



MIT Sloan School of Management

MIT Sloan Working Paper 4603-06
February 2006

What Makes You Click? — Mate Preferences and Matching Outcomes in Online Dating

Günter J. Hitsch, Ali Hortaçsu, Dan Ariely

© 2006 by Günter J. Hitsch, Ali Hortaçsu, Dan Ariely.

All rights reserved. Short sections of text, not to exceed two paragraphs, may be quoted without explicit permission, provided that full credit including © notice is given to the source.

This paper also can be downloaded without charge from the
Social Science Research Network Electronic Paper Collection:
<http://ssrn.com/abstract=895442>

What Makes You Click? — Mate Preferences and Matching Outcomes in Online Dating*

Günter J. Hitsch	Ali Hortaçsu	Dan Ariely
<i>University of Chicago</i>	<i>University of Chicago</i>	<i>MIT</i>
<i>Graduate School of Business</i>	<i>Department of Economics</i>	<i>Sloan School of Management</i>

February 2006

Abstract

This paper uses a novel data set obtained from an online dating service to draw inferences on mate preferences and to investigate the role played by these preferences in determining match outcomes and sorting patterns. The empirical analysis is based on a detailed record of the site users' attributes and their partner search, which allows us to estimate a rich preference specification that takes into account a large number of partner characteristics. Our revealed preference estimates complement many previous studies that are based on survey methods. In addition, we provide evidence on mate preferences that people might not truthfully reveal in a survey, in particular regarding race preferences. In order to examine the quantitative importance of the estimated preferences in the formation of matches, we simulate match outcomes using the Gale-Shapley algorithm and examine the resulting correlations in mate attributes. The Gale-Shapley algorithm predicts the online sorting patterns well. Therefore, the match outcomes in this online dating market appear to be approximately efficient in the Gale-Shapley sense. Using the Gale-Shapley algorithm, we also find that we can predict sorting patterns in actual marriages if we exclude the unobservable utility component in our preference specification when simulating match outcomes. One possible explanation for this finding suggests that search frictions play a role in the formation of marriages.

*We thank Babur De los Santos, Chris Olivola, and Tim Miller for their excellent research assistance. We are grateful to Derek Neal, Emir Kamenica, and Betsey Stevenson for comments and suggestions. Seminar participants at the 2006 AEA meetings, the Choice Symposium in Estes Park, Northwestern University, the University of Pennsylvania, the 2004 QME Conference, UC Berkeley, the University of Chicago, the University of Toronto, Stanford GSB, and Yale University provided valuable comments. This research was supported by the Kilts Center of Marketing (Hitsch) and a John M. Olin Junior Faculty Fellowship (Hortaçsu). Please address all correspondence to Hitsch (guenter.hitsch@chicagogsb.edu), Hortaçsu (hortacsu@uchicago.edu), or Ariely (ariely@mit.edu).

1 Introduction

Starting with the seminal work of Gale and Shapley (1962) and Becker (1973), economic models of marriage markets predict how marriages are formed, and make statements about the efficiency of the realized matches. The predictions of these models are based on a specification of mate preferences, the mechanism by which matches are made, and the manner in which the market participants interact with the mechanism. Accordingly, the empirical literature on marriage markets has focused on learning about mate preferences, and how people find their mates. Our paper contributes to this literature using a novel data set obtained from an online dating service. We provide a description of how men and women interact in this dating market, and utilize detailed information on the search behavior of site users to infer their revealed mate preferences. Our data allows us to estimate a very rich preference specification that takes into account a large number of partner attributes, including detailed demographic and socioeconomic information, along with physical characteristics. We use the preference estimates to investigate the empirical predictions of the classic Gale-Shapley model, especially with regard to marital sorting patterns.

The revealed preference estimates presented in this paper complement a large literature in psychology, sociology, and anthropology investigating marital preferences. This literature has yielded strong conclusions, in particular regarding gender differences in marital preferences (see Buss 2003 for a detailed survey of these findings). However, the extent to which these findings on preferences can be used to make quantitative predictions regarding marital sorting patterns has not been explored. Since these studies typically do not provide information on the tradeoffs between different mate attributes, it is difficult to use their results as inputs in an economic model of match formation. Moreover, much of the prior literature utilizes survey methods. Relying on stated rather than revealed preferences might not yield reliable results for certain dimensions of mate choice, such as race preferences.¹

An important motivation to studying marital preferences is to understand the causes of marital sorting. Marriages exhibit sorting along many attributes such as age, education, income, race, height, weight, and other physical traits. These empirical patterns are well documented (see Kalmijn 1998 for a recent survey). However, as pointed out by Kalmijn (1998) and others, several distinct mechanisms can account for the observed sorting patterns, and it is difficult to distinguish between the alternative explanations. For example, sorting on educational attainment (highly educated women date or marry highly educated men) may be the result of a preference for a mate with a similar education level. Alternatively, the same outcome can arise in equilibrium (as a stable matching) in a market in which all

¹In this light, our focus on inferring revealed preferences from the actions of dating site users may be seen as akin to implicit association tests (IATs) used in social psychology to study racial attitudes and stereotyping.

men and women prefer a highly educated partner to a less educated one. The participants in this market have very different preferences than in the first example, and the correlation in education is caused by the market mechanism that matches men and women. Another possible explanation for sorting is based on institutional or search frictions that limit market participants' choice sets. For example, if people spend most of their time in the company of others with a similar education level (in school, at work, or in their preferred bar), sorting along educational attainment may arise even if education does not affect mate preferences at all.²

Online dating provides us with a market environment where the participants' choice sets and actual choices are observable to the researcher.³ Our preference estimation approach relies on the well-defined institutional environment of the dating site, where a user first views the posted "profile" of a potential mate, and then decides whether to contact that mate by e-mail. This environment allows us to use a straightforward estimation strategy 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). Our analysis is based on a data set that contains detailed information on the attributes and online activities of approximately 22,000 users in two major U.S. cities. The detailed information on the users' traits allows us to consider preferences (and sorting) over a much larger set of attributes than in the extant studies that are based on marriage data.

Our revealed preference estimates corroborate several salient findings of the stated preference literature. For example, while physical attractiveness is important to both genders, women have a stronger preference for the income of their partner than men. We also document preferences to date a partner of the same ethnicity. Our estimation approach allows us to examine the preference tradeoffs between a partner's attributes. For example, we calculate the additional income that black, Hispanic, and Asian men need to be as desirable to a white woman as a white man.

In order to examine the quantitative importance of the estimated preferences in determining marital sorting, we simulate equilibrium (stable) matches between the men and women in our sample using the Gale-Shapley (1962) algorithm. The simulations are based on the estimated preference profiles. The Gale-Shapley framework is not only a seminal theoretical benchmark in the economic analysis of marriage markets, but it also provides an approximation to the match outcomes from a realistic search and matching model that resembles the environment of an online dating site (Adachi 2003).

²An analysis of an alumni database of a prestigious West Coast university reveals that 46% of all graduates are married to another graduate of the same school (which could be explained by all three mentioned theories of sorting). — We thank Oded Netzer of Columbia University for pointing out this result to us.

³To be precise, we do not observe the site users' opportunities outside the dating site. However, we observe them browsing multiple alternatives on the site and their choices, which allows us to infer their relative rankings of these potential mates.

Our simulations show that the preferences estimates can explain many of the salient sorting patterns among the users of the dating site. For example, compared to a world with color-blind preferences, the race preferences that we estimate lead to sorting within ethnic groups. Perhaps more surprisingly, our preference estimates, coupled with the Gale-Shapley model, can also replicate sorting patterns in actual marriages quite well when we ignore the idiosyncratic, unobservable error term that is part of our preference specification. One explanation for this finding interprets the error term as “noise” in the users’ behavior: the searchers sometimes make mistakes when they decide who to approach by e-mail. The second explanation interprets the error term as a utility component that is observed by the site users but unobserved to us, the analysts. For example, these utility components could represent personality traits. Finding a partner along such traits may be easier using the technology of online dating than in traditional marriage markets, where—due to search frictions, for example—partner search may be directed along easily observed attributes, such as age, looks, and education.

Most closely related and complementary to our analysis, both in terms of the focus on revealed preferences and the methodological approach, are two studies by Fisman, Iyengar, Kamenica and Simonson (2005, 2006) that utilize data from speed-dating experiments conducted at Columbia University. Their results on gender differences and in particular same-race preferences are remarkably similar to ours, which is especially surprising given the different samples employed in our and their studies (Fisman et al. use a subject pool composed of graduate students). The research design of Fisman et al. has the advantage of eliciting information regarding match-specific components of utility (e.g. the perceived degree of shared interests) that are not observable in our data. In contrast to our work, Fisman et al. do not explore the consequences of their preference estimates for sorting.

Our work is also related to an important literature that estimates mate preferences based on marriage data (Choo and Siow 2006, Wong 2003). In comparison to these papers, our data contains more detailed information about mate attributes; measures of physical traits, for example, are not included in U.S. Census data. Our setting also allows us to observe the search process directly, providing us with information regarding the choice sets available to agents. On the other hand, although we do not find stark differences between the observed characteristics of the dating site users and the general population in the same geographic areas, our sample is not as representative as the samples employed by Choo and Siow (2006) and Wong (2003). Also, by design marriage data are related to preferences over a marriage partner. In contrast, we can only indirectly claim that our preference estimates relate to marriages by examining how well these estimates predict marriage sorting patterns in the general population.

A potential methodological drawback of our estimation approach, compared to Choo and

Siow (2006) and Wong (2003) is that we do not allow for strategic behavior. For example, a man with a low attractiveness rating may not approach a highly attractive woman if the probability of forming a match with her is low, such that the *expected* utility from a match is lower than the cost of writing an e-mail or the disutility from a possible rejection. In that case, his choice of a less attractive woman does not reveal his true preference ordering. A priori, we expect that strategic behavior or fear of rejection should be most pronounced with respect to physical attractiveness. However, our analysis in Section 4 does not reveal much evidence for such strategic behavior. In particular, we find that regardless of their own physical attractiveness rating, users are more likely to approach a more attractive mate than a less attractive mate. We thus believe that the assumption of no strategic behavior is justified, although we cannot ultimately reject the possibility that some strategic behavior is present in the data. Note that the analysis in Choo and Siow (2006) and Wong (2003) is based on final match outcomes only. Such data can be interpreted as choices under an extreme form of strategic behavior, where the market participants choose only their final match partner. The identification of preferences in these papers is achieved through structural assumptions on the market mechanism by which the final matches are achieved; thus the bias introduced by strategic behavior is corrected by an explicit specification of the equilibrium of the matching game and the incorporation of the equilibrium restrictions in the estimation procedure.⁴ Our paper, on the other hand, is based on a comparatively straightforward analysis of choices among potential mates. We believe that both our and the extant approaches have their relative merits, and should be seen as complementary.

The paper proceeds as follows. Section 2 describes the online dating site from which our data were collected, and the attributes of the site users. Section 3 outlines the modeling framework. In Section 4, we address the question of whether users behave strategically. Section 5 presents the preference estimates from our estimation approaches. Section 6 compares the match predictions from our preference estimates with the structure of online matches and actual marriages. Section 7 concludes.

2 The Data and User Characteristics: Who Uses Online Dating?

Our data set contains socioeconomic and demographic information and a detailed account of the website activities of approximately 22,000 users of a major online dating service. 10,721 users were located in the Boston area, and 11,024 users were located in San Diego.

⁴Choo and Siow (2006) estimate a transferable utility model, while Wong (2003) estimates an equilibrium search model of a marriage market. Fox (2006) discusses nonparametric identification in the transferable utility model.

We observe the users' activities over a period of three and a half months in 2003. We first provide a brief description of online dating that also clarifies how the data were collected.

Upon joining the dating service, the users answer questions from a mandatory survey and create "profiles" of themselves.⁵ Such a profile is a webpage that provides information about a user and can be viewed by the other members of the dating service. 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 answer a question on 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 or not, and their political views. All this information is either numeric (such as age and weight) or an answer to a multiple choice question, and hence easily storable and usable for our statistical analysis. The users can also answer essay questions that provide more detailed information about their attitudes and personalities. This information is too unstructured to be usable for our analysis. Many users also include one or more photos in their profile. We have access to these photos and, as we will explain in detail later, used the 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 this detailed profile, the searcher 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.⁶ 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 also 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.⁷

In order to initiate a contact by e-mail, a user has to become a paying member of the

⁵Neither the names nor any contact information of the users were provided to us in order to protect the privacy of the users.

⁶We obtained this information in the form of a "computer log file."

⁷We do not see the full content of the e-mail, or the e-mail address or phone number that was exchanged.

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.

In summary, our data provide detailed user descriptions, and we know how the users interact online. The keyword searches provide some information on the progress of the online relationships, possibly to an offline, “real world” meeting. We now give a detailed description of the users’ characteristics.

