of any search frictions as an equilibrium consequence of preferences and the market mechanism. For example, men and women may prefer a match with a similar partner, in which case sorting is due to "horizontal" mate preferences. On the other hand, preferences might be purely "vertical," in the sense that each mate ranks all potential partners in the same way. In the equilibrium of a frictionless market, the ranks of the matched men and women will then be perfectly correlated. If the ranks are monotonically related to the mate's attributes, there will also be sorting along these attributes. These distinct causes of sorting are not mutually exclusive: the observed sorting patterns can be due both to search frictions and mate preferences.

Online dating provides us with a unique environment to study the causes of sorting. First, search frictions in online dating markets are minimal—indeed, a main reason for the existence

of online dating sites is to make the search for a partner as easy as possible. Hence, online dating markets allow us to study to what extent sorting is due to mate preferences and the market mechanism. Second, online dating allows us to observe user attributes and mate choices in great detail. Moreover, the information available to us, the researchers, is similar to the information available to the users of the online dating site. In particular, we observe the choice sets faced by the users and the choices they make from these choice sets. Thus, online dating data are particularly useful to estimate mate preferences over various mate attributes. Marriage market data, in contrast, usually lack information on crucial mate attributes such as physical traits, do not include information on the agents' choice sets, and record final matches (marriages) only. A potential disadvantage of our data, however, is that we do not observe if two users who met online went on an actual date or eventually got married. Instead, we define a match as an event where two users exchanged information that indicates if they were about to meet offline.

The second objective of this paper is to investigate if sorting along mate attributes is also prevalent in online dating, and whether an economic matching model, based on estimated mate preferences, can predict the sorting patterns among offline dating and married couples.

The steps we take toward achieving our research objectives are as follows. We first specify mate preferences as a function of own and partner attributes. To estimate these preferences, we exploit the detailed information on the site users' attributes and partner search behavior contained in our data. The estimation strategy is based on the assumption that a user contacts a partner if and only if the potential utility from a match with that partner exceeds a threshold value (a "minimum standard") for a mate. Such a rule is consistent with equilibrium behavior in a two-sided search model (Hiroyuki Adachi 2003, for example). We also allow for the possibility of "strategic behavior" that arises if there is a cost of sending an e-mail or a cost from rejection: if such costs are significant, a site user may not contact a desirable mate if he or she expects that the probability of a match with this mate is small. To that end, we include the user's expectation of being acceptable to the potential mate in the threshold rule. Assuming rational expectations (the subjective and actual probabilities of being accepted coincide), we can infer this expectation from the empirical probability of receiving a reply from the mate, which can be estimated from our data.

Based on the preference estimates, we predict who matches with whom using the algorithm of David Gale and Lloyd S. Shapley (1962).² Although the Gale-Shapley mechanism does not provide a literal description of how matches are formed online (or offline), it has some appealing features for which we consider it as a theoretical benchmark. First, Adachi (2003) shows that the stable (equilibrium) matching predicted by the Gale-Shapley algorithm can be seen as the limit outcome of a two-sided search and matching model with negligible search costs, which resembles the institutional environment of online dating more closely. Second, the Gale-Shapley model provides an efficiency benchmark, since the stable matching predicted by the algorithm is also the Pareto-optimal match, within the set of stable matches, for the side of the market (men or women) that makes the offers in the deferred acceptance procedure (Roth and Sotomayor 1990, Theorem 2.12).

We first document that the users of the dating site sort along various attributes, such as age, looks, income, education, and ethnicity. The Gale-Shapley model, based on the estimated mate preferences, is able to predict these sorting patterns: the observed correlations in user attributes largely agree with the correlations in the predicted stable matches. This finding provides an out-of-sample confirmation of the validity of the estimated preferences, which are based only on data

² We do not observe offline activities such as eventual marriages in our data. Instead, we define a match as an event where two users exchange a phone number or e-mail address, or an e-mail that contains certain keywords such as "get together" or "let's meet."

of the users' first-contact decisions, but not on information about the observed matches which are typically achieved only after several e-mails are sent back and forth. Furthermore, the result shows that in online dating, sorting can arise without any search frictions: mate preferences, rational behavior, and the equilibrium mechanism by which matches are formed generate sorting patterns that are qualitatively similar to those observed "offline."

The strong agreement between the observed and predicted matches suggests that the online dating market achieves an approximately efficient outcome within the set of stable matches. We further confirm this result by showing that, on average, the site users would not prefer the matches obtained under the Gale-Shapley mechanism to the actually achieved match. The conclusion from these findings is that the design of the online dating site that we study provides an efficient, frictionless market environment.

In the last part of this paper we explore to what extent our approach can also explain "offline" sorting patterns in marriages. There are two caveats to this exercise. First, mate preferences and sorting patterns in online dating may differ from the mate preferences and resulting sorting patterns in a marriage market. Previous studies, however, do not support this objection: Edward O. Laumann et al. (1994) demonstrate similar degrees of sorting along age, education, and ethnicity/race across married couples, cohabiting couples, and couples in a short-term relationship (we confirm these facts using data from the National Survey of Family Growth). Nonetheless, in order to make statements about marriage patterns based on the preference estimates obtained from our data, we need to assume that conditional on observable attributes, the users of our dating site do not differ in their mate preferences from the population at large. Second, the Gale-Shapley model is a nontransferable utility framework, which may be appropriate for dating, while marriages may be better described by a transferable utility framework.

We cannot directly compare the attribute correlations observed online to correlations in marriages due to differences between the online and offline populations along demographic characteristics. Thus, we reweight our sample of Web site users to match the offline population along key demographic attributes, and then predict equilibrium matches for this "synthetic" sample of men and women. We find that the Gale-Shapley algorithm predicts sorting patterns that are qualitatively, and for some attributes quantitatively, similar to sorting patterns in marriages. This suggests that preferences are one important cause of offline sorting; the prevalence of ethnically homogeneous marriages, for example, is unlikely to be due entirely to segregation. However, we underpredict the exact degree of sorting along some attributes, most prominently for education. One possible reason for the difference between the actual and predicted correlation in education (and other attributes) is that search frictions are also one cause of sorting in marriages.

Finally, we attempt to uncover the importance of different preference components in driving observed sorting patterns. As we discussed above, the horizontal and vertical components of preferences can, in principle, lead to the same matching outcomes. Our framework enables us to analyze the relative importance of these two different components by constructing counterfactual preference profiles that omit one of the components and recomputing the equilibrium matches. The result of the exercise indicates that "horizontal" preference components are essential in generating the sorting patterns observed in the data. A similar exercise suggests that same-race preferences are an essential source of the observed patterns of marriage within ethnic groups.

Our work is related to recent structural econometric work that estimates mate preferences based on marriage data (Linda Y. Wong 2003; Eugene Choo and Aloysius Siow 2006; Christopher J. Flinn and Daniela del Boca 2006). The common empirical strategy of these papers is to fit observed marriage outcomes using a structural model of equilibrium match formation, in which preferences are parametrized. While Flinn and Del Boca (2006) use the Gale-Shapley model for the marriage market, Choo and Siow (2006) use a frictionless transferable utility matching framework. Perhaps the paper that is closest to ours is Wong (2003),

which nests an equilibrium two-sided search model within a maximum likelihood procedure to explain marriage outcomes in the Panel Study of Income Dynamics (PSID). She also utilizes data on time-to-marriage to pin down the arrival rate of marriage opportunities. Compared to these papers, our data contain more detailed mate attribute information; measures of physical traits, for example, are not used by the papers noted above. Our setting also allows us to observe the search process directly, providing us with information regarding the choice sets available to agents, and enabling us to estimate mate preferences based on “first-contact” behavior alone. On the other hand, our data on online dating are, by design, less related to marital preferences than data on actual marriages.

Our work is also related to a recent series of papers utilizing data from “speed-dating” events by Robert Kurzban and Jason Weeden (2005), Raymond Fisman et al. (2006, 2008), and Paul W. Eastwick and Eli J. Finkel (2008). While the online dating sample we use is larger and, compared to most of these papers, more representative of the population at large, our revealed preference findings are similar. The main goal of our paper, however, is to characterize the equilibrium market outcomes in online dating and marriage markets, which is not attempted by the aforementioned papers.

I. Market and Data Description

We first provide a description of online dating that clarifies how the data were collected. We then provide a description of the data; a more detailed description can be found in Hitsch, Hortaçsu, and Ariely (2009).

A. Mechanics of Online Dating

When they first join the dating service, the users answer questions from a mandatory survey and create “profiles” of themselves.³ A profile is a Web page that provides information about a user and can be viewed by any other members of the dating site. The users indicate various demographic, socioeconomic, and physical characteristics, such as their age, gender, education level, height, weight, eye and hair color, and income. The users also indicate why they joined the service, for example, to find a partner for a long-term relationship, or, alternatively, a partner for a “casual” relationship. In addition, the users provide information that relates to their personality, life style, or views. For example, the site members indicate what they expect on a first date, whether they have children, their religion, whether they attend church frequently, and their political views. The users can also answer essay questions that provide additional details on their attitudes and personalities, but this information is too unstructured to be usable for our analysis. Many users also include one or more photos in their profile. We will explain in more detail later how we used these photos to construct a measure of the users’ physical attractiveness.

After registering, the users can browse, search, and interact with the other members of the dating service. Typically, users start their search by indicating an age range and geographic location for their partners in a database query form. The query returns a list of “short profiles” indicating the user name, age, a brief description, and, if available, a thumbnail version of the photo of a potential mate. By clicking on one of the short profiles, the searcher can view the full user profile, which contains socioeconomic and demographic information, a larger version of the profile photo (and possibly additional photos), and answers to several essay questions. Upon reviewing

³ Neither the names nor any contact information was provided to us in order to protect the privacy of the users.

TABLE 2—PHYSICAL CHARACTERISTICS OF DATING SERVICE MEMBERS
VERSUS GENERAL POPULATION

Variable	Men		Women	
	Dating service	General population	Dating service	General population
<i>Weight (pounds)</i>				
20–29 years	175.3	172.1	136.3	141.7
30–39 years	184.6	182.5	136.9	154.2
40–49 years	187.9	187.3	138.4	157.4
50–59 years	187.0	189.2	140.8	163.7
60–69 years	188.5	182.8	147.2	155.9
70–79 years	185.9	173.6	144.1	148.2
<i>Height (inches)</i>				
20–29 years	70.6	69.3	65.1	64.1
30–39 years	70.7	69.5	65.1	64.3
40–49 years	70.7	69.4	65.1	64.1
50–59 years	70.6	69.2	64.7	63.7
60–69 years	70.3	68.5	64.6	63.1
70–79 years	69.0	67.7	63.7	62.2
<i>BMI</i>				
20–29 years	24.7	25.2	22.6	24.3
30–39 years	25.9	26.5	22.7	26.3
40–49 years	26.4	27.3	23.0	27.0
50–59 years	26.3	27.8	23.6	28.4
60–69 years	26.7	27.3	24.8	27.6
70–79 years	27.7	26.7	25.0	26.9

Notes: BMI (body mass index) is calculated as weight (in kilograms) divided by height (in meters) squared.

Source: General population statistics obtained from the National Health and Nutrition Examination Survey, 1988–1994 Anthropometric Reference Data Tables.

this detailed profile, the searcher then decides whether to send an e-mail (a “first contact”) to the user. Our data contain a detailed, second-by-second account of all these user activities.⁴ We know if and when a user browses another user, views his or her photo(s), sends an e-mail to another user, answers a received e-mail, etc. We have additional information that indicates whether an e-mail contains a phone number, e-mail address, or keyword or phrase such as “let’s meet,” based on an automated search for special words and characters in the exchanged e-mails.⁵

In order to initiate a contact by e-mail, a user has to become a paying member of the dating service. Once the subscription fee is paid, there is no limit on the number of e-mails a user can send. All users can reply to an e-mail that they receive, regardless of whether they are paying members or not.

B. Data Description

The analysis in this paper is based on a sample of 3,004 men and 2,783 women located in Boston and San Diego. The sample was chosen from 22,000 users of an online dating service; see Section IIID for a description of how the sample was selected. We observe all user activities over a period of three and a half months in 2003.

⁴ We obtained this information in the form of a computer log file.

⁵ We do not see the full content of the e-mail, or the e-mail address or phone number that was exchanged.