Motivation for using the dating service 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 only seek a partner for casual sex. The majority of all users are “hoping to start a long term relationship” (36% of men and 39% of women), or are “just looking/curious” (26% of men and 27% of women). Perhaps not surprisingly, an explicitly stated goal of finding a partner for casual sex (“Seeking an occasional lover/casual relationship”) is more common among men (14%) than among women (4%).

More important than the number is the share of activities accounted for by users who joined the dating service for various reasons. Users who seek a long-term relationship account for more than half of all observed activities. For example, men who are looking for a long-term relationship account for 55% of all e-mails sent by men; among women looking for a long-term relationship the percentage is 52%. The corresponding numbers for e-mails sent by users who are “just looking/curious” is 22% for men and 21% for women. Only a small percentage of activities is accounted for by members seeking a casual relationship (3.6% for men and 2.8% for women).

We conclude that at least half of all observed activities is accounted for by people who have a stated preference for a long-term relationship and thus possibly for an eventual marriage. Moreover, it is likely that many of the users who state that they are “just looking/curious” chose this answer because it sounds less committal than “hoping to start a long-term relationship.” Under this assumption, about 75% of the observed activities are by users who joined the site to find a long-term partner.⁸

Demographic/socioeconomic characteristics We now investigate the reported characteristics of the site users, and contrast some of these characteristics to representative samplings of these geographic areas from the CPS Community Survey Profile (Table 2.1). In

⁸The registration also asks users about their sexual preferences. Our analysis focuses on the preferences and match formation among men and women in heterosexual relationships; therefore, we retain only the heterosexual users in our sample.

particular, we contrast the site users with two sub-samples of the CPS. The first sub-sample is a representative sample of the Boston and San Diego MSA's (Metropolitan Statistical Areas), and reflects information current to 2003. The second CPS sub-sample conditions on being an Internet user, as reported in the CPS Computer and Internet Use Supplement, which was administered in 2001.

A visible difference between the dating site and the population at large is the over-representation of men on the site. 54.7% of users in Boston and 56.1% of users in San Diego are men.⁹ Another visible difference is in the age profiles: site users are more concentrated in the 26-35 year range than both CPS samples (the median user on the site is in the 26-35 age range, whereas the median person in both CPS samples is in the 36-45 age range). People above 56 years are underrepresented on the site compared to the general CPS sample; however, when we condition on Internet use, this difference in older users diminishes.

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.¹⁰

The reported marital status of site users clearly represents the fact that most users are looking for a partner. About two-thirds of the users are never married. The fraction of divorced women is higher than the fraction of divorced men. Interestingly, the fraction of men who declare themselves to be "married but not separated" (6.3% in San Diego and 7.2% in Boston) is larger than women making a similar declaration. However, less than 1% of men's and women's activities (e-mails sent) is accounted for by married people. This suggests that a small number of people in a long term relationship may be using the site as a search outlet. Of course, one may expect the true percentage of otherwise committed people to be higher than reported.

The education profile of the site users shows that they are on average more educated than the general CPS population. However, the education profile is more similar to that of the Internet using population, with only a slightly higher percentage of graduate and professional degree holders.

The income profile reflects a pattern that is similar to the education profile. Site users have generally higher incomes than the overall CPS population, but not compared to the Internet-using population.

These comparisons show that the online dating site attracts users who are typically single,

⁹When we restrict attention to members who have posted photos online (23% of users in Boston and 29% of users in San Diego), the difference between male and female participation decreases slightly. 51% of users with a photo in Boston and 53% of such users in San Diego are men.

¹⁰We should note that we had difficulty in reconciling the "other" category in the site's ethnic classification with the CPS classification and that some of the discrepancy may be driven by this.

somewhat younger, more educated, and have a higher income than the general population. Once we condition on household Internet use, however, the remaining differences are not large. This suggests that during recent years, online dating has become an accepted and widespread means of partner search.

Reported physical characteristics of the users Our data set contains detailed (although self-reported) information regarding the physical attributes of the users. 27.5% post one or more photos online. For the rest of the users, the survey is the primary source of information about their appearance.

The survey asks the users to rate their looks on a subjective scale. 19% of men and 24% of women possess “very good looks,” while 49% of men and 48% of women have “above average looks.” Only a minority—29% of men and 26% of women—declare that they are “looking like anyone else walking down the street.” That leaves less than 1% of users with “less than average looks,” and a few members who avoid the question and joke that a date should “bring your bag in case mine tears.” Posting a photo online is a choice, and hence one might suspect that those users who post a photo are on average better looking. On the other hand, those users who do not post a photo might misrepresent their looks and give an inflated assessment of themselves. The data suggest that the former effect is more important. Among those users who have a photo online, the fraction of “above average” or “very good looking” members is about 7% larger compared to all site users.

The registration survey contains information on the users’ height and weight. We compared these reported characteristics with information on the whole U.S. population, obtained from the National Health and Examination Survey Anthropometric Tables (the data are from the 1988-1994 survey and cover only Caucasians). Table 2.2 reports this comparison. Among women, we find that the average stated weight is less than the average weight in the U.S. population. The discrepancy is about 6 lbs among 20-29 year old women, 18 lbs among 30-39 year old women, and 20 lbs among 40-49 year old women. On the other hand, the reported weights of men are only slightly higher than the national averages. The stated height of both men and women is somewhat above the U.S. average. This difference is more pronounced among men, although the numbers are small in size. For example, among the 20-29 year old, the difference is 1.3 inches for men and 1 inch for women. The weight and height differences translate into body mass indices (BMI) that are 2 to 4 points less than national averages among women, and about 1 point less than national averages among men.

Measured Physical Characteristics of the Users 26% of men (3174 users) and 29% of women (2811 users) post one or more photos online. To construct an attractiveness rating for these available photos, we recruited 100 subjects from the University of Chicago GSB

Decision Research Lab mailing list. The subjects were University of Chicago undergraduate and graduate students in the 18-25 age group, with an equal fraction of male and female recruits.

Each subject was paid \$10 to rate, on a scale of 1 to 10, 400 male faces and 400 female faces displayed on a computer screen. Each picture was used approximately 12 times across subjects. We randomized the ordering of the pictures across subjects to minimize bias due to boredom or fatigue.

Consistent with findings in a large literature in cognitive psychology, attractiveness ratings by independent observers appear to be positively correlated (for surveys of this literature, see Langlois et al. 2000, Etcoff 1999, and Buss 2003). Cronbach's alpha across 12 ratings per photo was calculated to be 0.80 and satisfies the reliability criterion (0.80) utilized in several studies that employed similar rating schemes.¹¹ To eliminate rater-specific mean and variance differences in rating choices, we followed Biddle and Hamermesh (1998) and 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 this standardized rating across the subjects rating a particular photo.

Table 2.3 reports the results of regressions of (reported) annual income on the attractiveness ratings. Our results largely replicate the findings of Hamermesh and Biddle (1994) and Biddle and Hamermesh (1998), although the cross-sectional rather than panel nature of our data makes it difficult to argue for a causal relationship between looks and earnings. Nevertheless, the estimated correlations between attractiveness ratings and reported income are significant. The coefficient estimates on the standardized attractiveness score imply that a one standard deviation increase in a man's attractiveness score is related to a 10% increase in his earnings, whereas for a woman, the attractiveness premium is 12%. Interestingly, there also appears to be a significant height premium for men: a one inch increase is related to a 1.4% increase in earnings. For women, the corresponding height premium is smaller (0.9%) and not statistically significant. We find no important relationship between earnings and weight.

3 A Modeling Framework for Analyzing User Behavior

Our data is 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 to. In order to interpret the data using a revealed preference framework, we make the following assumption:

¹¹Biddle and Hamermesh (1998) report a Cronbach alpha of 0.75.

Assumption 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' .

We will thus 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

$$U_M(m, w) \geq v_M(m) \tag{1}$$

and does not send an e-mail otherwise.

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

Adachi considers a discrete time model, 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 . The sampling is done according to the distribution F_W (the corresponding sampling distribution for women is F_M). We assume that the sampling distribution is stationary, and assigns positive probability of meeting each person on the opposite side of the market. A standard assumption (as in Morgan 1995, Burdett and Coles 1996, and 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:

$$\begin{aligned} \mathcal{A}_W(m, w) &= \mathbb{I}\{\text{woman } w \text{ accepts man } m\} = \mathbb{I}\{U_W(m, w) \geq v_W(w)\}, \\ \mathcal{A}_M(m, w) &= \mathbb{I}\{\text{man } m \text{ accepts woman } w\} = \mathbb{I}\{U_M(m, w) \geq v_M(m)\}. \end{aligned}$$

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

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

The first term in this expression is the utility from a mutual match, and the second and third terms capture the continuation utility from a mismatch.

The Bellman equations characterizing the optimal reservation values and search rules of

man m and woman w are given by:

$$\begin{aligned} v_M(m) &= \rho \int EU_M(m, w) dF_W(w), \\ v_W(w) &= \rho \int EU_W(m, w) dF_M(m). \end{aligned} \tag{2}$$

Adachi (2003) shows that the above system of equations 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.¹² 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.

3.1 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, $m \in \mathcal{M} = \{1, \dots, M\}$, $w \in \mathcal{W} = \{M + 1, \dots, W\}$. The preference orderings are generated by $U_M(m, w)$ and $U_W(w, m)$.¹³

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).

Either of these two extreme points can be reached through the use of Gale-Shapley’s *deferred-acceptance algorithm*. 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,

¹²The solution is not unique, but has a lattice structure in strong analogy to the Gale-Shapley model. See the next section for further details.

¹³We impose the restriction that the preferences are strict.

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. 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 who 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.

3.2 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:

$$\begin{aligned} v_M(m) &= \max_{W \cup \{m\}} \{U_M(m, w) | U_W(w, m) \geq v_W(w)\}, \\ v_W(w) &= \max_{M \cup \{w\}} \{U_W(w, m) | U_M(m, w) \geq v_M(m)\}. \end{aligned} \tag{3}$$

Furthermore, as $\rho \rightarrow 1$, the system of Bellman equations (2) becomes equivalent to the system of equations in (3). 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 (3) imply that in a stable match, man m is matched with the best woman who is willing to match with him, and vice versa.

Generally, Adachi's model has more than one equilibrium. Analogous to the result on men- and women-optimal matches in the Gale-Shapley model, Adachi shows that the set of solutions of the system of equations (2) has a lattice structure and possesses extreme points. At the men-optimal extreme, men are pickier (i.e., they have higher reservation utilities) and women are less picky than in any other solution.

3.3 Discussion

Of course, actual behavior in the online dating market that we study is not exactly described by the models of Adachi or Gale and Shapley. However, both models capture some basic mechanisms that apply to the workings of the dating market that we study. The Adachi model captures the search process for a partner, and the plausible notion that people have an understanding of their own dating market value, which influences their threshold or “minimum standard” for a partner. The Gale-Shapley model, especially their 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, since search frictions on the online dating site are likely to be low, the difference in matching outcomes as predicted by the two modeling frameworks is likely to be small, as suggested by Adachi’s equivalence result.

This motivates the following empirical hypothesis, which we will investigate in Section 6:

Hypothesis Given preference profiles $U_M(m, w)$ and $U_W(w, m)$ estimated using the threshold-crossing rule, matching outcomes obtained on the online dating site are close to those that would have been obtained as a stable match in a Gale-Shapley marriage market with the same preference profiles.

3.4 Costly Communication and Strategic Behavior

If sending e-mails is costly, the threshold rule we use to estimate preferences may lead to biased results. As an example, let us assume that there is a single dimension 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 in equilibrium.

This is a potentially serious source of bias in the preference estimates, and we are compelled to investigate whether strategic behavior is an important concern in our data (Section 4) before we estimate mate preferences. 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 non-trivial, 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. Furthermore, the fear of rejection should be mitigated by the anonymous environment provided by the dating site (in our data, 71% of men’s and 56% of women’s first-contact e-mails in our data are rejected, i.e. do not receive a reply).

Moreover, note that Adachi’s model is one without uncertainty regarding the potential partner’s preferences (i.e. the potential partner’s type is perfectly observed). In reality, these preferences are likely to have an unobservable component, such that initially a mate is uncertain as to how desirable he or she is to the potential partner. Then, if the expected benefit from *any match* within a mate’s acceptance set exceeds the marginal cost of sending an e-mail, the users will not strategically refrain from contacting mates they find acceptable.

We should also note that the presence of strategic behavior does not render the empirical investigation of the hypothesis stated above uninteresting. It merely changes our interpretation of the “preferences” that are estimated using the threshold crossing rule. I.e., even if we interpret the users’ e-mailing behavior as indicative of their expectations about their likely equilibrium match partners, a comparison between actual matches observed on the online dating site, and simulated matches obtained by the Gale-Shapley algorithm (that uses “preference” estimates based on the threshold crossing rule) may be seen as a test of whether the users have rational expectations.