The registration survey asks users why they are joining the site. It is important to know the users' motivation when we estimate mate preferences, because we need to be clear whether these preferences are with regard to a relationship that might end in a marriage or whether the users seek a partner only for casual sex. More than half of all observed activities (e-mails sent) are due to users who have a stated preference for a long term relationship. Furthermore, more than 20 percent of all activities are due to users who state they are "just looking/curious," an answer that is likely, given that it sounds less committal than "hoping to start a long-term relationship." Activities by users who state they seek a casual relationship account for less than 4 percent of all activities.

To understand how representative our sample is compared to the general population, in Table 1, we contrast the reported characteristics of the site users with two samples of the population from the CPS: one sample is representative of the general population in Boston and San Diego, while the other sample only contains Internet users. Men are somewhat overrepresented among the dating service users. As expected, the site users are younger than the population at large: the median user is in the 26–35 age range, whereas the median person in both CPS samples is in the 36–45 age range. The majority of the dating service users are single; married users account for less than 1 percent of all activities (e-mails sent). Site users are more educated and have a higher income than the general population; once we condition on Internet use, however, the remaining differences are not large. The profile of ethnicities represented among the site users roughly reflects the profile in the corresponding geographic areas, especially when conditioning on Internet use, although Hispanics and Asians are somewhat underrepresented on the San Diego site and whites are overrepresented. Hence, along demographic and socioeconomic attributes, our sample appears to be fairly representative of the population at large.

Our dataset contains detailed (although self-reported) information regarding the physical attributes of the users, and 27.5 percent post one or more photos online. For the rest of the users, the primary source of information about their appearance is the survey, which asks users to rate their looks on a subjective scale. Users also report their height and weight. As we expected that these are among the main attributes not truthfully reported, in Table 2, we compare the reported characteristics with information on the whole US population, obtained from the National Health and Examination Survey Anthropometric Tables.⁶ We find that the average weight reported by women is lower (between 6 and 20 pounds) than the average weight in the population. Men report slightly higher weights than the national averages. Also, the stated height of both men and women is somewhat above the US average. Thus, overall the data provide evidence for only small levels of misrepresentation.

We constructed an attractiveness rating for the photos posted by the site users. This measure is based on the evaluations (on a scale from 1 to 10) provided by 100 students at the University of Chicago.⁷ Following Jeff E. Biddle and Daniel S. Hamermesh (1998), we standardized each photo rating by subtracting the mean rating given by the subject and dividing by the standard deviation of the subject's ratings. We then averaged these standardized ratings across subjects to obtain a single rating for each photo. Our data reveal that beauty is not entirely in the eye of the beholder: the attractiveness ratings are strongly correlated across observers.

II. A Modeling Framework for Analyzing User Behavior

Our data are in the form of user activity records that describe, for each user, which profiles were browsed, and to which profiles an e-mail was sent. In order to interpret the data using a

⁶ The data are from the 1988–1994 survey and only cover Caucasians.

⁷ Each photo was evaluated by 11 students on average.

revealed preference framework, we start with the following assumption (which we will relax below in Section E):

ASSUMPTION 1: *Suppose a user browses the profiles of two potential mates, w and w' , and sends an introductory e-mail to w but not to w' . Then the user prefers a potential match with w over a potential match with w' .*

This assumption enables us to interpret user actions as binary choices over potential mates. Let $U_M(m, w)$ be the expected utility that male user, m , gets from a potential match with woman w , and let $v_M(m)$ be the utility m gets from his outside option of not responding to the ad. If m browses w 's profile, he chooses to send an e-mail if and only if

$$(1) \quad U_M(m, w) \geq v_M(m).$$

A. The Adachi Model

The threshold-crossing rule above arises naturally in a search model. In particular, we consider a search model due to Adachi (2003), which we believe provides a useful stylized description of user behavior on the dating site.

Adachi's model is set in discrete time, with period discount factor ρ . In each period, there are M men and W women in the market. In each period, man m comes across a randomly sampled profile, w . This profile is drawn from the distribution F_W , whereas women sample profiles of men from the distribution F_M . We assume that F_M and F_W are stationary and assign a positive probability of meeting each person on the opposite side of the market. A standard assumption (as in Ken Burdett and Melvyn G. Coles 1996; Adachi 2003) that guarantees stationarity is that men and women who leave the market upon a match are immediately replaced by agents who are identical to them.

Let $v_M(m)$ and $v_W(w)$ be the *reservation utilities* of man m and woman w from staying single and continuing the search for a partner. Define the following indicator functions:

$$\mathcal{A}_W(m, w) = \mathbb{I}\{\text{woman } w \text{ accepts man } m\} = \mathbb{I}\{U_W(m, w) \geq v_W(w)\},$$

$$\mathcal{A}_M(m, w) = \mathbb{I}\{\text{man } m \text{ accepts woman } w\} = \mathbb{I}\{U_M(m, w) \geq v_M(m)\}.$$

We can then characterize the utility that man m gets upon meeting a woman w :

$$(2) \quad EU_M(m, w) = U_M(m, w) \mathcal{A}_M(m, w) \mathcal{A}_W(m, w) \\ + v_M(m)(1 - \mathcal{A}_W(m, w)) \mathcal{A}_M(m, w) + v_M(m)(1 - \mathcal{A}_M(m, w)).$$

The first term in this expression is the utility from a mutual match, and the second and third terms capture the continuation utility when no double-coincidence of wants has occurred.

The Bellman equations characterizing the optimal reservation values and search rules of man m and woman w are given by

$$(3) \quad v_M(m) = \rho \int EU_M(m, w) dF_W(w), \\ v_W(w) = \rho \int EU_W(m, w) dF_M(m).$$

Adachi (2003) shows that the system of equations above defines a monotone iterative mapping that converges to a profile of reservation utilities $(v_M^*(m), v_W^*(w))$ solving the system, and thus characterizing the stationary equilibrium in this market.⁸ The equilibrium reservation utilities $(v_M^*(m), v_W^*(w))$ can be thought of as person-specific “prices” that clear demand for, and supply of, that person.

B. The Gale-Shapley Model

Under some conditions, the predictions of who matches with whom from the Adachi model are identical to the predictions of the seminal Gale-Shapley (1962) matching model. Before explaining this result in detail, we briefly review the Gale-Shapley model.

The matching market is populated by the same set of men and women as in Adachi’s model, $m \in \mathcal{M} = \{1, \dots, M\}$, $w \in \mathcal{W} = \{M + 1, \dots, M + W\}$. The preference orderings are generated by $U_M(m, w)$ and $U_W(w, m)$.⁹

Let $\mu(m)$ denote the match of man m that results from a matching procedure, and let $\mu(w)$ be the match of woman w . Note that if $\mu(m) \notin \mathcal{W}$, then $\mu(m) = m$, and if $\mu(w) \notin \mathcal{M}$, then $\mu(w) = w$; i.e., agents may remain single.

The matching μ is defined to be *stable* (in the Gale-Shapley sense) if there is no man m and woman w such that $U_M(m, w) > U_M(m, \mu(m))$ and $U_W(w, m) > U_W(w, \mu(w))$. That is, in a stable matching it is not possible to find a pair (m, w) who are willing to abandon their partners and match with each other.

The set of stable matches in the Gale-Shapley model is not unique. However, the set of stable matches has two extreme points: the “men-optimal” and “women-optimal” stable matches. The men-optimal stable match is unanimously preferred by men and opposed by all women over all other stable matches, and vice versa (Roth and Sotomayor 1990, Theorem 2.12). Either of these two extreme points can be reached through the use of Gale-Shapley’s *deferred-acceptance algorithm*.¹⁰

C. Equivalence between Decentralized Search Outcomes and Gale-Shapley Stable Matches

A remarkable result obtained by Adachi (2003) is that, as search costs become negligible, i.e., $\rho \rightarrow 1$, the set of equilibrium matches obtainable in the search model outlined above is identical to the set of stable matches in a corresponding Gale-Shapley marriage model.

Adachi’s insight derives from an alternative characterization of Gale-Shapley stable matchings. In particular, let $v_M(m) = U_M(m, \mu(m))$, and let $v_W(w) = U_W(w, \mu(w))$ be the utility that m and w get from their match partners. Adachi shows that, in a stable match, $v_M(m)$ and $v_W(w)$ satisfy the following equations:

$$(4) \quad \begin{aligned} v_M(m) &= \max_{\mathcal{W} \cup \{m\}} \{U_M(m, w) \mid U_W(w, m) \geq v_W(w)\}, \\ v_W(w) &= \max_{\mathcal{M} \cup \{w\}} \{U_W(w, m) \mid U_M(m, w) \geq v_M(m)\}. \end{aligned}$$

⁸ The solution is not unique, but has a lattice structure in strong analogy to the Gale-Shapley model. See the next section for further details.

⁹ We impose the restriction that the preferences are strict.

¹⁰ The algorithm that arrives at the men-optimal match works as follows. Men make offers (proposals) to the women, and the women accept or decline these offers. The algorithm proceeds over several rounds. In the first round, each man makes an offer to his most preferred woman. The women then collect offers from the men, rank the men who made proposals to them, and keep the highest ranked men engaged. The offers from the other men are rejected. In the second round, those men who are not currently engaged make offers to the women who are next highest on their list.

Furthermore, as $\rho \rightarrow 1$, the system of Bellman equations (3) becomes equivalent to the system of equations in (4). That is, as agents become more and more patient, or, equivalently, as search costs decline to zero, the search process will lead to matching outcomes that are stable in the Gale-Shapley sense. This is intuitive, as the equations (4) imply that in a stable match, man m is matched with the best woman who is willing to match with him, and vice versa.

D. Discussion

Of course, actual user behavior in online dating markets is not exactly described by the models of Adachi or Gale and Shapley. However, both models capture some basic mechanisms that are at work in online dating. The Adachi model captures the search process for a partner, and the plausible notion that people have an understanding of their own market value, which influences their threshold or “minimum standard” for a partner. The Gale-Shapley model, especially the deferred-acceptance algorithm, captures the notion that stability can be attained through a protocol of repeated rounds of offer-making and corresponding rejections, which reflects the process of the e-mail exchanges between the site users. Moreover, if search frictions on the online dating site are low enough, the difference in matching outcomes as predicted by the two modeling frameworks will be small, as suggested by Adachi’s equivalence result. In this case, Gale-Shapley’s deferred acceptance algorithm provides a computationally convenient method to characterize the set of stable matches that can be obtained in this market. However, whether the equivalence holds (at least in an approximate sense) is an empirical question, which we will investigate further in Section IV.

E. Costly Communication and Strategic Behavior

If sending a first-contact e-mail is costly, the threshold rule (1) may not hold. As an example, let us assume that there is a single dimensional index of attractiveness in the market, and consider the decision by an unattractive man as to whether he should send an introductory e-mail to a very attractive woman. If composing the e-mail is costly, or the psychological cost of being rejected is high, the man may not send an e-mail, thinking that the woman is “beyond his reach,” even though he would ideally like to match with her. Thus, the estimated preferences based on the threshold crossing rule reflect not only the users’ true preferences, but also their expectations on who is likely to match with them.

Such strategic behavior can bias the preference estimates, and thus we investigate whether strategic behavior is an important concern in our data. Consider a slight modification of the Adachi search model, where a user incurs a cost k to contact a potential mate and cost r if he or she is rejected. Then, *conditional on contacting* that mate, the utility that man m expects to get from a match with woman w is

$$(5) \quad EU_M^c(m, w) = -k + U_M(m, w) P_W(m, w) + (v_M(m) - r)(1 - P_W(m, w)).$$

Again, women consider all men who made them proposals, including the currently engaged man, and keep the highest ranked man among these. In each subsequent round, those men who are not engaged make an offer to the highest ranked woman whom they have not previously made an offer to, and women engage the highest ranked man among all currently available partners. The algorithm ends after a finite number of rounds. At this stage, men and women either have a partner or remain single. The women-optimal match is obtained using the same algorithm, where women make offers and men accept or decline these proposals.

Here, $P_w(m, w)$ is the probability that w will accept m 's offer. Man m will then contact woman w if and only if $EU_M^c(m, w) \geq v_M(w)$, which simplifies to the threshold-crossing condition

$$(6) \quad U_M(m, w) \geq v_M(m) - r + \frac{k + r}{P_w(m, w)}.$$

If contacting a partner is costless, $k = r = 0$, we obtain the original threshold rule, $U_M(m, w) \geq v_M(m)$.