4 Some Preliminary Evidence on Partner Choice

As we discussed in Section 3, if the time cost of composing an e-mail or the psychological cost of rejection is significant compared to the *expected* benefit from an eventual match, a site user may not contact an otherwise desirable mate if that mate appears to be unattainable. For example, unattractive men may shy away from sending e-mails to very attractive women, and instead focus their efforts on women who are similar to their own attractiveness level. Such behavior can introduce bias in our estimates. In this Section, we examine whether there is any preliminary evidence pointing towards strategic behavior in our 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 profiles of potential mates that 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

$$\text{EMAIL}_{ij} = \beta \cdot \text{ATTRACTIVENESS}_j + u_i + \varepsilon_{ij}, \quad (4)$$

where EMAIL_{ij} equals 1 if browser i sends an e-mail to mate j . The term u_i indicates person-specific fixed effects (conditional logit estimates yield similar results). Within the context of a sequential search model, u_i can be interpreted as the (unobserved) optimal search threshold for sending an e-mail to profile j .

We first use our measure of physical attractiveness as a proxy for the overall attractiveness of a profile. We run the regression (4) separately for users in different groups of physical attractiveness. I.e., we segment the suitors according to their physical attractiveness, and allow for the possibility that users in different groups respond differently to the attractiveness of the profiles that they browse. Figure 4.1 shows the relationship between a browsed profile’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.

Figure 4.2 provides some evidence on the probability of receiving a reply to a first-contact e-mail. This figure shows the relationship between the physical attractiveness of the person sending a first-contact e-mail and the probability that the receiver replies. As expected, the relationship is monotonic in the attractiveness of the sender (there is no real concern regarding rejection here, since the responder knows that the person who initiated the contact is interested in him or her). Note that men appear much more receptive to first-contact e-mails than women. The median man (in terms of photo attractiveness) can expect to hear back from the median woman with an approximately 35% chance, whereas the median woman can expect to get a reply with a more than 60% chance. Figure 4.2 also provides evidence that more attractive men and women are “pickier.” The least attractive women are two to three times more likely to reply to a first-contact e-mail than the most attractive women. However, despite this difference in “pickiness,” we see that men in the bottom quintile of the attractiveness distribution can expect to hear back from the top quintile of women with more than 20% probability. This appears to be a good return to spending a few minutes on writing an introductory e-mail, or spending less than one minute using a “copy and paste” strategy.

These results provide some support for our assumption regarding the absence of signifi-

cant costs of e-mailing attractive users and (consequently) strategic behavior. Note that this 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 (the results in Section 5 confirm that physical attractiveness is one of the most important preference components). Still, we take the empirical evidence of this Section as suggestive, and leave a more detailed examination of the importance of strategic behavior for future research.

5 Mate Preference Estimation

We employ two approaches to estimating mate preferences. The first method, which we call the *outcome regression approach*, is mainly based on the assumption that all men and women have homogeneous preferences over their potential mates. The single-dimensional index that describes these preferences, and the relationship of the index to *all* observed user attributes, can then be estimated using regression analysis. This approach can be extended to the case where the source of preference heterogeneity is known a priori, for example in the case of ethnicity-based preferences. The second approach allows for preference heterogeneity in a more flexible way, and is based on a discrete choice estimator. While more general than the first approach, it is also computationally more costly, and therefore requires us to make a priori assumptions on what user attributes to include. The choice of these attributes is guided by the results from the first estimation approach.

5.1 Outcome Regressions: Homogeneous Preferences and A Priori Heterogeneity

Consider the following two assumptions, which we impose on the Adachi model (Section 3):

1. All men (women) agree on women’s (men’s) rankings: if $U_M(m, w) \geq U_M(m, w')$, then $U_M(m', w) \geq U_M(m', w')$ for all m' (an analogous condition is satisfied for women’s preferences). In particular, all men and women can be ranked according to a utility index $U_W(m)$ and $U_M(w)$.
2. All profiles are equally likely to be sampled during the search process.

The first assumption, which says that preferences are homogeneous, is critical to the approach in this Section. Under assumptions 1 and 2, *higher ranked women (men) receive e-mails at a higher rate*. The expected number of e-mails received is therefore monotonically related to a user’s rank. We assume that this rank or utility index is a function of various user attributes and a preference parameter that determines the valuation of mate attributes.

All women, for example, rank men according to the same utility index $U_W(X_m; \theta_W)$. We can then infer the relationship between the utility index and the mate attributes using regression analysis, where the number of unsolicited e-mails received is regressed on the user’s attributes.

The single index assumption can be relaxed if the source of preference heterogeneity is known a priori, such that all users can be segmented into a small number of distinct groups. Preferences within a group are assumed to be homogeneous, in which case all group members rank a potential mate according to the same index. Using the same reasoning as above, it is clear that the group-specific utility index is monotonically related to the number of first-contact e-mails that were received from the members of group g . Group g preferences can then be estimated using the following steps: (1) For any user in the data set, count the number of first-contacts received from the members of group g , and (2) regress this outcome measure on all user attributes. This approach is of course only practical for a small number of user segments, which, for example, rules out heterogeneity that is based on several segmentation variables.¹⁴

We note that if preferences are not homogeneous, our regressions still reveal what makes users click, and how dating outcomes or “success” are related to a user’s traits. Of course, to equate the quantity of e-mails received with success, it must also be true that there is no systematic relationship between the number of first-contacts and the average “type” of the users from who these e-mails originate.

We denote the number of first-contact, i.e. unsolicited e-mails that a user received by Y . Y is an integer outcome, and we therefore use Poisson regression, a count data model, to estimate the model parameters.¹⁵ The conditional expectation of the outcome variable is specified as $\mathbb{E}(Y|x) = \exp(x'\theta)$, where x is a vector of user attributes. Under the Poisson assumption, this conditional expectation fully determines the distribution of the outcome variable. The Poisson assumption places strong restrictions on the data. In particular, the conditional variance of a Poisson distributed outcome variable equals the conditional expectation, $\text{Var}(Y|x) = \mathbb{E}(Y|x)$. However, as long as the conditional expectation is correctly specified, the (quasi) maximum likelihood estimator associated with the Poisson regression model is consistent, even if the Poisson assumption is incorrect (Wooldridge 2001, pp. 648-649). We report robust (under distributional mis-specification) standard error estimates for

¹⁴Consider an example where preferences vary by income, education, looks, and age. Even if each of these variables takes only three values, the total number of segments that describe a homogeneous group is $3^4 = 81$.

¹⁵Alternatively, a linear regression model has the obvious disadvantage of predicting negative outcome values for some user attributes. A logarithmic transformation of the outcome variable avoids this problem, but would force us to drop many observations for which the outcome measure is zero. Furthermore, it is not clear how the estimated conditional expectation $\mathbb{E}(\log(Y)|x)$ is related to the object of our interest, $\mathbb{E}(Y|x)$. The same problem pertains to the transformation $\log(1 + Y)$, which is defined for outcome values of zero.

the regressions (Wooldridge 2001, p. 651).

In our application, all regressors are categorical variables indicating the presence of a specific user attribute. If two users A and B differ only by one attribute that is unique to A, with the associated regression coefficient θ_j , the ratio of expected outcomes is

$$\frac{\mathbb{E}(Y|x_A)}{\mathbb{E}(Y|x_B)} = \exp(\theta_j).$$

The *incidence rate ratio*, $\exp(\theta_j)$, measures the premium (or penalty) from a specific attribute in terms of an outcome multiple. For example, using the number of e-mails received as outcome variable, the coefficient associated with “some college” education is 0.21 for men. Hence, holding all other attributes constant, men with some college education receive, on average, $\exp(0.27) = 1.31$ as many e-mails as the baseline group, men who have not finished high school yet. Alternatively, we can calculate the “college premium” for men as $100 \times (\exp(0.27) - 1) = 31\%$.

Table 5.1 presents summary statistics of the outcome measures. Women are browsed more often, and receive more first-contact e-mails and e-mails containing a phone number or e-mail address than men. A first contact is therefore more likely to be initiated by a man. While men receive an average of 2.3 first-contact e-mails, women receive 11.4 e-mails. 56.4% of all men in the sample did not receive a first-contact e-mail at all, whereas only 21.1% of all women were never approached.

We estimate separate regressions for men and women. All 304 observed user attributes are used in the analysis. As the outcome numbers are only meaningful if measured with respect to a unit period of time, we include the (log) number of days a user was active on the dating site as a covariate. Also, we include a dummy variable for users who joined the dating service before the start of the sampling period. Below, we present the estimation results separately for different categories of user traits.¹⁶

Regression Results

Goodness of fit A preliminary analysis shows what fraction of the variability in the number of first contacts is explained by different user attributes. To that end, we present R^2 measures obtained from several OLS regressions using the transformed outcome measure $\log(1 + Y)$ as the dependent variable.¹⁷ A similar, straightforward goodness of fit measure is not available for the Poisson regressions employed in the remainder of this Section.

The results are displayed in Table 5.2. The full set of user attributes explains 28% of the outcome variability for men and 44% of the outcome variability for women. “Looks”

¹⁶The full regression results are available from the authors.

¹⁷The outcome Y is adjusted for the number of days a user was active during the sample period.

has the strongest explanatory power (30% for women and 18% for men), while income and education, if used as the only regressors, explain only a much smaller fraction of the outcome variance.

Stated “dating goals” The impact of the stated goals for joining the dating service on the number of first-contact e-mails received differs across men and women (Figure 5.1). Men who indicate a preference for a less than serious relationship or casual sex are contacted less often than men who state that they are “Hoping to start a long term relationship.” Women, on the other hand, are not negatively affected by such indications. To the contrary, women who are “Seeking an occasional lover/casual relationship” receive 17% more first-contact e-mails relative to the baseline, while men experience a 41% penalty. Men who are “Just looking/curious” receive 19% fewer first-contact e-mails, and the statement “I’d like to make new friends. Nothing serious” is associated with a 21% outcome penalty. Either indication is mostly unrelated to women’s outcomes.

Looks and physical attributes The users of the dating service describe many of their physical attributes, such as height and weight, in their profile. Also, about one third of all users post one or more photos online. We rated the looks of those members in a laboratory environment, as previously described in Section 2. We then classified the ratings into deciles, where the top decile was split again in two halves. This classification was performed separately for men and women. The looks of those member who did not post a photo online are measured using their self-descriptions, such as “average looks” or “very good looks.”

The relationship between the looks rating of the member who posted a profile and the number of first-contact e-mails received is shown in Figure 5.2. Outcomes are strongly increasing in measured looks. In fact, the looks ratings variable has the strongest impact on outcomes among all variables used in the Poisson regression analysis. Men and women in the lowest decile receive only about half as many e-mails as members whose rating is in the fourth decile, while the users in the top decile are contacted about twice as often. Overall, the relationship between outcomes and looks is similar for men and women. However, there is a surprising “superstar effect” for men. Men in the top five percent of ratings receive almost twice as many first contacts as the next five percent; for women, on the other hand, the analogous difference in outcomes is much smaller.

Having a photo online per se improves the members’ outcomes. Women receive at least twice as many e-mails, and men receive at least about 60% more e-mails than those users who did not post a photo and describe themselves as having “average looks.” Figure 5.3 also shows that outcomes are positively related to the user’s self assessment, although the effect sizes are small compared to the impact of looks on outcomes for those users who include a

photo in their profile.

Further evidence on the importance of physical attributes is provided by the members' description of their physique. Members who are "chiseled" and "toned" receive slightly more first-contact e-mails than "height-weight proportionate" users, while "voluptuous/portly" and "large but shapely" members experience a sizable penalty.

Height matters for both men and women, but mostly in opposite directions. Women like tall men (Figure 5.4). Men in the 6'3 - 6'4 range, for example, receive 65% more first-contact e-mails than men in the 5'7 - 5'8 range. In contrast, the ideal height for women is in the 5'3 - 5'8 range, while taller women experience increasingly worse outcomes. For example, the average 6'3 tall woman receives 42% fewer e-mails than a woman who is 5'5.

We examine the impact of a user's weight on his or her outcomes by means of the body mass index (BMI), which is a height-adjusted measure of weight.¹⁸ Figure 5.5 shows that for both men and women there is an "ideal" BMI at which success peaks, but the level of the ideal BMI differs strongly across genders. The optimal BMI for men is about 27. According to the American Heart Association, a man with such a BMI is slightly overweight. For women, on the other hand, the optimal BMI is about 17, which is considered under-weight and corresponds to the figure of a supermodel. A woman with such a BMI receives 90% more first-contact e-mails than a woman with a BMI of 25.