In the next section, we use both the threshold crossing condition (1) and the more general condition (6) to estimate mate preferences. In particular, in Section B, we show how one can estimate the combined cost of rejection and e-mailing, $k + r$. A priori, however, we do not anticipate that strategic behavior is important in the context of online dating. Unlike a conventional marriage market, where the cost of approaching a potential partner is often nontrivial, online dating is designed to minimize this cost. The main cost associated with sending an e-mail is the cost of composing it. However, the marginal cost of producing yet another witty e-mail is likely to be small since one can easily personalize a polished form letter, or simply use a ‘‘copy and paste’’ approach. The fear of rejection should be mitigated by the anonymity provided by the dating site. Furthermore, rejection is common in online dating: in our data, 71 percent of men’s and 56 percent of women’s first-contact e-mails do not receive a reply.

III. Mate Preference Estimation

Our estimation approach is based on a sequence of binary decisions, as in the Adachi (2003) search and matching model of Section A. For each user, we observe the potential mates that he or she browses, and we observe whether a first-contact e-mail was sent. As developed in the previous section, if the cost of e-mailing is zero, man m contacts woman w if and only if $U_M(m, w) > v_M(m)$, and woman w contacts man m if and only if $U_W(m, w) > v_W(w)$. If there is a positive cost of sending an e-mail and a positive rejection cost, the threshold rule takes on the more general form shown above in equation (6).

We first discuss the details of the discrete choice approach. We then discuss how to account for positive costs of e-mailing in our estimation framework, and examine if there is any preliminary evidence pointing toward strategic behavior. Finally, we present the estimation results.

A. Discrete Choice Model

Given the threshold-crossing decision rule, mate preferences can be estimated using discrete choice methods. We assume that mate preferences depend on observed own and partner attributes, and on an idiosyncratic preference shock: $U_M(m, w) = U_M(X_m, X_w; \theta_M) + \epsilon_{mw}$. We split the attribute vector and the parameter vector into separate components: $X_w = (x_w, d_w)$, $\theta_M = (\beta_M, \gamma_M^+, \gamma_M^-, \vartheta_M)$. The latent utility of man m from a match with woman w is parametrized as

$$(7) \quad U_M(X_m, X_w; \theta_M) = x_w' \beta_M + (|x_w - x_m|'_+)^{\alpha} \gamma_M^+ + (|x_w - x_m|'_-)^{\alpha} \gamma_M^- \\ + \sum_{k,l=1}^N \mathbb{I}\{d_{mk} = 1 \text{ and } d_{wl} = 1\} \cdot \vartheta_M^{kl} + \epsilon_{mw}.$$

The first component of utility is a simple linear valuation of the woman’s attributes. The other components relate the man’s preferences to his own characteristics. The difference between

the woman's and man's attributes is $|x_w - x_m|_+$ if this difference is positive, and $|x_w - x_m|_-$ denotes the absolute value of this difference is the difference is negative.¹¹ For example, consider the difference in age between man m and woman w . If the coefficients corresponding to the age difference in both γ_M^+ and γ_M^- are negative, it means that users prefer someone of their own age. Note that each component of the difference terms is taken to the power α . The fourth component in (7) relates preferences to categorical attributes of both mates. Dummy variables indicating that man m and woman w possess a certain trait are represented by d_{mk} and d_{wl} . For example, if $d_{mk} = 1$ and $d_{wl} = 1$ indicate that m is white and that w is Hispanic, then the parameter ϑ_M^{kl} expresses the relative preference of white men for Hispanic women.

We assume that ε_{mw} has the standard logistic distribution, and is i.i.d. across all pairs of men and women. The reservation values $v_M(m)$ and $v_W(w)$ are estimated as person-specific fixed effects, denoted by c_m and c_w . Assuming that contacting a mate is costless and that there is no cost of rejection ($k = r = 0$), the choice probabilities then take the standard logit form:

$$(8) \quad \Pr\{m \text{ contacts } w \mid m \text{ browses } w\} = \frac{\exp(U_M(X_m, X_w; \theta_M) - c_m)}{1 + \exp(U_M(X_m, X_w; \theta_M) - c_m)}.$$

The preference parameters in this model can be recovered using a fixed effects logit estimator. Note that (8) is fully consistent with the equilibrium search rule in the Adachi model. In particular, the reservation values, which are the endogenously determined objects in the model, are controlled for using fixed effects. Thus, the model allows us to recover the structural parameters of the utility function (7).

We estimate the model with quadratic differences ($\alpha = 2$). Alternatively, if the attribute differences enter the utility function in linear form ($\alpha = 1$), the fixed effects logit model is not identified.¹² In order to check the sensitivity of our results with respect to this functional form assumption, we estimated a random effects probit model and compared the probit results with $\alpha = 1$ to the fixed effects logit results with $\alpha = 2$. Qualitatively, the results were very similar.

B. Accounting for Costly Communication and Strategic Behavior

As discussed in Section E, strategic behavior may bias our preference estimates. If there is a positive e-mailing cost k and a cost r from rejection, user behavior is described by the threshold crossing condition

$$U_M(m, w) \geq v_M(m) - r + \frac{k + r}{P_W(m, w)}.$$

To account for the effect of $P_W(m, w)$ on man m 's decision to send a first-contact e-mail, we estimate $P_W(m, w)$ directly from the data. To that end, we utilize the information contained in

¹¹ Formally, $|a - b|_+ = \max(a - b, 0)$ and $|a - b|_- = \max(b - a, 0)$.

¹² To see this, note that $x_w - |x_w - x_m|_+ + |x_w - x_m|_- = x_m$. Suppose the estimated fixed effect for man m is c_m . Let $e_k = (0, \dots, 1, \dots, 0)$ be a vector with 1 as the k^{th} component. Choose some arbitrary number α . Then the parameter vectors

$$\begin{aligned} \hat{\beta}_M &= \beta_M + (a/x_{mk}) e_k, \\ \hat{\gamma}_M^+ &= \gamma_M^+ - (a/x_{mk}) e_k, \\ \hat{\gamma}_M^- &= \gamma_M^- + (a/x_{mk}) e_k, \end{aligned}$$

and the fixed effect $\hat{c}_m = c_m + a$ will yield the same utility function as the one parametrized by $(\beta_M, \gamma_M^+, \gamma_M^-, c_m)$.

our data on whether woman w sent a reply to m 's first contact. In particular, we use a binary logit specification to model the probability of receiving a reply conditional on own and partner attributes. These attributes enter the reply probability using the utility function (7).¹³ We then substitute the reciprocal of the predicted reply probability as an additional term into the first-contact specification:

$$(9) \quad \Pr\{m \text{ contacts } w \mid m \text{ browses } w\} = \frac{\exp(U_M(X_m, X_w; \theta_M) - c_m - \frac{k+r}{\hat{P}_w(m, w)})}{1 + \exp(U_M(X_m, X_w; \theta_M) - c_m - \frac{k+r}{\hat{P}_w(m, w)})}.$$

We also experimented with using $\hat{P}_w(m, w)$ directly as a covariate, but this did not change our results.¹⁴ Note that in the specification above, the rejection cost r and e-mailing cost k are not separately identified, since r is absorbed into the fixed effect c_m . Hence, if $r > 0$, the fixed effect c_m will be a biased estimate of the reservation value $v_M(m)$. However, we can use the estimate of $k+r$, which is as an upper bound for r , to assess the importance of this bias.

We use the bootstrap to account for the effect of the estimation error in $\hat{P}_w(m, w)$ on the preference estimates. We choose 250 bootstrap samples, each of which contains resamples of both the first-contact and reply data. To obtain a resample of the first-contact data, we draw (with replacement) N user names from the original set of users, $\{1, \dots, N\}$, and then add the observed first-contact decisions of the randomly selected users to the resampled dataset. For each resample, we first estimate the probability of receiving a reply to a first contact, and then the preference parameters in the first-contact model (9). As bootstrap standard errors, we report the standard deviation across the 250 resampled parameter estimates centered around the point estimates.

C. Preliminary Evidence on Strategic Behavior

Before presenting the estimation results, we examine whether there is any direct evidence pointing toward strategic behavior in the data. We focus on decisions based on physical attractiveness, as we expect that strategic behavior would be most prevalent with regard to looks. In particular, we investigate how a user's propensity to send an e-mail is related to the attractiveness of a potential mate, and whether this propensity is different across attractive versus unattractive searchers.

We first construct a choice set for each user that contains all potential mates this user browses. We then construct a binary variable to indicate the choice of sending an e-mail. Our basic regression specification is a linear probability model of the form

$$(10) \quad e\text{-mail}_{ij} = \sum_{d=1}^{10} \beta_d \cdot \mathbb{I}\{\text{attractiveness}_j = d\} + u_i + \varepsilon_{ij}.$$

¹³ To be precise, the probability that m receives a reply from w is determined by the utility function $U_W(X_w, X_m)$, i.e., the preference of a woman with attributes X_w for a man with attributes X_m .

¹⁴ In Hitsch, Hortaçsu, and Ariely (2009), we perform an extra robustness check on this procedure. The selection correction literature points out that if one can find a variable that is included in the selection equation, but excluded from the outcome equation, then one can achieve semiparametric identification of the outcome equation without relying on the correct functional form specification for the selection correction term. In our application, this means that we have to find a variable that affects the reply probability of a potential mate, but does not enter into the utility from matching with that potential mate. (Moreover, this variable should be observable by the person sending the first-contact e-mail.) In Hitsch, Hortaçsu, and Ariely (2009), we experiment with several variables that may plausibly satisfy an exclusion restriction, and find that our utility estimates are unchanged.

Here, $e-mail_{ij} = 1$ if browser i sends an e-mail to mate j ; $attractiveness_j = d$ denotes that mate j is in the d^{th} decile of attractiveness ratings; \mathbb{I} is an indicator function; and u_i is a person-specific fixed effect, i.e., the optimal search threshold for sending an e-mail to mate j .¹⁵

The physical attractiveness measure is used as a proxy for the overall attractiveness of a profile. We run the regression (10) separately for users in different groups of physical attractiveness, i.e., we segment the suitors, indexed by i , according to their physical attractiveness, and allow for the possibility that users in different groups respond differently to the attractiveness of the mates that they browse. Figure 1 shows the relationship between a browsed user's photo rating and the estimated probability that the browser will send a first-contact e-mail. We see that, regardless of the physical attractiveness of the browser, the probability of sending a first-contact e-mail in response to a profile is monotonically increasing in the attractiveness of the photo in that profile. Thus, even if unattractive men (or women) take the cost of rejection and composing an e-mail into account, this perceived cost is not large enough such that the net expected benefit of hearing back from a very attractive mate would be less than the net expected benefit of hearing back from a less attractive mate.

These results suggest that there are no significant costs of e-mailing attractive users, and, consequently, that strategic behavior is of little importance in online dating. Note, however, that the evidence is not ultimately conclusive, in that multiple attributes enter into the perceived attractiveness of a given profile, while we focus only on a single dimension, physical attractiveness. Our estimation results below will be robust to positive communication and rejection costs, using the method outlined before in Section B.

D. Sample and Attribute Selection

Because we are not interested in the preferences and matching patterns of users who joined the site to find a partner for "casual sex," we keep only those users who state that they are single, divorced, or describe themselves as "hopeful" in the estimation sample. Furthermore, we eliminate users who indicate that they joined the site to start a "casual" relationship. This left us with users who want to "start a long-term relationship" (59 percent), are "just looking/curious" (28 percent), "like to make new friends" (9 percent), and users who state that "a friend put me up on this" (4 percent). The latter three groups are likely to comprise users who want to sound less committal than users who state that they are "hoping to start a long-term relationship." The estimation results and the predicted matching patterns in Sections IV and V change little if we keep only the users in our sample who explicitly state that they want to start a long-term relationship.¹⁶

Given the large number of observed characteristics, we can employ only a subset of all potential variables and potential variable interactions in the estimated model. To guide us in this choice, we use the results from regressions relating the number of unsolicited first-contact e-mails that a user receives in a given time period to his or her attributes.¹⁷ Note that if mate preferences were homogeneous, these regressions would reveal the (common) utility index by which users rank their potential mates.¹⁸ This approach can be extended to allow for some preference heterogeneity by considering only first-contacts received from users who are homogeneous along some observed attribute. The outcome regressions include all observed attributes. For the discrete

¹⁵ Conditional logit estimates yield similar results.

¹⁶ For the full sample, where we also included the users who may be seeking casual affairs, many parameter estimates were smaller in absolute value. The online behavior of these users appears "less focused" than the behavior of the site members who try to find a long-term partner.

¹⁷ The results are available from the authors upon request.

¹⁸ For this statement to be true, we also need to assume that users are uniformly sampled.