Finally, regarding hair color (using brown hair as the baseline), we find that men with red hair suffer a moderate outcome penalty. Blonde women have a slight improvement in their outcomes, while women with gray or "salt and pepper" hair suffer a sizable penalty. Men with long curly hair receive 18% fewer first-contact e-mails than men in the baseline category, "medium straight hair." For women, "long straight hair" leads to a slight improvement in outcomes, while short hair styles are associated with a moderate decrease in outcomes.

Income 65% of men and 53% of women report their income. Income strongly affects the success of men, as measured by the number of first-contact e-mails received (Figure 5.6). While there is no apparent effect below an annual income of \$50,000, outcomes improve monotonically for income levels above \$50,000. Relative to incomes below \$50,000, the increase in the expected number of first contacts is at least 34% and as large as 151% for incomes in excess of \$250,000. In contrast to the strong income effect for men, the online success of women is at most marginally related to their income. Women in the \$50,000-\$100,000 income range fare slightly better than women with lower incomes. Higher incomes, however, do not appear to improve outcomes, and—with the exception of incomes between \$150,000 and \$200,000—are not associated with a statistically different effect relative to the \$15,000-\$25,000 income range.

¹⁸The BMI is defined as $BMI = 703 \times w/h^2$, where w is weight in pounds and h is height in inches.

Educational attainment Figure 5.7 reveals only a slight relationship between outcomes and education. For men, higher levels of education are associated with a modest increase in first contacts; for women, the relationship is essentially flat. We find, however, that an interpretation of these results as preferences is misleading, due to the importance of preference heterogeneity with respect to education.

As a first look at education-based preference heterogeneity, we segment men and women into three groups, based on whether they have attained or are working towards a high school degree, college degree, or graduate degree. Figure 5.8 shows the relationship between education and outcomes, as measured with respect to the number of first-contact e-mails received from each group. The graph displays evidence for preference heterogeneity. Women, in particular, have a preference for men with equivalent education levels. For example, men with a master's degree receive 48% fewer first-contact e-mails from high school educated women than high school educated men. From college educated women, on the other hand, they receive 22% more e-mails, and from women with (or working towards) a graduate degree they receive 82% more e-mails. Similar to the behavior of women, high school educated men appear to avoid women with higher education levels. There is little evidence, however, that men with college or graduate degrees prefer women with a similar education level.

Occupation Online success also varies across different occupational groups. Here, all outcomes are measured relative to those of students, who are chosen as the baseline group. Holding everything else constant, the biggest improvement in outcomes is observed for men in legal professions (62% outcome premium), followed by fire fighters (45%), members of the military (38%), and health related professions (35%). The occupation of women, on the other hand, has little influence on their outcomes; in fact, most professions are associated with a slightly lower number of first contacts relative to students.

Same-race preferences The dating service allows the users to declare a preference for their own ethnicity in their profile. We find a striking difference across men and women in this stated preference: 38% of all women, but only 18% of men say that they prefer to meet someone of their own ethnic background. This stated ethnicity preference also varies across users of different ethnic backgrounds (Figure 5.9). For example, among Caucasians, 49% of all women and 22% of men declare a preference for Caucasian mates. On the other hand, only 30% of black women and 8% of black men state a preference for their own ethnicity.¹⁹

The question is whether ethnicity preferences also influence the interaction between users, and whether the stated ethnicity preferences are reflected in these users' online behavior. We create four groups of users, based on whether they declare their ethnicity as Caucasian,

¹⁹This, of course, could reflect self selection to a dating service with a majority of Caucasian users.

black, Hispanic, or Asian. We then construct first-contact e-mail outcome measures for all users, separately with respect to each segment, as we did before in the analysis of preference heterogeneity.

The regression results provide evidence that members of all four ethnic groups “discriminate” against users belonging to other ethnic groups (Figure 5.10). For example, relative to white men, African American and Hispanic men receive only about half as many first-contact e-mails from white women, while Asian men receive fewer than 25% as many first-contact e-mails. Note that these results fully control for all other observable user attributes, such as income and education. Also, note that these results are not due to a market size effect, as the outcomes reflect the relative success of the different ethnic groups with respect to the same population of potential mates. Overall, it appears that women discriminate more strongly against members of the different ethnicities than men. Also, Asian men and women seem to be least discriminating among the ethnicities, although the effect sizes are not precisely measured.

Figure 5.11 shows the estimated ethnicity preferences separately for users who declare that they only want to meet users of their own race and users who do not have a declared preference. Due to sample size issues, we consider only first-contact e-mails from Caucasians. It is evident that both members who declare a preference for their own ethnicity, and those who do not, discriminate against users who belong to different ethnic groups. However, discrimination is more pronounced for members of the former group, i.e. these users act in a manner that is consistent with their stated preferences. There is strong evidence, however, that the members of the latter group also have same-race preferences, which contradicts their statement that ethnicity “doesn’t matter” to them.

5.2 Discrete Choice Estimation: Heterogeneous Preferences

We now take an alternative, discrete choice based approach to estimating mate preferences, which allows us to control for preference heterogeneity in a more flexible way compared to the a priori segmentation approach pursued in Section 5.1. This approach is computationally more costly and hence forces us to limit the number of included attribute variables. We use the results from Section 5.1 to guide us in the choice of attributes and whether to allow for heterogeneity in a specific preference component.

The estimation approach is based on a sequence of binary decisions, as in the Adachi model of Section 3. For each user, we observe the potential mates that he or she browses, and we observe whether a first-contact e-mail was sent. Man m , for example, contacts woman w if and only if $U_M(m, w) > v_M(m)$. We assume that the utility function depends 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 parameterized as

$$U_M(X_m, X_w; \theta_M) = x_w' \beta_M + (|x_w - x_m|_+)^{\alpha} \gamma_M^+ + (|x_w - x_m|_-)^{\alpha} \gamma_M^- + \sum_{k,l=1}^N \mathbb{I}\{d_{mk} = 1 \text{ and } d_{wl} = 1\} \cdot \vartheta_M^{kl} + \epsilon_{mw}. \quad (5)$$

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. $|x_w - x_m|_+$ is the difference between the woman's and man's attributes if this difference is positive, and $|x_w - x_m|_-$ denotes the absolute value of this difference is the difference is negative.²⁰ For example, consider the difference in age between man m and woman w . If the coefficient corresponding to the age difference in γ_M^+ and γ_M^- is 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 (5) relates preferences to categorical attributes of both mates. d_{mk} and d_{wl} are dummy variables indicating that man m and woman w possess a certain trait. 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 employ two methods to estimating the discrete choice model. First, we use a fixed effects logit estimator, where ϵ_{mw} is assumed to have the standard logistic distribution. The reservation values $v_M(m)$ and $v_W(w)$ are estimated as person-specific fixed effects, denoted by c_m and c_w . Choice probabilities then take the standard logit form:

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

Using this approach, the model is not identified if the attribute differences enter the utility function in linear form ($\alpha = 1$).²¹

We instead estimate the model with quadratic differences ($\alpha = 2$). Our second estima-

²⁰Formally, $|a - b|_+ = \max(a - b, 0)$ and $|a - b|_- = \max(b - a, 0)$.

²¹To 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 a . 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 same utility function as the one parameterized by $(\beta_M, \gamma_M^+, \gamma_M^-, c_m)$.

tion approach allows us to check the sensitivity of the results with respect to this functional form assumption. We estimate a random effects probit model, where the reservation values are assumed to be independent of all observed covariates, independent across mates, and distributed $N(0, \sigma_c^2)$.²² This approach allows us to enter the difference terms in linear form, and thus the estimates may be less sensitive to large attribute difference values. By assumption, $\epsilon_{mw} \sim N(0, 1)$, and the choice probabilities take the form

$$\Pr \{m \text{ contacts } w | m \text{ browses } w\} = F(U_M(X_m, X_w; \theta_M) - c_m),$$

where F is the cdf of the standard normal distribution. A drawback of this approach is the treatment of the reservation values, which are assumed to be independent of the covariates. The reservation values are determined in equilibrium as a function of own attributes and the distribution of attributes of the other market participants (Section 3). Generally, therefore, the independence assumption will be violated.²³ How much bias this introduces in the estimates is unknown.

Because our final interest is in preferences over potential marriage partners, our estimation sample only includes observations on users who state that they are looking for a long term relationship and who are single, divorced, or describe themselves as “hopeful.” Also, we eliminated choices among potential mates who indicate a preference for a casual affair or who are either married or in some other relationship.²⁴

Estimation Results

Table 5.3 presents the maximum likelihood estimates of the binary logit and probit models. Recall that the logit model is estimated with squared attribute difference terms while the probit model is estimated with linear attribute differences. We also estimated the random effects probit model with squared difference terms and found that the results were similar to the logit estimates.

Overall, the results confirm the importance of the variables highlighted in Section 5.1, but qualify some of the main findings. The logit and probit estimates are mostly very

²²We estimate separate random effects variance parameters for men and women.

²³Following Chamberlain (1980), we could specify c_m to be conditionally normal with mean $\mu + x'_m \eta$, and thus allow the reservation values to depend on own attributes. However, because $x_w - |x_w - x_m|_+ + |x_w - x_m|_- = x_m$, the effect of own characteristics on the reservation utility is not separately identified from the effect of own characteristics on the valuation of mate attributes. Identification in this model fails for a similar reason as in the case of the fixed effects logit model with linear attribute differences.

²⁴We also estimated the model with the choices of users who are “just looking/curious.” The results were similar. 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.

similar.²⁵ However, the two approaches sometimes differ in the relative weight put on preferences for the level of a mate attribute versus the difference of two mates' attributes. The logit estimates, in particular, tend to put more weight on the attribute levels, while the probit estimates put more weight on preference heterogeneity (the attribute difference terms). This could indicate that the linear difference terms included in the probit model are more reasonable descriptions of preference heterogeneity than the squared terms in the logit model, which are sensitive to large attribute differences.

As expected, we find that the users of the dating service prefer a partner whose age is similar to their own. The probit estimates, in particular, indicate that men try to avoid older women, while women have a distaste for younger men.

Women who are single tend to avoid divorced men, while divorced women have a preference for a partner who is also divorced. Most corresponding parameter estimates for men are small and statistically insignificant; the exception is that according to the logit estimates, single men do not want to meet divorced women. 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, but not men, prefer a partner who indicates that he is seeking a long term relationship.

As we found previously in the outcome regression results, looks and physique are important determinants of preferences for both men and women. The utility weights on the looks rating variable differ only slightly across men and women. Also, as in the case of the outcome regressions, 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. The probit estimates strongly indicate that this is a relative effect, such that men do not want to meet taller women than themselves. According to the logit estimates, on the other hand, men generally prefer shorter to taller women, irrespective of their own height. Women's preferences over height are the exact opposite of men's preferences. According to the probit estimates, women have a strong aversion to men who are shorter than themselves, while the logit estimates imply that regardless of their own height, women prefer to meet tall men. As regards weight, men have a strong distaste for women with a large BMI, while women tend to prefer heavier men. Here, the quantitative significance of the heterogeneity components is overall small compared to the BMI level effect.

The estimates of men's and women's income preferences confirm the results in Section 5.1. Women, in particular, place about twice as much weight on income than men. There is little

²⁵The different distributional assumptions on the i.i.d. error term introduces differences in the scale of the estimated parameters. Therefore, one should only compare the relative size of the estimated coefficients across the two models.

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. According to the probit estimates, in particular, men avoid women who are more highly educated than themselves, while women avoid less educated men. The logit estimates attribute these gender differences more to level effects, whereby women generally prefer more to less educated men and men generally prefer less educated women.²⁶

The estimated same-ethnicity preferences also confirm the findings in the previous Section. Both men and women have a relative distaste for a partner of a different ethnicity.²⁷ Here, as before in Section 5.1, we find that women “discriminate” more against members of a different ethnicity than men.

Finally, we find that both men and women have a preference for a partner of the same religion.

Attribute Trade-Offs

In order to obtain a better understanding of the relative magnitude of the attribute preferences we consider the implied trade-offs between different traits. We focus on the trade-offs between income and several other attributes.

First, we look at the trade-off between looks and income. Consider a woman evaluating the profile of a man whose looks rating is in the n^{th} decile ($n < 10$) of all looks scores among men. We would like to know the amount of additional income this man would need to be as “successful” with the woman as another man whose looks rating is in the top decile. To that end, we calculate the income variation such that the woman’s utility index for either man is equal. Remember that the utility index allows for preference heterogeneity through attribute distance terms, and hence we also need to specify the income of the woman and the “baseline man” in the top looks decile. We assume (here and below) that the woman has an annual income of \$42,500 and that the man has an annual income of \$62,500. These are the median income levels for men and women among the dating site users in our data. Table 5.4 shows the income tradeoffs for all looks deciles. A man in the bottom decile, for example, needs an additional income of \$186,000 (a total annual income of \$248,500) to compensate for his poor looks. The table also shows that women cannot make up for their looks at all. The reason is that our preference estimates indicate that men’s marginal utility from income is approximately flat between income levels of \$100,000 and \$200,000 and declining for income levels higher than \$200,000. Hence, even for a woman in the 9th

²⁶For men, the estimated preference coefficient on the woman’s education level is not statistically significant.

²⁷In the case of the probit model, the estimates of some ethnicity coefficients are positive. Most of these estimates are not statistically significant, however.

decile of looks there is no amount of additional income that could make her as attractive in a man’s eyes as a woman in the top decile. Of course, these results should not be taken fully literally—functional form assumptions, distributional assumptions, and sampling error will generally influence the precise income compensation numbers. Hence, for example, our model will not be able to accurately predict how a man evaluates a woman with an annual income of \$2 million. However, the results strongly indicate two basic messages: preferences for looks are quantitatively important, and there are strong gender differences in the relative preference of looks versus income.

Table 5.5 shows the trade-offs between height and income. A man who is 5 feet 6 inches tall, for example, needs an additional \$175,000 to be as desirable as a man who is approximately 6 feet tall (the median height in our sample) and who makes \$62,500 per year.

Maybe the most striking numbers are with regard to income-ethnicity trade-offs, as shown in Table 5.6. For equal success with a white woman, an African-American man needs to earn \$154,000 more than a white man. Hispanic men need an additional \$77,000, and Asian men need an additional \$247,000 in annual income. In contrast to men, women mostly cannot compensate for their ethnicity with a higher income.

6 Predicting the Structure of Matches

As noted in the introduction, a very large empirical literature in sociology, psychology, and economics documents strong correlation patterns in the demographic, physical, and socioeconomic characteristics of married couples. In Table 6.1, column (I), we report some of these correlation patterns. To construct this table we utilized the 2000 Census IPUMS 5% sample for the two metropolitan areas covered by our online dating data set. We then located married couples in the sample and computed Pearson correlations of their age, education and income. Married couples exhibit a very strong degree of sorting in age ($\rho = 0.94$) and years of education ($\rho = 0.64$). There is less sorting along income ($\rho = 0.13$), although this measure does not take household production or “potential earnings” into account. Regarding empirical correlations in looks, height, and weight, we consulted several widely cited 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 to our study—between 0.34 and 0.54). Note that the studies reporting correlations in physical 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. Absent a better alternative, however, we take these results as our empirical benchmarks.

A natural question to ask in the online dating setting is whether the structure of matches that are facilitated by this technology is significantly different from the structure of matches formed through traditional channels. Traditionally, people find their marriage partners in the social and geographic environment they live in, such as the school they attend, at work, in their neighborhood, or in public places including bars, discos, parties and outings with activity groups.²⁸ Most people are therefore more frequently exposed to potential partners who are more similar to them in terms of their education, income, or ethnicity than a randomly drawn partner from the general population. Therefore, the empirically observed correlations in marriages along certain attributes, such as age, income and education, may be purely due to the social institutions that bring partners together and only partially due to the preferences that men and women have over their mates.²⁹ Compared to traditional marriage markets, online dating is characterized by only small search frictions, and the resulting matches should therefore be largely driven by preferences and the equilibrium mechanism that brings partners together.

Before we provide some evidence on the observed matches from our dating service, a clarification of what we mean by a match is in order. A main limitation of our data is that we can only track the users' online behavior. We therefore do not know whether two partners who met online ever went on a date or eventually got married. However, our data provides some information on the contents of the exchanged e-mails. We 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 define such 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 responds by sending her contact information).

Table 6.1, column (II) shows the correlation of several user attributes in the observed online matches defined in the above manner. Not surprisingly, age is strongly correlated across men and women ($\rho = 0.73$). Looks, as measured by the standardized photo rating, are also strongly correlated ($\rho = 0.33$). There are smaller but positive correlations in height ($\rho = 0.16$), BMI ($\rho = 0.13$), income ($\rho = 0.15$), and years of education ($\rho = 0.13$).

Although the correlations in online matches, especially for education, are smaller in

²⁸Unfortunately, as noted by Kalmijn (1998), systematic evidence on how couples find each other appears to be scarce. The most notable exception is a study by Bozon and Heran (1989), who survey the meeting places of French couples between 1914 to 1984. In their sample, school accounted for 8% of meetings, work 15%, while dances, parties and gatherings, night clubs, activity groups, outings, holiday clubs, and meetings in other "public" places excluding work accounted for 63%. Home visits (some arranged) and neighborhood encounters accounted for 13%, while personal ads and matrimonial bureaus arranged only 1% of the matches.

²⁹Some of these institutions, such as "upscale" or "dive" bars and clubs, or church "socials," may well have arisen endogenously to facilitate sorting along certain traits. Nonetheless, it is instructive to compare matching in environments with different degrees of search frictions.

magnitude than their offline counterparts, our results suggest that search frictions are not the sole reason for assortative matching. One factor that may explain our finding of smaller correlations than in marriage data is that our definition of a match is much more indicative of a first date than a marriage. One may expect daters to experiment with a wide variety of individuals, and this experimentation may lead to attribute correlations among dating couples being lower than attribute correlations among married couples. We discuss this hypothesis further in Section 6.2.

Note also that sorting patterns in online matches may differ from offline marriages due to selection: the users in our sample who chose to join the dating site may be different from the general population in terms of attributes and tastes. Although our results in Section 2 do not suggest strong differences between our sample and the offline population with respect to observed attributes, we acknowledge that selection on unobservables such as tastes and goals might play a significant role. However, such a discrepancy would lead us to expect different match patterns than those observed in traditional marriages. Yet even with this self-selected sample of individuals, many of the previously documented correlation patterns hold up at least in a qualitative manner.

6.1 Can the Gale-Shapley Model Predict the Correlation Structure of Online Matches?

We next examine whether the observed correlation in online matches can be predicted from the preference estimates in Section 5.2 and a specific assumption on the equilibrium mechanism by which matches are formed. For both geographic markets in our data set, we use the preference estimates shown in Table 5.3 (the fixed-effect logit specification) to construct user-specific preference orderings over members of the opposite sex. Based on these preference profiles, we use the Gale-Shapley algorithm to compute the male- or female-optimal stable matchings in both dating markets. We then compute the Pearson correlations between the attributes of the matched couples. Remember that the specification of preferences also includes an i.i.d. type I extreme value term, ϵ_{mw} , for each pair of potential mates.³⁰ We use random draws of these utility terms to construct a profile of preference orderings. We repeat the process of drawing random utility terms, calculating preference profiles, and running the Gale-Shapley algorithm 50 times, and report the average and standard deviation of the attribute correlations across these 50 repetitions.

In principle, the male- and female-optimal stable matchings in a market can be very different, since one is the stable matching that is unanimously most-preferred by men, and the other is the most-favored stable matching by women (Roth and Sotomayor 1990). Ac-

³⁰To be precise: Man m 's taste shock for woman w is different from woman w 's taste shock for man m , $\epsilon_{mw} \neq \epsilon_{wm}$.

cordingly, Table 6.1, columns (III a) and (III b) show the correlations in user attributes of the predicted stable matches under these two extreme outcomes. Our results indicate that, at least in terms of the attribute correlations, the predicted male- and female-optimal matching outcomes are virtually identical.

We next compare the attribute correlations obtained by the Gale-Shapley matching procedure to those observed in “keyword matches.” We see a wide degree of agreement between the Gale-Shapley predictions in columns (III a) and (III b) and the observed correlations in column (II). The age correlation in observed ($\rho = 0.73$) and predicted ($\rho = 0.71$) matches are very similar. Although the Gale-Shapley predicted correlation in looks is somewhat smaller than the looks correlation in keyword matches ($\rho = 0.19$ versus $\rho = 0.33$), the height and weight correlations predicted by the Gale-Shapley algorithm are higher.³¹ Finally, we get similar correlations in income and education across the keyword and Gale-Shapley matches.

Overall, it appears that the Gale-Shapley algorithm, coupled with our preference estimates, can predict the correlation structure of observed keyword matches quite well. Note that our preference estimation procedure did not assume a particular match formation process beyond the threshold rule. Indeed, the estimates in Table 5.3 were obtained without utilizing any information beyond browsing and first-contact e-mails.

The close correspondence between the attribute correlations across keyword matches and simulated Gale-Shapley matches suggests that the decentralized process of match formation on the online dating site provides a close approximation to the outcome that would have been achieved by Gale-Shapley’s protocol. This is perhaps not too surprising, given the theoretical result that the matching outcome in the Adachi model under negligible search costs is stable in the Gale-Shapley sense (Section 3). While the users of the dating site are of course not literally matched as in the Gale-Shapley procedure, the Adachi model provides what we consider a plausible description of actual search and matching behavior. Furthermore, the technology of online dating was invented precisely for the purpose of minimizing search costs, and the institution provided by the dating site may be considered a close approximation to the environment in the theoretical limit investigated by Adachi.

6.2 The Role of “Unobservables” in Matching

Since the Gale-Shapley algorithm can predict the correlation structure of observed online matches well, we now utilize it to simulate matching outcomes under counterfactual scenarios in which we change some aspects of the user preferences.

Our first counterfactual exercise gives some insight regarding the role played by “unobservables” in the match formation process. Although our preference estimation results show that observable factors such as looks, height, ethnicity, and income play a very important

³¹The looks correlation was only computed when both members of the matched couple had posted photos.

role in the dating market, these observable user attributes appear to account only for part of the preferences for a potential partner.

One interpretation of the unobservable term in the utility function is that it represents noise in the users' behavior. I.e., the users who browse through the listings make random mistakes when choosing which mates to contact by e-mail. An alternative interpretation is that the "unobservable" is a utility component observed by the site users, but unobservable to us, the analysts. The results of Fisman et al. (2005) provide some insight as to what such an unobservable utility component may be. In their study, Fisman et al. asked speed-daters to rate each other in terms of "shared interests." They show that the importance of certain observable attributes, such as race/ethnicity, decline if the "shared interest" factor is accounted for in their revealed preference analysis.

In order to examine the role of these unobserved factors, we match the users again under the assumption that their preferences are only over observable attributes (i.e., we took the ϵ out of the latent utility specification (5)). Table 6.1, column (IV) shows the match correlations under this assumption. We observe that now the correlations in age, and in particular in income, education, and in all the physical attributes is much higher than under the original matching, where some attributes are unobserved or match-specific. For example, the correlation in income is now 0.33 (previously 0.15), and the correlation in looks is 0.51 (previously 0.19). Hence, unobservable factors appear to play a very important factor in the formation of matches.

Another very interesting pattern emerges when we compare the correlations in column (IV) and column (I). The attribute correlations of matches that are made when the users care only about observable attributes are very similar to the correlations in offline marriages.

We offer two alternative explanations for this finding, based on the two different interpretations of the unobservable term in the utility function that we discussed before. Under the first interpretation, this term represents noise in the site users' behavior. In that case, the error term in the econometric specification is not part of the utility from a potential marriage partner. Therefore, when we simulate matches using the observable utility components only, we predict the match outcomes in a market where the participants do not make mistakes, which leads to more sorting along observable attributes.

The other explanation is that the unobservable term reflects a utility component that is not captured by the set of covariates that we include in our preference specification. Traditional marriage markets may be characterized by search frictions. Due to search frictions, it may be difficult to find a partner along attributes that are not easily observed, such as "shared interests," and therefore people may direct their partner search more along easily observed attributes, such as looks, income, and education. Moreover, many traditional institutional settings where people commonly meet their mates facilitate matching along such

observable attributes. These institutional features can also preclude a wider search over potential partners who might compensate for their lack of observable qualities through other, more difficult to observe traits, some of which may be easier to convey on a webpage than in an offline environment. For example, in a noisy bar physical attractiveness is probably the main attribute along which search and matching takes place. An online dating profile, on the other hand, can contain much information regarding a person’s personality that would have been unobservable in a bar.

This explanation leads to an interesting hypothesis: if more and more people find their mates online, we may begin to see less sorting along easily observable attributes, and more sorting along attributes such as “shared interests” that are difficult to measure. Of course, it could be the case that the users of the dating site have a stronger preference for such attributes than the general population, which is precisely the reason why they try to find a partner online. A further examination of our prediction is left for future research on the impact of the online dating technology on marriage markets.

6.3 Absolute versus Relative Preferences

As we already indicated in the introduction, sorting can be driven both by preferences over the absolute value of a partner’s trait, or by preferences for a partner with similar traits. For example, consider a market where people are only distinguished by their beauty. Suppose everyone prefers a more beautiful mate to a less beautiful one. In a stable match, the most beautiful woman marries the best looking man, the second most beautiful woman marries the second most beautiful man, and so forth. Thus, there will be perfect correlation along beauty in this matching. Alternatively, assume that everyone tries to find a partner who is exactly as beautiful as themselves. Perfect correlation along beauty will also arise under this alternative preference structure.