FIGURE 1. EVIDENCE FOR/AGAINST STRATEGIC BEHAVIOR

Notes: The graphs display the results of OLS regressions where the dependent variable is an indicator variable for whether a user sends a first-contact e-mail after browsing the profile of a potential mate. The independent variables are indicators for the photo rating of the user being browsed. The regressions control for browser fixed effects. The vertical axis plots the estimated mean probability of sending a first-contact e-mail to a browsed profile, while the horizontal axis indicates the photo rating of the browsed profile. The regressions were estimated separately for different groups of suitors. The first group comprises users who fall within percentiles 1–20 of the photo ratings distribution within their gender, etc. The estimates shown are from a sample of users in the 30–39 age range.

choice model, we select user characteristics that are of interest to us and are quantitatively (and statistically) significant in the outcome regressions.

TABLE 3—BINARY LOGIT ESTIMATES

	Preference of men				Preference of women			
	(1)		(2)		(3)		(4)	
	Estimate	SE	Estimate ^a	SE ^a	Estimate	SE	Estimate ^a	SE ^a
Age	-0.0598	0.0023	-0.0605	0.0041	-0.0098	0.0034	-0.0095	0.0077
Age difference (+)	-0.0007	0.0002	-0.0007	0.0004	-0.0016	0.0002	-0.0016	0.0006
Age difference (-)	-0.005	0.0001	-0.0051	0.0003	-0.0063	0.0004	-0.0064	0.0011
Single; mate divorced ^b	-0.0461	0.0231	-0.0446	0.0273	-0.0718	0.0316	-0.0688	0.033
Both divorced	0.0959	0.0275	0.0961	0.0285	0.1728	0.0305	0.1789	0.0392
Both "long term"	0.0177	0.0178	0.0191	0.0199	0.2388	0.0258	0.2398	0.0322
Both have children	0.1874	0.0271	0.187	0.0532	0.2039	0.0298	0.1973	0.0366
Neither has children	-0.2649	0.0224	-0.264	0.0333	-0.3636	0.0334	-0.3681	0.0423
Has photo	-0.0657	0.0341	-0.0623	0.0522	0.1318	0.0457	0.1365	0.0576
Looks rating	0.5604	0.0144	0.5631	0.0201	0.5848	0.0211	0.5842	0.0269
"Very good" looks	0.5719	0.0396	0.5763	0.0545	0.5516	0.0555	0.5578	0.0688
"Above average" looks	0.2738	0.0363	0.2773	0.0412	0.1733	0.0495	0.1761	0.0627
"Other" looks	0.1742	0.2044	0.1682	0.2096	0.0842	0.2073	0.0519	0.2263
Height	-0.1421	0.0066	-0.1423	0.0101	0.1831	0.0093	0.1826	0.0149
Height difference (+)	-0.0018	0.0037	-0.0044	0.0095	-0.0096	0.0006	-0.0098	0.0011
Height difference (-)	-0.0099	0.0005	-0.0099	0.0008	-0.0227	0.0093	-0.0296	0.0186
BMI	-0.3962	0.028	-0.3932	0.0474	0.1332	0.0499	0.1354	0.0618
BMI ²	0.0043	0.0006	0.0042	0.0009	-0.0007	0.001	-0.0006	0.0013
BMI difference (+)	0.0034	0.0008	0.0034	0.0011	-0.0103	0.0008	-0.0108	0.0013
BMI difference (-)	-0.0101	0.0005	-0.01	0.0012	0.0022	0.0009	0.0025	0.0011
Education (years)	-0.0031	0.0056	-0.0037	0.0067	0.047	0.0076	0.0472	0.0095
Education difference (+)	-0.0039	0.001	-0.0039	0.0011	-0.0086	0.0012	-0.0087	0.0016
Education difference (-)	-0.0026	0.0008	-0.0027	0.001	-0.0022	0.0013	-0.0021	0.0016
Income (\$ 1,000)	0.0053	0.0012	0.0054	0.0013	0.0164	0.0029	0.0163	0.0031
Income (>50) ^c	-0.0027	0.0019	-0.0028	0.0019	-0.0062	0.0035	-0.006	0.0035
Income (>100) ^c	-0.0047	0.0021	-0.0046	0.0021	-0.0082	0.0016	-0.0082	0.0016
Income (>200) ^c	-0.0018	0.0034	-0.0018	0.0037	0.0074	0.0018	0.0075	0.0019
Income difference (+)	6.31E-06	4.07E-06	6.01E-06	4.21E-06	-1.20E-05	3.15E-06	-1.28E-05	3.90E-06
Income difference (-)	1.17E-08	2.53E-06	-5.11E-08	3.39E-06	1.04E-05	6.00E-06	1.21E-05	6.73E-06
Income "Only accountant knows"	0.3332	0.0453	0.3349	0.0516	1.0913	0.1285	1.085	0.1418
Income "What, me work?"	0.2838	0.0542	0.2825	0.0541	0.7155	0.1439	0.7064	0.1564

E. Estimation Results

Table 3 presents the maximum likelihood estimates of the fixed effects binary logit models. Columns 1 and 3 show the results for men and women under the assumption that the first-contact and rejection costs are zero. Columns 2 and 4 show the results for the more general model, where we introduce the estimated (inverse) probability of receiving a reply as a correction term.¹⁹ The estimate of $k + r$ (the coefficient on the reciprocal of the reply probability) is small and statistically insignificant, both for men and for women. Correspondingly, the preference coefficient estimates barely differ across the two model versions. These results, together with the previous findings in Section C, provide strong evidence that strategic behavior due to e-mailing or rejection costs is of little importance in the online dating market studied in this paper.

Before exploring the matching predictions implied by the preference estimates, we provide a brief discussion of the results. For a comparison to alternative estimation approaches, robustness

¹⁹ In these columns, we report the means and standard deviations across 250 bootstrap estimates, as discussed in Section B.

TABLE 3—BINARY LOGIT ESTIMATES (*Continued*)

	(1)		(2)		(3)		(4)	
	Estimate	SE	Estimate ^a	SE ^a	Estimate	SE	Estimate ^a	SE ^a
White; mate black	-0.8301	0.0861	-0.831	0.1051	-0.743	0.1195	-0.7426	0.1529
White; mate Hispanic	-0.2821	0.0367	-0.2873	0.04	-0.5752	0.0897	-0.5749	0.0924
White; mate Asian	-0.4952	0.0436	-0.4983	0.0604	-1.5952	0.2408	-1.6153	0.2854
White; mate other	-0.135	0.0375	-0.1397	0.0408	-0.5677	0.0742	-0.5624	0.0806
Black; mate white	-0.235	0.3701	-0.2214	0.5134	-1.5937	0.3806	-1.1607	0.4257
Black; mate Hispanic	-0.2358	0.4211	-0.2251	0.4657	-1.6185	0.8779	-2.7724	2.5201
Black; mate Asian	-0.6856	0.4609	-0.6981	0.5075				
Black; mate other	0.1764	0.4215	0.1793	0.5399	-0.8192	0.5738	-0.9328	0.8192
Hispanic; mate white	-0.3843	0.1436	-0.351	0.19	-0.6522	0.2303	-0.4896	0.2645
Hispanic; mate black	-0.3787	0.3549	-0.6907	0.6551	-0.8487	0.5082	-0.6407	0.5446
Hispanic; mate Asian	-0.3161	0.2548	-0.2811	0.2799				
Hispanic; mate other	-0.1886	0.2058	-0.1591	0.2493	-0.6777	0.3829	-0.5726	0.3771
Asian; mate white	-0.4617	0.3055	-0.3412	0.3569	-0.0291	0.4627	0.284	0.4246
Asian; mate black					-0.7563	0.9058	-0.4601	0.738
Asian; mate Hispanic	-0.0645	0.421	-0.0475	0.3277	-0.4781	0.5994	-0.228	0.4573
Asian; mate other	0.0383	0.4442	0.1108	0.5107	-0.374	0.5701	-0.1002	0.5644
Same religion	0.1792	0.0218	0.1799	0.0236	0.2918	0.0264	0.2846	0.0306
1/Pr(get reply)			0.0008	0.0007			0.0333	0.0763
Log-likelihood	-72,073.70		-72,093.10 (2,401.7)		-48,998.90		-49,041.40 (1,434.4)	
Observations	242,478				196,363			
Individuals	3,004				2,783			

Notes: The dependent variable is the 0/1 choice to contact another previously “browsed” user. The model includes fixed effects for each individual making first-contact decisions. The results are based on a subsample of users who state that they are single, divorced, or describe themselves as “hopeful.” Furthermore, we eliminate users who indicate that they joined the site to start a “casual” relationship and keep only users who want to “start a long-term relationship,” are “just looking/curious,” “like to make new friends,” and users who state that “a friend put me up on this.” In the full sample we observe more choices for men than for women. In order to make the sample sizes for both genders similar, we took a random sample of the men’s choices (i.e., we kept all men, but randomly discarded some of their observed choices).

To account for the possibility of strategic behavior, we include the reciprocal of the probability of receiving a reply to a first-contact e-mail as an additional variable (in columns 2 and 4). The probabilities of receiving a reply were estimated in a first step using a binary logit model that includes own and partner attributes as explanatory variables. To account for the standard error of the reply probability variable in the first-contact estimates (step two), we employ bootstrapping. We perform 250 bootstrap replications, and report the means and standard deviations of the parameter estimates across these replications. We also report the mean of the log-likelihood across all bootstrap replications and the standard deviation in parentheses.

^aMeans and standard deviations across 250 bootstrap samples.

^bThe user who makes the first-contact choice is single, and the potential mate is divorced.

^cIncome ($> x$) is the amount of income (in thousands of dollars) above the income level x .

checks, and a discussion of the implied trade-offs among various attributes, we refer the reader to Hitsch, Hortacsu, and Ariely (2009). Our results are similar to the revealed preference findings of Kurzban and Weeden (2005), Fisman et al. (2006, 2008), and Eastwick and Finkel (2008) using speed dating experiments, and survey studies in psychology (Nancy Etcoff 1999; David M. Buss 1995).

As expected, we find that the users of the dating service prefer a partner whose age is similar to their own. Women who are single try to avoid divorced men, while divorced women have a preference for a partner who is also divorced. Similarly, single men avoid divorced women, but divorced men have no particular preference to date a divorced woman. Both men and women who have children prefer a partner who also has children. Members with children, however, are much less desirable to both men and women who themselves do not have children. Also, women seeking a long-term relationship prefer a partner who indicates that he is also seeking a long-term relationship, but no such preference is apparent for men. Looks and physique are important deter-

minants of preferences for both men and women. The utility weights on the looks rating variable differ only slightly across men and women. Also, men and women have a stronger preference for mates who describe their looks as “above average” than for average looking members, and they have an even stronger preference for members with self-described “very good looks.” Regarding height, we find that men typically avoid tall women, while women have a preference for tall men. Men have a strong distaste for women with a large BMI, while women tend to prefer heavier men. The estimates of income preferences show that women place about twice as much weight on income than men. There is little evidence for preference heterogeneity here—the absolute value of the distance coefficients is small, and hence own income matters only slightly in the evaluation of a partner’s earnings. Regarding education, we find that both men and women want to meet a partner with a similar education level. While women have an overall strong preference for an educated partner, but also have a relatively small tendency to avoid men who are more educated than themselves, men generally shy away from educated women. The estimated same-race preferences show that both men and women have a preference for a partner of their own ethnicity. Finally, we find that both men and women have a preference for a partner of the same religion.

IV. The Structure and Efficiency of Online Matches

The purpose of an online dating site is to arrange dates and ultimately matches—marriages in particular. As described in Section I, the particular site we investigate uses a very decentralized mechanism, in which browsers communicate and interact without any intervention. In this section, we investigate the structure and efficiency of the matches that result from this mechanism. In particular, we discuss whether we can predict the resulting structure of matches using a well-known economic model of match formation—the Gale-Shapley (1962) deferred-acceptance algorithm. We choose the Gale-Shapley model, not because it provides a literal description of the market, but because it allows us to easily predict the equilibrium matches of the Adachi (2003) two-sided search and matching model with negligible search costs, which provides a more realistic description of online dating.

A. Definition of a “Match”

Our estimation strategy in Section III was based on the initial first-contact e-mail decisions, which may or may not lead to a match. To investigate matching patterns, we first have to define what we mean by a “match” in our data. Since we can track only the users’ online behavior, we do not know whether two partners ever met offline or eventually got married. However, our data do allow us to observe whether users exchange a phone number or e-mail address, or whether an e-mail contains certain keywords or phrases such as “get together” or “let’s meet.” We therefore have some indirect information on whether the online meeting resulted in an initial match, i.e., a date between the users. We therefore define a match as a situation where both mates exchange such contact information (i.e., for a match it is not enough for a man to offer his phone number; we also require that the woman respond by sending her contact information). For some users we observe multiple matches. We use each of these as a separate event when we describe the sorting patterns below.