In order to assess the importance of absolute and relative preference components, we set the “distance terms” in the utility specification (5) equal to zero, and also exclude unobserved utility components, as before. We then simulate the stable matches under this alternative preference specification. Table 6.1, column (V) shows that the correlations in looks, age, and in particular in income and education become small relative to the specification with heterogeneous preferences, similar to the results in column (IV). The correlations of height and BMI become negative, due to the “opposite” tastes of men and women over the level of these attributes. These results indicate that according to our preference estimates, sorting is largely driven by preference heterogeneity.

6.4 Physical versus Non-Physical Characteristics

We now try to assess how much of the observed sorting patterns is driven by physical characteristics, such as looks, height and weight, as opposed to non-physical characteristics, such as income and education.

Note that although we do find evidence of sorting along income and education, this might be due to these factors being positively correlated with looks (as we documented in Table 2.3). Thus, to isolate the role of income and education preferences, we simulate stable matches by setting the utility components related to looks, height, and weight equal to zero. The results from this matching exercise are presented in Table 6.1, column (VI). The correlations in education and income still persist, which argues against the presence of spurious correlation based on matching solely on looks. Instead, it appears that some of the correlation in looks may be driven by a preference for income and education: even though looks does not enter into the utility specification, the resulting matches still exhibit positive correlation in looks, possibly due to the correlation of looks with income and education (Table 2.3).

6.5 Interracial matching patterns

In a recent survey of intermarriage patterns, Kalmijn (1998) reports that in the U.S. “virtually all ethnic subgroups marry within their group more often than can be expected under random mating.” We replicate this finding of strong endogamy (marrying within a group) for the two metropolitan areas covered by our sample using the 2000 Census IPUMS 5% sample. Table 6.2, Panel (I) illustrates the matching patterns in the Census data. Whereas random matching would predict, for example, 74.6% of black men to be matched with white women (as 74.6% of women in the respective population are white), only 12.2% of black men are actually married to white women, and 75.6% are married to black women, who comprise only 3.3% of the sample. Endogamy is also present for Hispanics, Asians, and for whites.

One of the many explanations for endogamy is the presence of same-race preferences. It is, of course, very difficult to assess whether preferences or search frictions and institutional factors determine the structure of offline matches, or whether both factors are in effect (possibly with different quantitative importance).

Our sample of online daters, however, provides some insights towards answering this question, subject to the caveat that the sample may not be fully representative of the population at large. Our estimates in Section 5 indicate preferences for a partner of the same ethnicity. To explore how much endogamy is implied by these estimated same-race preferences, we again utilize the Gale-Shapley model to predict matches in our sample.

Panels (II) and (III) of Table 6.2 report the results of the simulations. In Panel (II),

we simulate matches using the fixed effect logit preference estimates reported in Table 5.3. In Panel (III), we set the coefficients of the ethnicity terms equal to zero—i.e., we simulate matches in a world where preferences for race do not exist. Thus, the difference in the intermarriage patterns in these two simulated environments indicates the importance of race preferences in generating endogamy patterns.

Our results show that while some amount of endogamy can be generated by same-race preferences, the predicted level is comparatively less than the endogamy observed in offline marriages. For example, in our simulations, 28.7% of black men match with black women, even though black women account for only 3.3% of all women in the online population. However, 28.7% is much smaller than the 75.6% figure we find in offline marriages. The simulation results indicate similarly low levels of endogamy for all ethnic groups.

Interestingly, however, our simulations can replicate the endogamy patterns in Panel I much closer when we set the unobservable utility term in the preference specification equal to zero. Panel (IV) reports the result of this simulation. Note that now 74.2% (99.2%) of black men (women) are predicted to match with black women (men), in close agreement with the numbers (75.6% and 89.6%) in the 2000 Census. Similarly, 87.4% (70.3%) of Hispanic men (women) are predicted match with Hispanic women (men), compared with 82.7% (76.5%) in the Census data, and 88.8% (93.2%) of white men (women) are predicted to match with white women (men), compared with 93.5% (95.5%) in the Census data. The only ethnic group for which we systematically under-predict endogamy is Asians.

As in Section 6.2, we can provide two alternative explanations for these findings, based on two different interpretations of the unobservable term in our preference specification. The first explanation is based on random noise in the site users' behavior. The second explanation is that unobservable attributes (such as “shared interests,” as found by Fisman et al. 2005) may play a more important role in determining matches in the online dating world, since the online technology reduces the cost of searching over such attributes, compared to traditional search and matching channels.

7 Conclusions

This paper investigates mate preferences, match formation, and the resulting attribute correlation and sorting patterns using a novel data set from an online dating site. Our analysis is based on unusually detailed data on the attributes and interactions of men and women, which are available to us due to the well-defined institutional rules of the online dating market. Our analysis of *revealed preferences*, and the relationship of these preferences to user attributes, confirms many findings in psychology, anthropology, and sociology studies, which are based on *stated preference* data. For example, we find a stronger emphasis on

a partner's income among women than among men. Revealed preference data allow us to investigate mate preferences that people might not truthfully reveal, in particular their behavior towards potential mates of different ethnicities. Regarding such preferences, we find that the users of the dating site prefer to match with a partner of their own ethnicity, and that such same-race preferences are more pronounced for women than for men.

Based on our preference estimates, we use the Gale-Shapley algorithm to predict equilibrium (stable) matches and the resulting correlations along attributes such as age, looks, income, education, and race. These predictions are made under the assumption of no search frictions, which we believe characterizes online dating well compared to the traditional "real world" way of finding a partner. In line with the theoretical results by Adachi (2003), the Gale-Shapley algorithm predicts the structure of online matches quite well. We can tentatively conclude, therefore, that the online dating market that we study yields outcomes that are close to the efficient benchmark.

Surprisingly, we find that the Gale-Shapley algorithm also predicts the structure of actual (offline) marriages well, in particular if we exclude the unobservable utility component from the users' preferences. A possible explanation for this finding is that the behavior of the site users exhibits some noise, which is attributed to the error term in our econometric analysis. A second explanation regards the unobservable as a true utility component (which reflects shared interests, for example), that is not captured by the set of covariates we condition on. Search frictions in traditional dating or marriage markets may result in a form of partner search behavior that is predominantly based on observable characteristics, such as looks, income, and education. The online dating market that we study, in contrast, might make it easier to find a partner who possesses traits, such as shared interests, that are difficult to observe in more traditional settings. Our simulations indicate that by matching on attributes that are unobserved to us, the analysts, the degree of sorting along observed attributes declines.

Based on the results obtained in this study, we believe that an analysis of online dating markets can yield important insights on the workings of dating and also marriage markets. Many 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 results in a marriage, which is one outcome that we are interested in. This issue could be addressed through exit/follow up surveys of dating site users. A methodological drawback of our analysis concerns the issue of strategic behavior. A more structural estimation approach, such as in Choo and Siow (2006) and Wong (2003), could address this caveat to our estimation approach.

References

- [1] Adachi, H. (2003): "A search model of two-sided matching under non-transferable utility," *Journal of Economic Theory*, vol. 113, 182-198.
- [2] Becker, G. S. (1973): "A Theory of Marriage: Part I," *Journal of Political Economy*, vol. 81 (4), 813-846.
- [3] Biddle J. and D. Hamermesh (1998): "Beauty, Productivity, and Discrimination: Lawyers' Looks and Luchre," *Journal of Labor Economics*, vol. 16 (1), 172-201.
- [4] Bozon M and F. Heran (1989): "Finding a Spouse: A Survey of How French Couples Meet, " *Population*, 44, 91-121.
- [5] Buss, D. (2003): *The Evolution of Desire: Strategies of Human Mating*. Basic Books.
- [6] Chamberlain, G. (1980): "Analysis of Covariance with Qualitative Data," *Review of Economic Studies*, 47, 225-238.
- [7] Choo, E. and A. Siow (2006): "Who Marries Whom and Why," *Journal of Political Economy*, 114 (1), 175-201.
- [8] Etcoff, N. (1999): *Survival Of The Prettiest*. Doubleday.
- [9] Fisman, R., S. S. Iyengar, E. Kamenica, and I. Simonson (2005): "Racial Preferences in Dating: Evidence from a Speed Dating Experiment," manuscript, Columbia University.
- [10] Fisman, R., S. S. Iyengar, E. Kamenica, and I. Simonson (2006): "Gender Differences in Mate Selection: Evidence From a Speed Dating Experiment," *Quarterly Journal of Economics* (forthcoming).
- [11] Fox, J. T. (2006): "Estimating Matching Games with Transfers," manuscript, University of Chicago.
- [12] Gale, D. and L. S. Shapley (1962): "College Admissions and the Stability of Marriage," *American Mathematical Monthly*, 69, 9-14.
- [13] Hamermesh, D., J. Biddle (1994): "Beauty and the Labor Market," *American Economic Review*, 84 (5), 1174-1194.
- [14] Hinsz, V. B. (1989), "Facial Resemblance in Engaged and Married Couples," *Journal of Social and Personal Relationships*, 6, 223-229.
- [15] Kalmijn, M. (1998): "Intermarriage and Homogamy: Causes, Patterns, Trends" *Annual Review of Sociology*, 24, 395-421.

- [16] Roth, A. and M. Sotomayor (1990): *Two-Sided Matching: A Study in Game-Theoretic Modeling and Analysis*. Cambridge: Cambridge University Press.
- [17] Langlois, J. H., Kalakanis, L., Rubenstein, A. J., Larson, A., Hallam, M., and Smoot, M. (2000): "Maxims or Myths of Beauty? A Meta-Analytic and Theoretical Review," *Psychological Bulletin*, 126, 390-423.
- [18] Spuhler, J.N. (1968), "Assortative Mating with Respect to Physical Characteristics," *Social Biology*, 15 (2), 128-140.
- [19] Stevens, G. and D. Owens and E. C. Schaefer (1990): "Education and Attractiveness in Marriage Choices," *Social Psychology Quarterly*, 53 (1), 62-70.
- [20] Wong, L. Y. (2003): "Structural Estimation of Marriage Models," *Journal of Labor Economics*, 21 (3), 699-727.

Table 2.1 – Dating Service Members and County Profile of General Demographic Characteristics

Variable	San Diego			Boston		
	Dating Service	General Population	Internet User	Dating Service	General Population	Internet User
<i>General Information</i>						
No. of Members/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 (1)</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>						
<i>Men</i>						
Never married	65.6	31.8	28.5	67.2	35.3	36.8
Married & not separated	6.3	57	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
<i>Women</i>						
Never married	62.2	20.2	23.9	65.9	28.0	32.7
Married & not separated	2.6	57	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
<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 (2-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

Income (2)

<i>Individuals with Income information</i>	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

Source. Estimates from CPS Internet and Computer use Supplement, September 2001. All the CPS estimates are weighted. We only consider individuals who are at least 18 years old. The percentages for the column "Internet user" are from the CPS sample, restricted to individuals who declare to use the Internet.

Notes. The geographic information regards two MSA's. The Boston PMSA includes a New Hampshire portion. San Diego geographic information corresponds to the San Diego MSA. The site member information is from 2003.

(1) The figures for whites, blacks, Asians and "other" ethnicities for the CPS data correspond to those with non-Hispanic ethnicity.

(2) The income figures from the CPS data were adjusted to 2003 dollars.

Table 2.2 – Physical Characteristics of Dating Service Members vs. General Population

Variable	Men		Women	
	Dating Service	General Population	Dating Service	General Population
Weight (lbs)				
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
Height (inches)				
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
BMI**				
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

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

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

Table 2.3 – Log Earnings and Photo Ratings

	Men		Women	
	(1)	(2)	(1)	(2)
Years of Education	0.0809 (0.0055)*	0.0808 (0.0055)*	0.0762 (0.0065)*	0.0756 (0.0066)*
Standardized Photo Rating	0.0988 (0.0230)*	0.0974 (0.0237)*	0.1244 (0.0246)*	0.1175 (0.0263)*
Weight (lbs)		-0.0002 (0.0006)		-0.0004 (0.0007)
Height (inches)		0.0140 (0.0054)*		0.0085 (0.0061)
Observations	1,665	1,665	1,136	1,136
R-squared	0.51	0.52	0.45	0.45

Note: The dependent variable in each regression is the log of reported annual income. Each regression also includes indicator variables controlling for occupation, ethnicity, marital status, and the city (Boston or San Diego) where the user lives. We also included a “years in the workforce” variable, defined as the age of the user minus the years of education minus five. The square of this variable is also included. Standard errors are reported in parentheses.