Matches are very different events from first-contact decisions: Table 4 shows that of all first-contact e-mails sent by men, 4.3 percent lead to an eventual match. Typically, a match is achieved only after multiple rounds of e-mail exchanges: the median and mean numbers of e-mails sent back and forth before a match is realized for conversation initiated by men are 6 and 11.6, respectively. Figure 2 provides detailed information on the distribution of the number of e-mail exchanges that took place until a match was achieved.

TABLE 4—USER BEHAVIOR SUMMARY STATISTICS

	Men	Women
Users	3,004	2,783
<i>First-contact behavior</i>		
Profiles browsed	385,470	172,946
First-contact e-mails (Percentage of browses)	49,223 12.7	14,178 8.2
<i>Matching</i>		
First contacts that lead to match (Percentage of first contacts)	2,130 4.3	914 6.4
<i>E-mails exchanged until match is achieved</i>		
Mean	11.6	12.6
Median	6	6
SD	22.8	26.3

Notes: The summary statistics apply to the sample of users employed in the estimation and matching sections of this paper. In particular, we report only summary statistics on user interactions within this sample. (The binary logit estimates reported in Table 3, however, are based on all observations where user *A* browses or contacts user *B*, even if user *B* is not a member of the subsample.)




FIGURE 2. NUMBER OF E-MAILS EXCHANGED UNTIL A MATCH IS ACHIEVED

Because matches and first contacts are different events, our sorting predictions below are out-of-sample predictions, not in-sample predictions. Remember that we used data on first-contacts only to estimate mate preferences; the estimation did not use any information on a match taking place.

B. Sorting: Actual and Predicted

Correlations in the attributes of couples have been studied widely in sociology, psychology, and economics. We find that matching outcomes in online dating also exhibit strong sorting

TABLE 5—ATTRIBUTE CORRELATIONS IN OBSERVED AND PREDICTED ONLINE MATCHES

	Observed online matches			Simulated Gale-Shapley matches from estimated preferences	
	All (1a)	Users who became inactive after 14 days (1b)	Users who became inactive after 21 days (1c)	Men-optimal (2a)	Women-optimal (2b)
Age	0.70 (0.01) [3,034]	0.71 (0.02) [479]	0.71 (0.02) [830]	0.70 (0.01) [1,769]	0.70 (0.01) [1,769]
Height	0.15 (0.02) [3,034]	0.13 (0.04) [479]	0.12 (0.03) [830]	0.25 (0.02) [1,769]	0.25 (0.02) [1,769]
BMI	0.13 (0.02) [3,034]	0.12 (0.06) [479]	0.12 (0.04) [830]	0.18 (0.04) [1,769]	0.18 (0.04) [1,769]
Looks rating	0.31 (0.02) [1,729]	0.22 (0.06) [265]	0.25 (0.04) [444]	0.16 (0.04) [482]	0.16 (0.04) [482]
Income	0.15 (0.05) [670]	0.10 (0.11) [98]	0.17 (0.09) [189]	0.10 (0.06) [422]	0.10 (0.06) [422]
Education	0.12 (0.02) [3,034]	0.14 (0.05) [479]	0.14 (0.04) [830]	0.15 (0.02) [1,769]	0.15 (0.02) [1,769]

Notes: The table displays Pearson correlation coefficients between mate attributes. Column 1a reports the matching patterns observed in the keyword matches in our online dating sample. Bootstrap standard errors are reported in parentheses. Columns 1b and 1c display the results from those keyword matches that were achieved at least 8 weeks prior to the end of the observation period in our sample, and where both partners became inactive on the site within 14 or 21 days after the match date. Columns 2a and 2b report attribute correlations in the predicted Gale-Shapley matches, based on the logit preference estimates. To simulate the matches, we first drew preference and reservation threshold parameters from the asymptotic distribution of the estimates, simulated the random utility terms, and constructed the preference orderings for each user. We then ran the Gale-Shapley algorithm. We repeated this process 100 times, and report the average and standard deviation of the attribute correlations across these 100 repetitions. The figures in square brackets indicate the median number of matches on which the correlations are based. For some users we do not have looks or income information; therefore, the matches involving these users are not included in the reported attribute correlation numbers.

patterns. Table 5, column 1a, shows the correlation of several user attributes in the observed matches. Age is strongly correlated across men and women ($\rho = 0.70$). Looks, as measured by the standardized photo rating, are also strongly correlated ($\rho = 0.31$). There are smaller but positive correlations in height ($\rho = 0.15$), BMI ($\rho = 0.13$), income ($\rho = 0.15$), and years of education ($\rho = 0.12$). In order to check the sensitivity of these results to the exact definition of a match, we also examine a subset of the keyword matches, where a match was formed at least 8 weeks before the end of the observation period, and both matched users became inactive on the dating site within 14 or 21 days after the date of their initial match. This more stringent match definition attempts to capture users who started dating offline and therefore stopped using the online dating site. Columns 1b and 1c show that the correlation results from these subsets of matches largely agree with the original results based on all keyword matches.

We report results using an alternative sorting measure in Table 6. The first two columns displays the means and standard deviations of the attribute differences between two matched partners (the differences are defined as $x_m - x_w$, where m is the man and w is the woman in a given match). We find that in a typical match, men are 1.3 years older, 5.9 inches taller, have a larger BMI, and a lower looks rating than women. Also, men earn \$49,000 more in annual income,

TABLE 6—ATTRIBUTE DIFFERENCES IN OBSERVED AND PREDICTED ONLINE MATCHES

	Observed online matches		Simulated Gale-Shapley matches from estimated preferences	
	Mean difference	SD of difference	Mean difference	SD of difference
Age	1.3 (0.11)	6.34 (0.08)	1.26 (0.15)	7.48 (0.12)
Height	5.91 (0.06)	3.38 (0.04)	6.05 (0.07)	3.3 (0.05)
BMI	4.28 (0.06)	3.36 (0.09)	2.27 (0.18)	6.77 (0.36)
Looks rating	-0.22 (0.01)	0.64 (0.01)	-0.32 (0.02)	0.7 (0.02)
Income	48.96 (2.44)	65.41 (2.81)	22.77 (1.25)	51.95 (1.92)
Education	0.49 (0.06)	3.3 (0.05)	0.18 (0.05)	3.06 (0.03)

Notes: The table displays means and standard deviations of the differences between men's and women's attributes for the observed keyword matches and the predicted Gale-Shapley matches (see the note in Table 5 for the exact details of the simulation procedure). Bootstrap standard errors for the observed matches are reported in parentheses. Standard errors across 100 simulations are reported for the predicted matches.

and have achieved half a year more in schooling. There is, however, much heterogeneity in these attribute differences. The standard deviation of the age differences, for example, is 6.3 years, and the standard deviation of the education differences is 3.3.

Having documented the sorting patterns in our data, we turn to the question of whether these patterns can be explained by an economic matching model and the mate preferences of the site users. However, it is difficult to provide a literal model of user interactions and match formation on the dating site. Adachi's search equilibrium model outlined in Section A provides a useful stylized description of search behavior, but there are many aspects of behavior in this market that are not captured by this model.

We thus pursue a different strategy, and compare the observed matching outcomes to the matching outcomes that would have been generated by a benchmark mechanism under the preference profiles we estimated. The benchmark in this case is the Gale-Shapley mechanism, which yields stable and efficient (within the class of stable matches) matching outcomes which we can compare to the observed matches. Aside from providing a descriptive vantage point, running the Gale-Shapley algorithm using the estimated preference profiles helps us answer an interesting market design question: given the revealed preferences of the users, would significant utility gains have been achieved if a centralized matchmaking algorithm (possibly the Gale-Shapley algorithm) had been used by the Web site? A lively debate on a similar issue exists in the online dating industry: while some major dating services, such as Match.com and Yahoo!Personals, use a decentralized approach allowing users to freely browse and communicate, a "planner" approach is pursued by other major dating sites. For example, eHarmony.com does not allow users to search and communicate freely, but instead limits users to choose among a subset of all potential partners selected by the company's proprietary algorithm.

To compute the Gale-Shapley matches, we use the estimated preference profiles of the site users (Table 3) and run the deferred-acceptance algorithm for both geographic markets. We calculate both the man-optimal and woman-optimal stable solutions to assess whether the extreme points on the lattice of stable matches yield significantly different matching outcomes. The exact simulation algorithm is as follows:

- (i) Draw preference parameters and reservation thresholds from their estimated asymptotic distribution;
- (ii) Draw random utility terms, ϵ_{mw} , for each pair of potential mates;
- (iii) Compute the man-optimal or woman-optimal stable matching using the deferred-acceptance algorithm.

We repeat this process 100 times and report the average and standard deviation of the attribute correlations and differences across these 100 repetitions.

Table 5, columns 2a and 2b, report the attribute correlations across matched couples for the man-optimal and woman-optimal Gale-Shapley matches. The correlations in the man-optimal and woman-optimal predictions are virtually identical, which we also found for the other predicted matches reported below; hence, we will no longer separately report the man-optimal and woman-optimal results.

We find that the predicted correlations largely agree with the observed keyword correlations in column 1a. The predicted and observed age correlations are identical, ($\rho = 0.70$). Although the predicted correlation in looks is somewhat smaller than the looks correlation in keyword matches ($\rho = 0.16$ versus $\rho = 0.31$), this could be caused by the fact that many users in our sample do not post pictures.²⁰ All users report their height and weight, and the correlations predicted by the Gale-Shapley algorithm are slightly higher than the observed correlations. We also get similar correlations in income and education across the keyword and Gale-Shapley matches.

Table 6 also reveals strong agreement between the predicted and actual mean attribute differences and the standard deviation of the attribute differences among matched couples. For example, the predicted mean age difference is 1.26, while the observed mean age difference is 1.3. The largest discrepancies between the predicted and actual difference measures are with regard to BMI and income, where we underpredict the average difference between men and women.

In summary, we observe strong sorting patterns in online dating, and we are largely able to predict the sorting patterns using the estimated mate preferences and the Gale-Shapley algorithm. This is especially reassuring, because our predictions are out of sample and not based on any information about the actually achieved matches.

C. Efficiency

The finding that the observed and predicted attribute correlations and differences are largely similar suggests that the online dating market achieves an approximately efficient matching within the class of stable matches. However, these correlations are aggregate measures that might mask important differences in the allocations obtained by the actual (“decentralized”) and Gale-Shapley (“planner”) protocols. Alternatively, we could measure the difference between actual and predicted matches by comparing the identities of the matched couples. However, since the error terms, ϵ_{mw} and ϵ_{wm} , change the stable matches across each run of the deferred-acceptance algorithm, it is unreasonable to expect a one-to-one correspondence between the two match results. Therefore, we choose a different method to evaluate the efficiency of the observed match outcomes. As before, we construct each user’s preference profile over all potential mates, based on the preference estimates and the simulated error terms ϵ . For each user i , we then record the

²⁰ The looks correlation was computed only when both members of the matched couple had posted photos.

rank of the match that is achieved in the data, and the rank of the match achieved by the Gale-Shapley algorithm.²¹

To implement this method, we first construct each user's preference ordering, based on a draw from the estimated distribution of preferences and the simulated error terms ϵ . This preference ordering assigns a rank to all potential mates. If there are N potential partners, then the user with rank N is the most desirable mate, $N - 1$ the second most desirable mate, and so forth. For each user i , let R_i^{data} be the rank of the match observed in the data, and R_i^{GS} the rank of the match predicted by the Gale-Shapley algorithm.²² The difference in achieved ranks between the Gale-Shapley outcome and the actually achieved match is represented by $\Delta R_i = R_i^{GS} - R_i^{data}$. Let $\Delta \bar{R} = J^{-1} \sum_{i=1}^J \Delta R_i$ be the average difference in ranks over all J users. If $\Delta \bar{R}$ is approximately zero, then a typical user of the online dating site could not have achieved a better outcome had the matches been generated by the Gale-Shapley model. If, on the other hand, $\Delta \bar{R}$ is positive, a centralized market mechanism could have improved, on average, on the allocation achieved by the current market design.

A caveat to the approach just outlined is that the random components of the users' mate preferences are not observed by us, the analysts, which leads to some "bias" in the comparison of actual and Gale-Shapley matches. While not observed to us, the ϵ_{mw} terms are observed by the site users before they form a match. Hence, in the event that m matches with w , $\mathbb{E}(\epsilon_{mw} | m \text{ matches with } w) \geq \mathbb{E}(\epsilon_{mw})$ and vice versa, with m and w exchanged.²³ This means that if we observe a match between m and w in the data, these two partners must have received unusually high utility draws. Therefore, as we draw the random utility components from the unconditional distribution of ϵ_{mw} , we will typically underpredict the rank of a match that was achieved in the data, and thus overpredict ΔR_i and the efficiency gains of the Gale-Shapley protocol.

In order to evaluate the magnitude of this bias, we conduct the following simulation exercise. We first construct the site users' preference orderings based on draws of the preference parameters and the random utility terms, ϵ_{mw} . We then compute a match according to the Gale-Shapley algorithm, and treat the predicted matches as "data." We redraw the ϵ terms, construct the corresponding preference orderings, and compute a new Gale-Shapley match. Obviously, if we calculated the mate preferences according to the original ϵ_{mw} terms from which the simulated data were created, the predicted rank differences would all be zero. Correspondingly, the average difference in achieved ranks, $\Delta \bar{R}$, based on the new, "incorrect" ϵ_{mw} terms is a measure of the bias of our method.

Table 7 Panel A shows the mean, median, and standard deviations of the predicted rank differences across all matches, averaged across 100 Gale-Shapley predictions. The table displays information both on the rank differences in absolute terms and on the rank differences expressed as a percent of the highest achievable rank, $100 \times \Delta R_i / N$. The mean difference in ranks between the Gale-Shapley matches and the observed keyword matches is 64 for men and 57 for women. Expressed in percent of the highest rank achievable (the users' "first choice"), the improvement is 4.6 percent for men and 3.8 percent for women. Although these numbers suggest that the Gale-Shapley protocol could have improved on the outcomes achieved by the dating site, we know from our discussion above that our approach tends to overstate the potential efficiency gains. Table 7 Panel B also shows the predicted rank differences using the simulation exercise where both the "data matches" and the predicted matches are obtained by the Gale-Shapley protocol. The predicted

²¹ If a user achieved more than one match in the data, we record the highest rank within the set of matched mates and the option of staying single.

²² If a user achieved more than one match in the data, we record the highest rank within the set of matched mates and the option of staying single.

²³ This result is a straightforward consequence of the asymptotic equivalence between the matching resulting from the Adachi and Gale-Shapley models. In the Adachi model, a match arises only if the threshold rule (1) is satisfied. Therefore, ϵ_{mw} will necessarily be left-truncated conditional on a match.

TABLE 7—RANK DIFFERENCES OF “KEYWORD” AND SIMULATED MATCH OUTCOMES

	Men			Women		
	Mean	Median	SD	Mean	Median	SD
<i>Panel A. Observed keyword matches</i>						
Rank differences						
Gale-Shapley versus observed match	63.5 (3.5)	13.4 (2.5)	124.8 (6.6)	57.3 (2.9)	12.3 (1.9)	111.4 (5.7)
First choice versus observed match	99.0 (3.5)	51.6 (2.8)	129.7 (6.5)	84.8 (3.1)	42.7 (2.2)	116.1 (5.5)
First choice versus Gale-Shapley	35.5 (1.5)	20.2 (1.2)	46.1 (3.9)	27.5 (1.2)	15.1 (0.8)	37.1 (3.1)
Rank differences as percentage of maximal ranking						
Gale-Shapley versus observed match	4.6 (0.2)	1.0 (0.2)	9.0 (0.5)	3.8 (0.2)	0.8 (0.1)	7.4 (0.4)
First choice versus observed match	7.1 (0.3)	3.7 (0.2)	9.3 (0.5)	5.7 (0.2)	2.8 (0.2)	7.7 (0.4)
First choice versus Gale-Shapley	2.6 (0.1)	1.5 (0.1)	3.3 (0.3)	1.8 (0.1)	1.0 (0.0)	2.5 (0.2)
<i>Panel B. Simulated data: keyword matches generated using Gale-Shapley algorithm</i>						
Rank differences						
Gale-Shapley versus observed match	115.7 (3.4)	41.5 (3.1)	183.5 (5.0)	114.3 (4.1)	42.4 (2.9)	180.0 (6.4)
First choice versus observed match	160.8 (3.5)	89.8 (3.6)	189.6 (4.8)	145.1 (4.1)	76.1 (3.2)	184.1 (6.5)
First choice versus Gale-Shapley	45.0 (1.8)	24.7 (1.1)	65.0 (5.6)	30.8 (1.3)	16.9 (0.8)	42.4 (3.2)
Rank differences as percentage of maximal ranking						
Gale-Shapley versus observed match	8.3 (0.2)	3.0 (0.2)	13.2 (0.4)	7.6 (0.3)	2.8 (0.2)	12.0 (0.4)
First choice versus observed match	11.6 (0.3)	6.5 (0.3)	13.7 (0.3)	9.7 (0.3)	5.1 (0.2)	12.3 (0.4)
First choice versus Gale-Shapley	3.2 (0.1)	1.8 (0.1)	4.7 (0.4)	2.1 (0.1)	1.1 (0.1)	2.9 (0.2)

Notes: The table shows summary statistics on the preference rank difference across several actual and potential match outcomes. The reported means, medians, and standard deviations are averages across 100 simulations, with standard errors in parentheses. All numbers are calculated for site users who achieved at least one “keyword match.” If a user achieved more than one match in the data, we record the highest rank within the set of matched mates and the option of staying single. “First choice” denotes the outcome where a user is matched to his or her top-ranked partner.

improvement in ranks in the simulation exercise is actually larger than the predicted improvement in ranks based on the keyword matches observed in our data. Hence, the bias in our procedure can account for all of the predicted rank improvements achieved by the Gale-Shapley protocol. We conclude that the decentralized search protocol utilized by the dating site does not result in significant efficiency losses compared to what could have been achieved using the Gale-Shapley procedure.

An intuitive rationalization for this finding can be proposed based on Adachi’s asymptotic equivalence result discussed in Section A: it appears that the search frictions in this online dating site are small enough that the decentralized search outcomes approximate the stable matches obtained from the centralized Gale-Shapley procedure. For the purpose of designing online dating sites, it would be interesting to characterize how large search frictions have to be to lead to significant departures from efficiency.²⁴ We leave this question for future research.

²⁴ Along with the cognitive costs of using a particular search interface and site layout, some online dating sites introduce explicit search costs by charging a browsing fee per profile (e.g., rightstuffdating.com).

Table 7 also compares the rank differences between the users' "first choice" (the maximum achievable rank) and either the rank of the observed match or the predicted rank by the Gale-Shapley algorithm. The average difference between the users' first choice and the rank achieved by the Gale-Shapley algorithm is 36 for men and 28 for women; this corresponds to 2.6 and 1.8 percent of the highest achievable rank. It hence appears that the estimated preference profiles give rise to equilibrium matches that are close, on average, to the user's first choice. This is likely a result of the considerable heterogeneity in preferences, which "softens" the competition for mates. To illustrate, consider the extreme situation in which all men and women have identical preferences over the opposite sex. With approximately 1,500 men and women on each side of the market, the resulting perfect assortative matching would lead to the average user being matched to a person who is 750 ranks away from his or her first choice.

V. Predicting "Offline" Sorting

A large empirical literature in sociology, psychology, and economics documents strong assortative mating (homogamy) patterns in the demographic, physical, and socioeconomic characteristics of married couples. Similar sorting patterns along observable attributes have been documented at earlier stages of a relationship: Laumann et al. (1994), using data from the National Health and Social Life Survey, show that the differences in sorting along age, education, and ethnicity between marriages, cohabitations, and short-term relationships are small.²⁵ Here, following Debra L. Blackwell and Daniel T. Lichter (2004), we document sorting patterns among dating, cohabiting, and married couples using data from the National Survey of Family Growth (NSFG), Cycle 6, which contains information on the relationship history of a sample of 7,643 women.²⁶ Table 8 compares the correlations in age and education, and the interethnicity matching patterns observed in the NSFG sample. The matching patterns among the three types of relationships largely agree; the age and education correlations among "dating" couples are even somewhat larger than the corresponding correlations in marriages.²⁷

In the previous section, we saw that the Gale-Shapley algorithm, based on mate preference estimates, predicts the matching patterns among the users of our online dating site well. We now explore if the same approach can also generate some aspects of the sorting patterns observed among couples who are dating (offline) and married couples. Comparing the predictions of the Gale-Shapley model with offline matching patterns is particularly interesting because Gale-Shapley, as the limit of the Adachi model, describes a market without search costs, while such search frictions may play an important role in the formation of marriages.

A. *Sorting Patterns in Marriages*

Table 9, column 1, reports attribute correlations among married couples. To construct this table we utilized the 2000 Census IPUMS 5 percent sample for the two metropolitan areas covered by our online dating data. In order to be able to compare the results to the predicted matches based on

²⁵ See Table 6.4 in Laumann et al. (1994). While married couples are slightly more similar in their ethnic background, couples in a short-term relationship are slightly more similar along age and education.

²⁶ The survey was conducted between 2002 and 2003, and the women sampled were 15 to 44 years of age. We classify women as "dating" if they were never married, do not live with a partner, and, regarding their partner, indicate that they are "going with him or going steady," "going out with him once in a while," are "just friends," or "had just met him." We classify women as "cohabitating" if they were never married and live with a partner of the opposite sex.

²⁷ Note that the correlations from the NSFG (Table 8) cannot be directly compared with the sorting patterns obtained from the Integrated Public Use Microdata Series (IPUMS) data (Tables 9 and 10), as the two samples cover couples within different age ranges and regions.

TABLE 8—SORTING PATTERNS AMONG DATING, COHABITING, AND MARRIED COUPLES

	Dating couples	Cohabiting couples	Married couples		
<i>Panel A. Attribute correlations</i>					
Age	0.81	0.76	0.78		
Education	0.56	0.57	0.54		
Percentage of women					
	White	Black	Hispanic	Asian	Population
<i>Panel B. Sorting along race/ethnicity</i>					
(1) Dating couples					
Matched with:					
White	83.3	2.9	21.9	11.7	51.5
Black	6.2	94.8	10.2	0.0	34.2
Hispanic	7.1	2.3	67.1	0.0	11.3
Asian	3.0	0.0	0.0	88.3	1.5
Other	0.4	0.0	0.8	0.0	1.5
(2) Cohabiting couples					
Matched with:					
White	80.1	7.1	8.9	13.4	56.5
Black	6.1	88.9	5.8	7.3	16.3
Hispanic	11.9	2.8	83.4	26.5	23.8
Asian	0.4	0.0	1.1	52.8	1.7
Other	1.6	1.2	0.9	0.0	1.8
(3) Married couples					
Matched with:					
White	92.4	4.2	17.0	14.6	72.2
Black	1.0	88.5	2.5	0.7	8.0
Hispanic	4.7	3.9	79.4	2.7	14.6
Asian	1.0	1.1	0.3	80.8	3.6
Other	0.8	2.4	0.8	1.2	1.5

Notes: The tables are based on data from the NSFG, Cycle 6, which contains information on the relationship history of a sample of 7,643 women. The survey was conducted between 2002 and 2003, and the women sampled were 15 to 44 years of age. We classify women as “dating” if they were never married, do not live with a partner, and, regarding their partner, indicate that they are “going with him or going steady,” “going out with him once in a while,” are “just friends,” or “had just met him.” We classify women as “cohabitating” if they were never married and live with a partner of the opposite sex.

users in our online dating sample, we use information on men and women only in the 18–55 age range. Note that the cross-sectional nature of census data restricts us to compute the correlations on the stock of marriages, not the flow of newlyweds. However, as discussed above, sorting patterns among dating couples are very similar to sorting patterns between married couples, which suggests that the measurement error from using stock correlations is unlikely to be large.

Married couples exhibit a strong degree of sorting in age ($\rho = 0.85$) and years of education ($\rho = 0.60$). There is less sorting along income ($\rho = 0.11$), which increases somewhat ($\rho = 0.17$) when we consider the correlation in the imputed income which accounts for the potential earnings of spouses who do not participate in the labor force.²⁸ Regarding the empirical correlations in looks, height, and weight, we consulted several empirical studies in sociology and psychology, which report high degrees of correlation among these characteristics as well (height has a ρ between 0.31 and 0.63, weight between 0.08 and 0.32, and looks—measured in a similar manner as in our study—between 0.34 and 0.54). Note that the studies reporting correlations in physical

²⁸ We impute income using a regression that relates the logarithm of a person’s income to years of education, years in the workforce (and the square of this variable), occupation dummies, and ethnicity dummies. We ran the regressions separately for men and women.