Table 5.1 – Description of Outcome Measures

	Browsets ^a	First Contacts ^b	Keywords ^c
<i>Men</i>			
<i>All Observations</i>			
No. Obs.	12,042	12,042	12,042
Median	10	0	0
Mean	40.8	2.3	1.3
SD	76.7	5.7	5.6
Min	1	0	0
Max	1,059	88	269
% Obs. = 0	0.0	56.4	68.3
<i>Observations > 0</i>			
No. Obs.	12,042	5,255	3,819
Median	10	2	2
Mean	40.8	5.3	4.2
SD	76.7	7.6	9.3
<i>Women</i>			
<i>All Observations</i>			
No. Obs.	9,703	9,703	9,703
Median	34	4	1
Mean	116.2	11.4	3.4
SD	182.0	20.3	7.5
Min	1	0	0
Max	1,649	202	378
% Obs. = 0	0.0	21.1	43.9
<i>Observations > 0</i>			
No. Obs.	9,703	7,657	5,447
Median	34	6	3
Mean	116.2	14.5	6.1
SD	182.0	21.9	9.2

^a Number of times user was browsed by unique users

^b Number of first-contact e-mails received

^c Number of e-mails containing a phone number or e-mail address received.

Table 5.2 – Explanatory Power of Different User Attribute Categories for First-Contact E-mails Received

<u>Attributes</u>	Men	Women
Looks	0.18	0.30
Income	0.07	0.04
Education	0.05	0.02
Looks, Income	0.21	0.320
Looks, Income, Education	0.22	0.325
All	0.28	0.44

Note: The table reports R-squared measures from different OLS regressions using $\log(1+Y)$ as the dependent variable. Y is defined as the number of first-contact e-mails received per day active. The regressions in the last row of the table (“All”) include all user attributes that were used in the Poisson regressions of Section 5.

Table 5.3 – Binary Logit and Probit Estimates

	Men				Women			
	Fixed Effects Logit (Squared Difference Terms)		Random Effects Probit (Linear Difference Terms)		Fixed Effects Logit (Squared Difference Terms)		Random Effects Probit (Linear Difference Terms)	
	Estimate	SE	Estimate	SE	Estimate	SE	Estimate	SE
Age	-0.0465	0.0034	-0.0028	0.0013	0.0034	0.0043	-0.0039	0.0013
Age Difference (+)	-0.0015	0.0003	-0.0420	0.0020	-0.0024	0.0003	-0.0143	0.0017
Age Difference (-)	-0.0046	0.0002	-0.0226	0.0013	-0.0055	0.0005	-0.0457	0.0021
Single, Mate: Divorced ^a	-0.0761	0.0324	-0.0192	0.0169	-0.0554	0.0388	-0.0466	0.0186
Divorced, Mate: Divorced	0.0100	0.0353	0.0182	0.0186	0.1872	0.0355	0.1076	0.0175
“Long Term”, Mate: “Long Term”	0.0137	0.0205	0.0004	0.0111	0.2194	0.0261	0.1165	0.0131
Has Children, Mate: Has Children	0.2122	0.0350	0.1048	0.0182	0.1943	0.0357	0.0992	0.0174
No Children, Mate: No Children	-0.3054	0.0316	-0.1711	0.0167	-0.3798	0.0403	-0.1775	0.0192
Has Photo	-0.0554	0.0479	-0.0196	0.0266	0.1843	0.0552	0.0948	0.0272
Looks Rating	0.5630	0.0204	0.2774	0.0108	0.5432	0.0271	0.2647	0.0136
Own Looks Rating			-0.0729	0.0196			-0.1347	0.0169
“Very Good” Looks	0.5995	0.0557	0.3299	0.0311	0.5276	0.0672	0.2790	0.0342
“Above Average” Looks	0.3769	0.0509	0.1937	0.0283	0.1871	0.0594	0.0968	0.0298
“Other Looks”	0.3718	0.3049	0.0312	0.1760	0.0903	0.2565	0.0941	0.1216
Height	-0.1637	0.0092	-0.0280	0.0040	0.1888	0.0115	-0.0044	0.0039
Height Difference (+)	0.0013	0.0047	-0.1120	0.0125	-0.0100	0.0008	0.0309	0.0033
Height Difference (-)	-0.0108	0.0007	-0.0154	0.0037	-0.0420	0.0149	-0.2155	0.0190
BMI	-0.3826	0.0390	-0.2423	0.0155	0.1285	0.0616	0.1182	0.0260
BMI ²	0.0037	0.0008	0.0031	0.0003	-0.0006	0.0012	-0.0021	0.0005
BMI Difference (+)	0.0051	0.0010	0.0038	0.0083	-0.0103	0.0010	-0.0035	0.0033
BMI Difference (-)	-0.0099	0.0008	-0.0409	0.0037	0.0006	0.0014	-0.0229	0.0055
Education (Years)	-0.0058	0.0079	0.0068	0.0044	0.0550	0.0094	-0.0002	0.0043
Education Difference (+)	-0.0042	0.0014	-0.0284	0.0043	-0.0094	0.0016	-0.0100	0.0040
Education Difference (-)	-0.0031	0.0012	-0.0053	0.0049	-0.0005	0.0016	-0.0389	0.0051
Income (\$ 1,000)	0.0068	0.0017	0.0033	0.0009	0.0171	0.0041	0.0071	0.0019
Income (>50) ^b	-0.0037	0.0027	-0.0015	0.0014	-0.0058	0.0048	-0.0014	0.0022
Income (>100) ^b	-0.0039	0.0029	-0.0027	0.0016	-0.0094	0.0020	-0.0054	0.0010
Income (>200) ^b	-0.0031	0.0047	-0.0013	0.0025	0.0076	0.0022	0.0027	0.0011
Income Difference (+)	0.0000	0.0000	0.0004	0.0003	0.0000	0.0000	-0.0007	0.0003
Income Difference (-)	0.0000	0.0000	-0.0006	0.0002	0.0000	0.0000	0.0000	0.0003
Income “Only Accountant Knows”	0.4022	0.0677	0.2226	0.0356	1.1888	0.1856	0.5177	0.0832

Income “What, Me Work?”	0.2970	0.0802	0.1500	0.0423	0.7916	0.2034	0.3289	0.0927
White, Mate: Black	-0.9584	0.1226	-0.4105	0.0602	-0.7159	0.1565	-0.3387	0.0702
White, Mate: Hispanic	-0.3229	0.0522	-0.1701	0.0283	-0.5029	0.1119	-0.2524	0.0541
White, Mate: Asian	-0.5299	0.0600	-0.2363	0.0320	-1.6735	0.3253	-0.6927	0.1290
White, Mate: Other	-0.1082	0.0317	-0.0443	0.0170	-0.0674	0.0376	-0.0297	0.0190
Black, Mate: White	-0.3836	0.6132	0.4860	0.0530	-1.4085	0.5851	0.1913	0.1390
Black, Mate: Hispanic	-0.4503	0.6920	0.2301	0.1904	-1.0706	1.4034	0.4956	0.5010
Black, Mate: Asian	-2.8566	1.1949	-0.3961	0.3518				
Black, Mate: Other	-0.1442	0.6400	0.6141	0.1197	-1.3328	0.6994	0.1299	0.2285
Hispanic, Mate: White	-0.3182	0.2716	0.0615	0.0559	-0.4681	0.3423	0.0094	0.0478
Hispanic, Mate: Black	-0.7339	0.7136	0.4245	0.2986	-0.3414	0.6344	0.2090	0.2738
Hispanic, Mate: Asian	-0.2713	0.4929	-0.0529	0.1859				
Hispanic, Mate: Other	-0.7986	0.3274	-0.0046	0.1050	-0.4798	0.3909	0.0289	0.1088
Asian, Mate: White	-0.3946	0.4391	-0.3498	0.1069	0.3452	0.7258	0.1978	0.0599
Asian, Mate: Black					-0.0004	1.0966	-0.1778	0.4213
Asian, Mate: Hispanic	0.6889	0.5462	0.0783	0.2244	-0.3197	0.9124	-0.1612	0.2639
Asian, Mate: Other	0.3477	0.5058	-0.0771	0.1545	-0.0868	0.7655	-0.0138	0.1238
Same Religion	0.2314	0.0298	0.0927	0.0158	0.3181	0.0319	0.1678	0.0156
Log-likelihood	-37,590.02		-40,017.34		-31,962.80		-35,203.29	
Observations	123,181		143,533		122,575		143,179	
Individuals	1,681		3,148		1,599		2,730	

^a The user who makes the choice is single, and the potential mate is divorced.

^b Income ($> X$) is the amount of income (in \$1,000) above the income level X .

Note: The dependent variable is the “0/1” choice to contact a previously “browsed” user. The model includes fixed effects or random effects for each user. The estimation is based on a sub-sample of users who state that they are “looking for a long-term relationship.” In the full sample we observe more choices for men than for women. In order to reduce the computation time, we took a random sample of the men’s choices (i.e., we kept all men, but randomly discarded some of their observed choices). Also, note that the sample size is smaller for the fixed effects logit than for the random effects probit model. The reason is that for some users, we observe a sequence of only “0” or “1” contact choices, which makes it impossible to include a fixed effect in the estimation. In the case of the logit model, all such users were eliminated from the sample.

Table 5.4 – Looks/Income Trade-Offs

Looks Rating	Additional	Additional
	Income Needed by Men (\$1,000)	Income Needed by Women (\$1,000)
Average in 1 st Decile	186	Not Feasible
Average in 2 nd Decile	169	Not Feasible
Average in 3 rd Decile	159	Not Feasible
Average in 4 th Decile	151	Not Feasible
Average in 5 th Decile	143	Not Feasible
Average in 6 th Decile	128	Not Feasible
Average in 7 th Decile	86	Not Feasible
Average in 8 th Decile	37	Not Feasible
Average in 9 th Decile	25	Not Feasible
Average in 10 th Decile	0	0

Note: The table shows the additional annual income that a man or woman needs to be as successful as a man or woman whose looks rating equals the average rating in the (upper) 10th decile. The baseline incomes are \$62,500 for men and \$42,500 for women. For example, consider a man whose looks rating is average among the 41 to 50 percent best looking men in the population. In order to be as desirable to a woman as a man whose rating is the average in the top decile and who earns \$62,500 per year, he needs to have an additional income of \$143,000 (i.e., he needs to make \$205,500 per year).

Table 5.5 – Height/Income Trade-Offs

Height	Additional Income Needed by Men (\$1,000)	Additional Income Needed by Women (\$1,000)
5' 0''	317	-43
5' 2''	269	-43
5' 4''	221	-34
5' 6''	175	16
5' 8''	138	Not Feasible
5' 10''	24	Not Feasible
6' 0''	-8	Not Feasible
6' 2''	-30	Not Feasible
6' 4''	-51	Not Feasible
6' 6''	-63	Not Feasible
6' 10''	-63	Not Feasible

Note: The table shows the additional annual income that a man or woman needs to be as successful as a 5' 11.5'' tall man or a 5' 5.5'' tall woman (the median heights in our online dating population). The baseline incomes are \$62,500 for men and \$42,500 for women. For example, consider a man who is 5' 2'' tall. In order to be as desirable to a woman as a man who is 5' 11.5'' tall and who earns \$62,500 per year, he needs to have an additional income of \$269,000 (i.e., he needs to make \$331,500 per year).

Table 5.6 – Ethnicity/Income Trade-Offs

For Equal Success With:	Additional Income Needed by White Men	Additional Income Needed by Black Men	Additional Income Needed by Hispanic Men	Additional Income Needed by Asian Men
	(\$1,000)	(\$1,000)	(\$1,000)	(\$1,000)
White Women	0	154	77	247
Black Women	220	0	184	-
Hispanic Women	59	30	0	-
Asian Women	-24	0	28	0

	Additional Income Needed by White Women	Additional Income Needed by Black Women	Additional Income Needed by Hispanic Women	Additional Income Needed by Asian Women
	(\$1,000)	(\$1,000)	(\$1,000)	(\$1,000)
White Men	0	Not Feasible	Not Feasible	Not Feasible
Black Men	Not Feasible	0	Not Feasible	Not Feasible
Hispanic Men	Not Feasible	Not Feasible	0	Not Feasible
Asian Men	Not Feasible	-	-43	0

Note: The table shows the additional annual income that a man or woman needs to compensate for his own ethnicity if he wants to date a partner who is a member of a different ethnic group. The baseline incomes are \$62,500 for men and \$42,500 for women. For example, consider an Asian man who would like to date a White woman. In order to be as desirable to her as a White man who earns \$62,500 per year, he needs to have an additional income of \$247,000 (i.e., he needs to make \$309,500 per year).