TABLE 9—OBSERVED AND PREDICTED ATTRIBUTE CORRELATIONS IN MARRIAGES

	Marriages (1)	Gale-Shapley predictions		
		With random utility (2)	Random utility terms set to zero (3)	Difference terms set to zero (4)
Age	0.85	0.73 (0.01) [3,240]	0.90 (0.02) [434]	0.09 (0.04) [3,888]
Height	0.31–0.63	0.28 (0.02) [3,240]	0.44 (0.10) [434]	0.01 (0.02) [3,888]
Weight (1); BMI (2–4)	0.08–0.32	0.05 (0.06) [3,240]	0.13 (0.19) [434]	–0.01 (0.03) [3,888]
Looks rating	0.34–0.54	0.18 (0.04) [649]	0.61 (0.12) [38]	0.03 (0.04) [788]
Income	0.11	0.13 (0.03) [1,042]	0.33 (0.10) [158]	0.04 (0.03) [1,191]
Imputed income	0.17	0.36 (0.01) [3,240]	0.58 (0.05) [434]	0.09 (0.02) [3,888]
Education	0.60	0.19 (0.02) [3,240]	0.53 (0.05) [434]	0.06 (0.02) [3,888]

Notes: The table displays Pearson correlation coefficients between mate attributes. The age, income, and education numbers in column 1 are from the 2000 IPUMS 5 percent sample, restricted to Boston and San Diego, the metropolitan areas for which we also have online dating data. We restricted the sample to men and women in the 18–55 age range, and transformed the age and education variables to conform with the online dating data. The height and weight entries in column 1 report the range of results obtained by anthropometric studies (no. obs. = 46–984) surveyed by James N. Spuhler (1968). The entries for looks correlations are from Verlin B. Hinsz (1989) and Gillian Stevens, Dawn Owens, and Eric C. Schaefer (1990), who construct the attractiveness of engaged and married couples whose engagement, marriage, and twenty-fifth anniversary announcements were published in local newspapers. Columns 2–4 report attribute correlations in the predicted Gale-Shapley matches, based on the logit preference estimates. To obtain these matches, we reweighted our sample to match the joint distribution of age, education, and ethnicity (separately for men and women in each market) observed in the IPUMS sample. To simulate the matches, we first drew preference and reservation threshold parameters from the asymptotic distribution of the estimates, simulated the random utility terms, and constructed the preference orderings for each user. We then ran the Gale-Shapley algorithm. We repeated this process 100 times to account for randomness in the preference orderings, and report the average and standard deviation of attribute correlations across these 100 repetitions. The figures in square brackets indicate the median number of matches on which the correlations are based. For some users we do not have looks or income information; therefore, the matches involving these users are not included in the reported attribute correlation numbers. Column 2 reports attribute correlations that are obtained in the men-optimal stable matches; the results for the women-optimal matches are almost identical. In column 3, we set the unobserved utility terms to zero. In column 4, we set the difference or distance terms in the utility specification, which account for preference heterogeneity, to zero.

attributes typically use much smaller and more selective samples than the census, and may not reflect the correlations in the metropolitan areas we are considering.

We also observe strong assortative mating or *endogamy* (marrying within a group) along ethnic-racial lines.²⁹ Panel A of Table 10 displays the percentage of men and women within an ethnic group matching with a mate of another ethnicity, and Table 11 illustrates the same

²⁹ In a recent survey of intermarriage patterns, Kalmijn (1998) reports that in the United States “virtually all ethnic subgroups marry within their group more often than can be expected under random mating.”

TABLE 10—OBSERVED AND PREDICTED SORTING PATTERNS ALONG RACE/ETHNICITY

	Percentage of men				Percentage of women			
	White	Black	Hispanic	Asian	White	Black	Hispanic	Asian
<i>Panel A. IPUMS</i>								
Population	70.3	6.7	14.7	6.6	69.9	6.7	14.7	6.6
Matched with:								
White	92.7	16.2	16.0	10.4	94.6	7.9	20.6	23.6
Black	0.4	72.2	1.0	0.4	1.0	87.5	2.1	2.3
Hispanic	3.9	6.2	81.4	2.3	2.9	3.3	75.8	2.1
Asian	2.5	4.0	1.3	86.5	1.0	0.7	1.1	71.6
Other	0.4	1.4	0.4	0.4	0.5	0.7	0.5	0.5
<i>Panel B. Gale-Shapley, random utility terms included</i>								
Population	71.8	6.1	15.3	6.0	70.8	6.0	15.5	7.0
Matched with:								
White	79.4	38.0	30.4	26.9	88.5	38.3	47.7	60.8
Black	3.3	43.0	6.4	9.1	3.6	41.5	5.2	3.6
Hispanic	10.8	14.0	54.9	20.7	5.8	12.7	42.0	13.4
Asian	6.0	4.2	7.7	41.5	1.5	5.2	4.6	21.2
Other	0.5	0.8	0.6	1.9	0.6	2.3	0.5	1.0
<i>Panel C. Gale-Shapley, random utility terms set to zero</i>								
Population	71.8	6.1	15.3	6.0	70.8	6.0	15.5	7.0
Matched with:								
White	86.1	3.4	3.5	–	98.5	5.0	26.7	61.3
Black	1.2	87.0	9.0	–	0.7	79.0	7.3	3.0
Hispanic	5.7	7.9	84.3	–	0.5	6.7	63.1	4.6
Asian	6.5	1.6	3.0	–	0.0	5.0	1.9	27.6
Other	0.5	0.2	0.3	–	0.2	4.3	0.9	3.4
<i>Panel D. Gale-Shapley, race terms set to zero</i>								
Population	71.8	6.1	15.3	6.0	70.8	6.0	15.5	7.0
Matched with:								
White	69.3	64.7	51.4	56.7	77.7	73.3	63.4	67.0
Black	7.8	9.0	7.9	7.2	6.3	7.4	6.3	6.9
Hispanic	15.4	17.6	30.4	23.5	10.4	13.5	22.7	16.9
Asian	7.0	8.1	9.7	11.8	4.7	5.0	7.1	8.4
Other	0.5	0.6	0.5	0.7	0.8	0.8	0.5	0.8

Notes: Panel A reports the intermarriage ethnicity patterns observed in the 2001 IPUMS 5 percent sample, restricted to Boston and San Diego. We also restricted the sample to men and women in the 18–55 age range. The columns display the percentage of men (women) in each ethnic group that are matched with women (men) of the ethnicities given in the rows. Panel B reports the intermarriage patterns that are generated using the (male-optimal) Gale-Shapley procedure, based on the fixed effects logit preference estimates. The simulation results are averaged across 100 runs of the Gale-Shapley algorithm using different preference draws. Panel C reports the marriage patterns if the random utility term is discarded from the preference specification, and in Panel D the race preferences are set to zero.

matching patterns using the odds-ratios of matches. For example, in Table 10 Panel A we see that 72.2 percent of black men are married to a black woman. In Table 11 Panel A the corresponding same-race matching odds ratio is 10.7, which means that black men are 10.7 times more likely to marry a black woman than they would under random matching. As can be seen from the diagonal entries, strong endogamy patterns are also present for Hispanic and Asian men and women.

B. Gale-Shapley Predictions

To predict the equilibrium matches based on our preference estimates, we first reweight our sample of dating site users to match the joint distribution of age, education, and ethnicity in the IPUMS data, separately for men and women in each market. We consider only users in the 18–55

TABLE 11—OBSERVED AND PREDICTED SORTING PATTERNS ALONG RACE/ETHNICITY: ODDS RATIOS

	Percentage of men				Percentage of women			
	White	Black	Hispanic	Asian	White	Black	Hispanic	Asian
<i>Panel A. IPUMS</i>								
Matched with:								
White	1.3	0.2	0.2	0.1	1.3	0.1	0.3	0.3
Black	0.1	10.7	0.1	0.1	0.2	13.0	0.3	0.3
Hispanic	0.3	0.4	5.6	0.2	0.2	0.2	5.2	0.1
Asian	0.4	0.6	0.2	13.1	0.1	0.1	0.2	10.9
Other	0.2	0.8	0.2	0.3	0.3	0.4	0.2	0.3
<i>Panel B. Gale-Shapley, random utility terms included</i>								
Matched with:								
White	1.1	0.5	0.4	0.4	1.2	0.5	0.7	0.8
Black	0.6	7.2	1.1	1.5	0.6	6.8	0.9	0.6
Hispanic	0.7	0.9	3.5	1.3	0.4	0.8	2.7	0.9
Asian	0.9	0.6	1.1	6.0	0.2	0.9	0.8	3.5
Other	0.6	1.0	0.8	2.3	0.8	3.1	0.7	1.3
<i>Panel C. Gale-Shapley, random utility terms set to zero</i>								
Matched with:								
White	1.2	0.0	0.0	–	1.4	0.1	0.4	0.9
Black	0.2	14.5	1.5	–	0.1	12.9	1.2	0.5
Hispanic	0.4	0.5	5.4	–	0.0	0.4	4.1	0.3
Asian	0.9	0.2	0.4	–	0.0	0.8	0.3	4.6
Other	0.6	0.2	0.3	–	0.3	5.7	1.2	4.6
<i>Panel D. Gale-Shapley, race terms set to zero</i>								
Matched with:								
White	1.0	0.9	0.7	0.8	1.1	1.0	0.9	0.9
Black	1.3	1.5	1.3	1.2	1.0	1.2	1.0	1.1
Hispanic	1.0	1.1	2.0	1.5	0.7	0.9	1.5	1.1
Asian	1.0	1.2	1.4	1.7	0.8	0.8	1.2	1.4
Other	0.6	0.8	0.7	0.9	1.1	1.1	0.6	1.1

Notes: Panel A reports the intermarriage ethnicity patterns observed in the 2001 IPUMS 5 percent sample, restricted to Boston and San Diego. We also restricted the sample to men and women in the 18–55 age range. The columns display odds ratios. On the left-hand side of the table, these ratios are defined as the percentage of men in each ethnic group matched with women of the ethnicities given in the rows relative to the percentage of women of a given ethnicity in the population. The right-hand side of the table shows the corresponding odds ratios for women matching with men. Panel B reports the intermarriage patterns that are generated using the (male-optimal) Gale-Shapley procedure, based on the fixed effects logit preference estimates. The simulation results are averaged across 100 runs of the Gale-Shapley algorithm using different preference draws. Panel C reports the marriage patterns if the random utility term is discarded from the preference specification, and in Panel D the race preferences are set to zero.

age range, who constitute more than 95 percent of all dating site members. Then, as in Section IV, we calculate the equilibrium matches among men and women using the deferred-acceptance algorithm, based on the estimated mate preferences.

Table 9, column 2, displays the predicted sorting patterns. Overall, as in Section IV, we predict sorting along many attributes. The predicted correlations of age, income, and height are similar to the correlations observed in marriages. The predicted correlations in looks, BMI, and in particular education, however, are smaller than the observed ones.

Panel B of Tables 10 and 11 reports the predicted interracial sorting patterns. The results reveal strong endogamy patterns, but we underpredict the full extent to which men and women marry within their own ethnic group.

Note that our predictions are based on the reservation values that we estimated using our online dating data. One may argue that these reservation values would change if applied to the marriage markets in the two metropolitan areas, because the distribution of potential mates and their attributes differs between the online and offline markets (see equation (3)). However, as noted in

Section I, our online sample is fairly representative of the population at large along demographic and socioeconomic attributes. Furthermore, the site users should rationally incorporate the distribution of potentially available mates in the general population into their reservation values, as finding a partner offline is an alternative to matching online. Nevertheless, we performed a sensitivity analysis where we ignored the reservation values and hence assumed that no market participant would want to stay single. The predicted correlations were very similar to the results reported in Table 9.

C. Discussion

Our results demonstrate that the preference estimates of Section III can lead to sorting patterns that are qualitatively, and for some attributes quantitatively, similar to sorting patterns in marriages. This suggests that preferences alone can explain some of the homogamy/endogamy patterns observed in marriages. Conversely, search frictions (or segregation) cannot be the only cause of sorting—since the Gale-Shapley model predicts matching outcomes in a frictionless matching environment. However, we cannot predict the exact degree of sorting for several attributes, such as education and ethnicity. Below, we discuss several reasons for the difference between the actual and predicted homogamy and endogamy patterns.

Selection.—An obvious explanation is that, controlling for observed attributes, the mate preferences of the online dating site users may differ from the preferences of the population at large. We cannot rule out this possibility, but hope that future research will be able to test this hypothesis by matching online dating data with offline marriage outcomes.

Transferable versus Nontransferable Utility.—The Gale-Shapley model is a nontransferable utility framework and hence appears appropriate to describe dating, where explicit transfers do not play a large role. However, such transfers may be important in marriages, and hence marriages may be better described by a transferable utility framework.

Interpretation of the Error Term.—The standard interpretation regards the error term in the utility function (equation (7)) as a random utility component that is observed by the site users, but unobserved to us, the analysts. But, alternatively, this term could simply represent noise in the users' behavior: while browsing through the listings and deciding whom to contact, the site users make occasional mistakes in their choices. Under the latter interpretation, we should not include the error terms in the mate preferences on which the Gale-Shapley simulations are based, because these errors will eventually be corrected before a final match is achieved.

Table 9, column 3, displays the predicted attribute correlations when we disregard the logit error term. The predicted and actual correlations are now very close along the majority of attributes, although we overpredict the income and imputed income correlations. Panel C of Tables 10 and 11 reports the endogamy patterns if we omit the error term. The predicted endogamy patterns are for the most part very similar to the matching patterns observed among married couples. However, we overpredict endogamy for black men and underpredict endogamy for Asian women.³⁰

To summarize these results, if we interpret the econometric error term as representing mistakes that are eventually corrected, we can roughly predict the sorting patterns in marriages based on preferences alone. But, of course, based on available evidence it is difficult to argue which interpretation of the error term is more appropriate.

³⁰ Note that, in every run of the Gale-Shapley algorithm, all Asian men stayed single, and hence we cannot predict the endogamy patterns for them.

Search Frictions.—The observed correlation patterns might be larger than the predicted correlation patterns due to search frictions. This interpretation of our findings implies that search frictions are one cause of sorting in marriages. The different degrees by which we mispredict correlations in marriages support this interpretation. In particular, we strongly underpredict the correlation in education. Education, however, is the one attribute where we might expect that search frictions are most important, given that we spend much of our lifetime surrounded by people with relatively similar education levels. For example, Laumann et al. (1994, Table 6.1) document that 60 percent of all married couples in their sample met in school, at work, at a private party, in church, or at a gym/social club. An interesting consequence, therefore, is that sorting along education could decrease if online dating became a common means of finding a partner.

D. How Do Different Preference Components Contribute to Sorting?

Our approach, which is based on an explicit model of preferences and a theoretical model to predict match outcomes, enables us to answer questions regarding how different aspects of men's and women's mate preferences affect the equilibrium matching patterns. We now present the results from two counterfactual exercises, in which we turn on or shut off different preference components, and use the Gale-Shapley algorithm to evaluate the contribution of these components to equilibrium matching patterns.

Vertical versus Horizontal Preferences.—As discussed in the introduction, sorting can be driven both by preferences over the absolute value of a partner's trait, and by preferences for a partner with similar traits. We will call the first example an instance of "vertical" preference, and the second one with "horizontal" preferences. In order to assess the importance of vertical and horizontal preference components, we set the "distance terms" in the utility specification (7) equal to zero, and then simulate the stable matches under this alternative preference specification. Table 9, column 4, shows that all attribute correlations become small in size. Hence, according to our preference estimates, sorting is largely driven by preference heterogeneity.

Matching in a "Color-Blind" World.—We saw that both actual marriages and the predicted matches from the Gale-Shapley model exhibit strong intra-ethnicity sorting patterns. However, even in the absence of search frictions it is not obvious whether same-race preferences alone cause sorting within ethnicity groups: because other attributes, such as income and education, are correlated with race, preferences over these other attributes could lead to intra-ethnicity sorting as well. To investigate what fraction of the predicted endogamy patterns is driven by same-race preferences, we remove these preferences from the utility specification and simulate the corresponding stable matches. Panel D of Tables 10 and 11 shows that intra-ethnicity sorting is strongly diminished, although small degrees of endogamy remain. These results suggest that racial sorting is mostly due to the same-race preferences; preferences for nonrace attributes can not quantitatively account for the small fraction of marriages across different ethnic groups.

VI. Conclusions

Using a novel dataset from an online dating site, we first estimate mate preferences and then use the classic Gale-Shapley algorithm to predict matching outcomes. Online dating exists to facilitate the search for a partner. Our results suggest that the particular site that we study leads to approximately efficient matching outcomes (within the set of stable matches), and that search frictions are mostly absent. Hence, the site appears to be efficiently designed.

We document strong sorting patterns along various attributes among the matched users of the dating site. Qualitatively, these patterns resemble the assortative mating (homogamy) patterns in marriages. We are largely able to predict these sorting patterns based on the estimated mate preferences and the Gale-Shapley algorithm. Hence, in online dating, assortative mating arises in the absence of search frictions, due primarily to preferences and the specific market mechanism by which matches are formed.

Finally, we use the same approach of predicting match outcomes based on estimated mate preferences to extrapolate to marriages. We first reweight our sample of online daters to match the population at large along observable attributes, and then predict matching outcomes for this “synthetic” sample. The goal of this exercise is to assess how much of the assortative mating patterns in marriages can be attributed to preferences alone. We predict sorting patterns that qualitatively, and for some attributes also quantitatively, are similar to the sorting patterns in marriages. This suggests that preferences are one important cause of sorting in marriages. However, the degree to which we underpredict sorting along some attributes—education in particular—is consistent with search frictions. For example, such search frictions may be due to the institutions by which men and women are mostly exposed to a potential partner with a similar education level.

Based on the results obtained in this study, we conclude that online dating data provide valuable insights toward understanding the economic mechanisms underlying match formation and the formation of marriages. Some important issues are left for future research. For example, an obvious drawback of our analysis is that we cannot observe whether an online meeting finally resulted in a marriage. A recent working paper by Soohyung Lee (2009) attempts to close this gap by utilizing a dataset from a South Korean matchmaking service that tracks participants across a sequence of dates, some of which result in marriages. Lee builds on the identification framework introduced in this paper, and augments our revealed preference specification to account for learning about unobservable match components.

REFERENCES

- Abdulkadiroğlu, Atila, and Tayfun Sönmez.** 2003. “School Choice: A Mechanism Design Approach.” *American Economic Review*, 93(3): 729–47.
- Adachi, Hiroyuki.** 2003. “A Search Model of Two-Sided Matching under Nontransferable Utility.” *Journal of Economic Theory*, 113(2): 182–98.
- Becker, Gary S.** 1973. “A Theory of Marriage: Part I.” *Journal of Political Economy*, 81(4): 813–46.
- Biddle, Jeff E., and Daniel S. Hamermesh.** 1998. “Beauty, Productivity, and Discrimination: Lawyers’ Looks and Lucre.” *Journal of Labor Economics*, 16(1): 172–201.
- Blackwell, Debra L., and Daniel T. Lichter.** 2004. “Homogamy among Dating, Cohabiting, and Married Couples.” *Sociological Quarterly*, 45(4): 719–37.
- Bozon, Michel, and Francois Heran.** 1989. “Finding a Spouse: A Survey of How French Couples Meet.” *Population*, 44(1): 91–121.
- Browning, Martin, Pierre-André Chiappori, and Yoram Weiss.** 2008. “The Economics of the Family.” http://www.tau.ac.il/~weiss/fam_econ/.
- Burdett, Ken, and Melvyn G. Coles.** 1997. “Marriage and Class.” *Quarterly Journal of Economics*, 112(1): 141–68.
- Buss, David M.** 1995. *The Evolution of Desire: Strategies of Human Mating*. New York: Basic Books.
- Chamberlain, Gary.** 1980. “Analysis of Covariance with Qualitative Data.” *Review of Economic Studies*, 47(1): 225–38.
- Choo, Eugene, and Aloysius Siow.** 2006. “Who Marries Whom and Why.” *Journal of Political Economy*, 114(1): 175–201.
- Comscore.** 2006. “Europeans Nearly Fifty Percent More ‘Engaged’ in Online Dating Compared to North Americans.” http://www.comscore.com/Press_Events/Press_Releases/2006/09/European_Online_Dating_Habits

- Eastwick, Paul W., and Eli J. Finkel. 2008. "Sex Differences in Mate Preferences Revisited: Do People Know What They Initially Desire in a Romantic Partner?" *Journal of Personality and Social Psychology*, 94(2): 245–64.
- Etcoff, Nancy. 1999. *Survival of the Prettiest: The Science of Beauty*. New York: Doubleday Books.
- Fisman, Raymond, Sheena S. Iyengar, Emir Kamenica, and Itamar Simonson. 2006. "Gender Differences in Mate Selection: Evidence from a Speed Dating Experiment." *Quarterly Journal of Economics*, 121(2): 673–97.
- Fisman, Raymond, Sheena S. Iyengar, Emir Kamenica, and Itamar Simonson. 2008. "Racial Preferences in Dating." *Review of Economic Studies*, 75(1): 117–32.
- Flinn, Christopher J., and Daniela del Boca. 2006. "Household Time Allocation and Modes of Behavior: A Theory of Sorts." <http://www.econ.nyu.edu/user/flinn/papers/hhsorting4.pdf>.
- Fréchette, Guillaume R., Alvin E. Roth, and M. Utku Ünver. 2007. "Unraveling Yields Inefficient Matchings: Evidence from Post-season College Football Bowls." *RAND Journal of Economics*, 38(4): 967–82.
- Gale, David, and Lloyd S. Shapley. 1962. "College Admissions and the Stability of Marriage." *American Mathematical Monthly*, 69(1): 9–15.
- Hamermesh, Daniel S., and Jeff E. Biddle. 1994. "Beauty and the Labor Market." *American Economic Review*, 84(5): 1174–94.
- Hinsz, Verlin B. 1989. "Facial Resemblance in Engaged and Married Couples." *Journal of Social and Personal Relationships*, 6(2): 223–29.
- Hitsch, Günter J., Ali Hortaçsu, and Dan Ariely. 2009. "What Makes You Click?—Mate Preferences in Online Dating." http://home.uchicago.edu/~ghitsch/Hitsch-Research/Guenter_Hitsch_Files/Mate-Preferences.pdf
- Kalmijn, Matthijs. 1998. "Intermarriage and Homogamy: Causes, Patterns, Trends." *Annual Review of Sociology*, 24: 395–421.
- Kurzban, Robert, and Jason Weeden. 2005. "HurryDate: Mate Preferences in Action." *Evolution and Human Behavior*, 26(3): 227–44.
- Langlois, Judith H., Lisa Kalakanis, Adam J. Rubenstein, Andrea Larson, Monica Hallam, and Monica Smoot. 2000. "Maxims or Myths of Beauty? A Meta-Analytic and Theoretical Review." *Psychological Bulletin*, 126(3): 390–423.
- Laumann, Edward O., John H. Gagnon, Robert T. Michael, and Stuart Michaels. 2000. *The Social Organization of Sexuality: Sexual Practices in the United States*. Chicago: University of Chicago Press.
- Lee, Soohyung. 2009. "Marriage and Online Mate-Search Services: Evidence from Korea." Unpublished.
- Madden, Mary, and Amanda Lenhart. 2006. "Online Dating." Pew Internet & American Life Project. http://www.pewinternet.org/~media/Files/Reports/2006/PIP_Online_Dating.pdf
- Niederle, Muriel, and Alvin E. Roth. 2003. "Unraveling Reduces Mobility in a Labor Market: Gastroenterology with and without a Centralized Match." *Journal of Political Economy*, 111(6): 1342–52.
- Roth, Alvin E. 1984. "The Evolution of the Labor Market for Medical Interns and Residents: A Case Study in Game Theory." *Journal of Political Economy*, 92(6): 991–1016.
- Roth, Alvin E., and Marilda A. O. Sotomayor. 1990. *Two-Sided Matching: A Study in Game-Theoretic Modeling and Analysis*. Econometric Society Monographs Number 18. Cambridge: Cambridge University Press.
- Roth, Alvin E., Tayfun Sönmez, and M. Utku Ünver. 2004. "Kidney Exchange." *Quarterly Journal of Economics*, 119(2): 457–88.
- Spuhler, James N. 1968. "Assortative Mating with Respect to Physical Characteristics." *Social Biology*, 15: 128–40.
- Stevens, Gillian, Dawn Owens, and Eric C. Schaefer. 1990. "Education and Attractiveness in Marriage Choices." *Social Psychology Quarterly*, 53(1): 62–70.
- Wong, Linda Y. 2003. "Structural Estimation of Marriage Models." *Journal of Labor Economics*, 21(3): 699–727.