Table 6.1 – Attribute Correlations in Marriages, and Observed and Predicted Online Matches

	Marriages	Observed Online	Simulated Gale-Shapley Matches from Estimated Preferences					
		Matches	Men-Optimal	Women-Optimal	Random Utility Term = 0	Random Utility and Difference Terms = 0	Looks, Height, and Weight Terms = 0	Race Terms = 0
	(I)	(II)	(III a)	(III b)	(IV)	(V)	(VI)	(VII)
Age	0.94*	0.73	0.71 (0.00)	0.71 (0.00)	0.74 [7247]	0.14 [7247]	0.71 (0.00)	0.71 (0.01)
		[3660]	[7247]	[7247]	[7247]	[7247]	[7247]	[7247]
Height	0.31 - 0.63**	0.16	0.28 (0.01)	0.28 (0.01)	0.69 [7247]	-0.16 [7247]	0.01 (0.01)	0.27 (0.01)
		[3660]	[7247]	[7247]	[7247]	[7247]	[7247]	[7247]
Weight (I) BMI (II-VII)	0.08 - 0.32**	0.13	0.19 (0.01)	0.19 (0.01)	0.19 [7247]	-0.22 [7247]	0.01 (0.01)	0.18 (0.01)
		[3660]	[7247]	[7247]	[7247]	[7247]	[7247]	[7247]
Looks Rating	0.34 - 0.54***	0.33	0.19 (0.03)	0.19 (0.01)	0.51 [842]	0.24 [804]	0.14 (0.04)	0.19 (0.03)
		[1991]	[813]	[812]	[842]	[804]	[811]	[805]
Income	0.13*	0.15	0.10 (0.02)	0.10 (0.02)	0.33 [2437]	0.05 [2267]	0.08 (0.03)	0.10 (0.02)
		[839]	[2295]	[2296]	[2437]	[2267]	[2291]	[2295]
Education	0.64*	0.13	0.14 (0.01)	0.14 (0.01)	0.57 [7247]	0.05 [7247]	0.14 (0.01)	0.14 (0.01)
		[3660]	[7247]	[7247]	[7247]	[7247]	[7247]	[7247]

Note: The table displays Pearson correlation coefficients between mate attributes. Entries marked with (*) in column (I) come from data on actual marriages in Boston and San Diego, obtained from the 2000 IPUMS 5% sample. Entries marked with (**) report the range of results obtained by anthropometric studies (No. obs. = 46 - 984) surveyed by Spuhler (1968). The entries for looks correlation in this column, marked with (***), come from Hinsz (1989) and Stevens et al. (1990) who construct the attractiveness of engaged and married couples whose engagement, marriage, and 25th anniversary announcements were published in local newspapers.

(continued on next page)

In column (II), we classify two users as matched if they exchanged e-mails containing contact information (a phone number or e-mail address) or if the e-mails contain certain phrases such as “let’s meet.”

Columns (III) - (VII) report attribute correlations in simulated matches, based on the fixed effect logit preference estimates. We first draw the random utility terms, and construct the preference orderings for each user. We then run the Gale-Shapley algorithm. We repeat this process 50 times to account for randomness in the preference orderings, and report the average and standard deviation of attribute correlations across these 50 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 including these users are not included in the reported attribute correlation numbers. Note that the predicted correlations are not run for the full sample of all users, but only for a select sample of users who are single or divorced and do not indicate that they joined the site for a “casual” relationship.

Column (III a) reports attribute correlations that are obtained in the Gale-Shapley “men-optimal” stable match, and column (III b) reports the correlations for the Gale-Shapley “women-optimal” stable match. Columns (IV) – (VII) report men-optimal stable matches. In column (IV), we set the unobserved utility terms to zero. In (V), we additionally set the difference or distance terms in the utility specification, which account for preference heterogeneity, to zero. In column (VI) we set utility components that pertain to the looks rating, height, and weight equal to zero. In column (VII), we set utility components pertaining to race to zero.

Table 6.2 – Sorting Patterns Along Race/Ethnicity

(I) Actual marriage patterns (2000 Census IPUMS 5% sample)

Married with:	% of Men					% of Women				
	White	Black	Hispanic	Asian	Population	White	Black	Hispanic	Asian	Population
White	93.5	12.2	14.4	9.6	74.6	95.5	5.8	19.6	20.4	75.8
Black	0.2	75.6	0.6	0.2	3.3	0.6	89.6	1.7	1.8	3.9
Hispanic	3.2	5.3	82.7	2.0	12.2	2.2	2.2	76.5	1.9	11.3
Asian	2.2	3.9	1.4	86.8	8.3	0.9	0.5	1.2	74.8	7.1
Other	0.9	2.9	0.9	1.3	1.6	0.8	1.8	1.1	1.1	1.9

(II) Simulated matches predicted by Gale-Shapley algorithm and fixed effect logit estimates

Matching with:	% of Men					% of Women				
	White	Black	Hispanic	Asian	Population	White	Black	Hispanic	Asian	Population
White	82.6	49.5	57.0	43.0	78.7	87.4	43.2	66.3	79.2	80.6
Black	1.7	28.7	5.4	7.4	3.3	2.7	37.3	6.9	1.9	4.5
Hispanic	5.5	11.2	23.8	8.8	7.0	3.8	8.6	17.7	6.0	6.2
Asian	5.6	2.6	6.8	21.4	5.9	1.0	4.1	2.3	6.5	2.9
Other	4.5	8.0	6.9	19.4	5.2	5.2	6.9	6.9	6.4	5.8

(III) Simulated matches when race terms set equal to zero

Matching with:	% of Men					% of Women				
	White	Black	Hispanic	Asian	Population	White	Black	Hispanic	Asian	Population
White	79.8	76.6	70.3	72.2	78.7	83.6	79.7	77.4	77.7	80.6
Black	3.2	4.2	3.6	3.8	3.3	4.3	5.5	4.7	4.5	4.5
Hispanic	6.5	7.5	10.9	8.5	7.0	4.8	5.8	8.4	7.2	6.2
Asian	5.6	6.0	8.0	8.8	5.9	2.3	2.8	3.0	3.6	2.9
Other	4.9	5.5	7.2	6.7	5.2	5.1	6.0	6.5	7.0	5.8

(IV) Simulated matches when when unobserved utility components set equal zero

Matching with:	% of Men					% of Women				
	White	Black	Hispanic	Asian	Population	White	Black	Hispanic	Asian	Population
White	88.8	5.3	12.3	25.3	78.7	93.2	0.0	13.1	91.4	80.6
Black	0.0	74.2	0.0	0.7	3.3	0.3	99.2	6.5	0.0	4.5
Hispanic	1.1	10.4	87.4	8.2	7.0	0.9	0.0	70.3	0.0	6.2
Asian	6.5	0.0	0.0	16.4	5.9	0.6	0.4	2.4	5.6	2.9
Other	3.6	10.1	0.2	49.3	5.2	5.0	0.4	7.7	3.09	5.8

Note: Panel (I) reports intermarriage ethnicity patterns across married couples in the 2000 Census IPUMS 5% sample. The columns in the left half of the table display the percentage of men in each ethnic group that are matched with women of the ethnicities given in the rows. Those who declared themselves to be of “Hispanic” ethnicity are considered Hispanic regardless of race classification. Panel (II) reports the intermarriage patterns that are generated using the (male-optimal) Gale-Shapley procedure, based on the fixed effect logit preference estimates. The simulation results are averaged across 50 draws of the unobservable utility term. Panel (III) sets the race preferences to zero and reports the average of 50 Gale-Shapley simulations, as in Panel (II). Panel (IV) reports the result of Gale-Shapley simulations in which we ignore the unobservable utility term and only take the deterministic term of preferences into account.

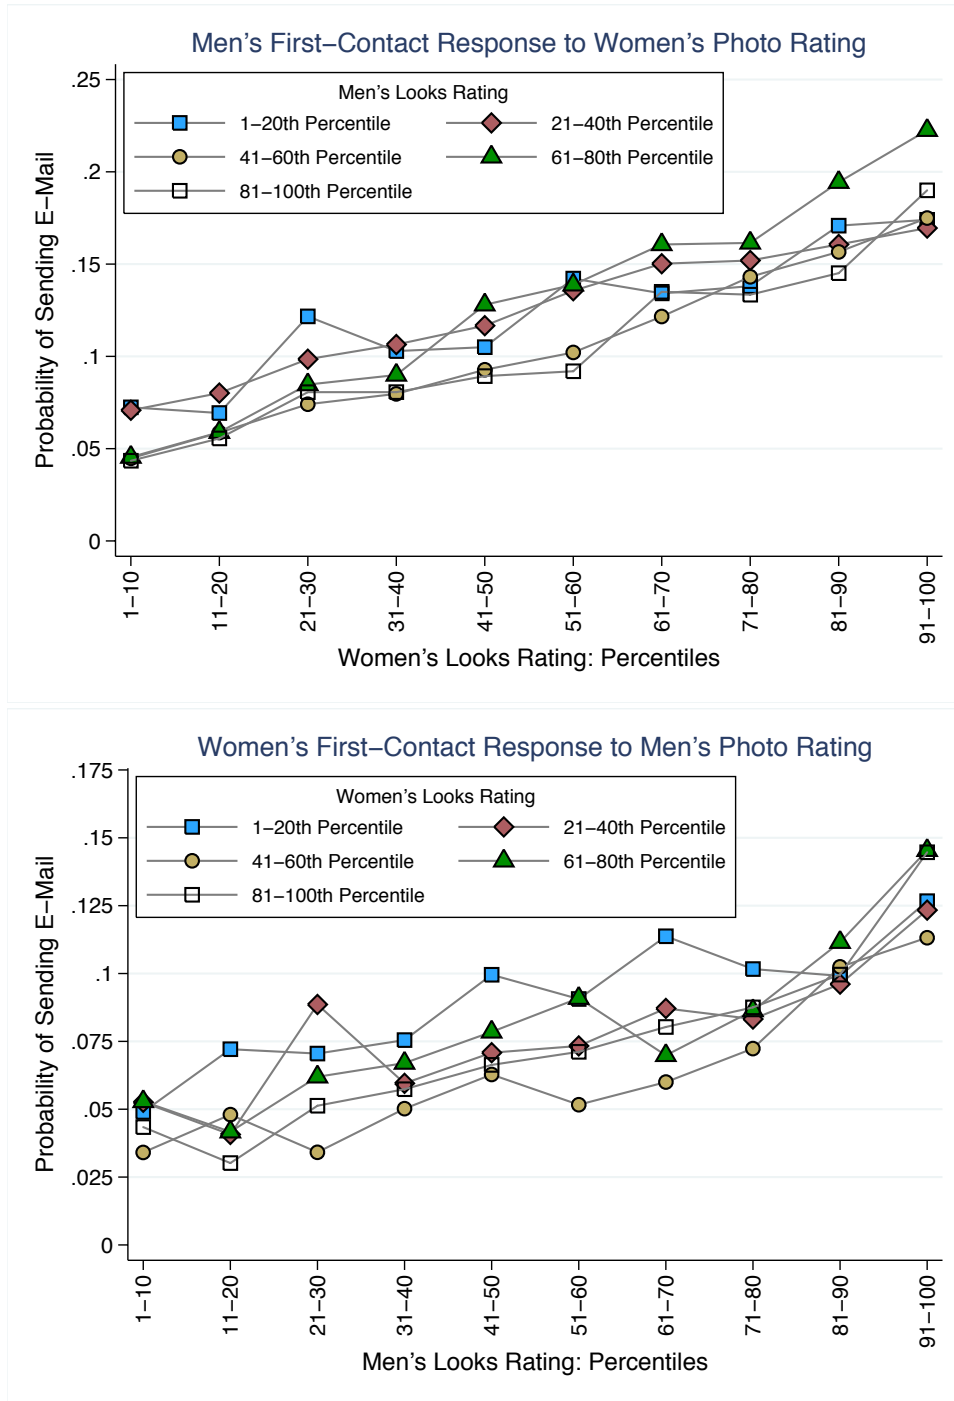


Figure 4.1 – Note: The figures report the result of an OLS regression where the dependent variable is an indicator variable for whether a user sends a first-contact e-mail after browsing the profile of a potential mate. The independent variables are indicators for the photo rating of the user being browsed. The regressions also control for browser fixed effects. The vertical axis plots the estimated mean probability of sending a first-contact e-mail to a browsed profile. 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 the 1-20th percentile of the photo ratings distribution within their gender, etc. The estimates shown are from a sample of users in the 30-39 age range.




Figure 4.2 – Note: The figures report the result of an OLS regression where the dependent variable is an indicator variable for whether a user replied to a first-contact e-mail. The independent variables are indicators for the photo rating of the person sending the first-contact e-mail. The regressions also control for responder fixed effects. The vertical axis plots the estimated mean probability of sending a reply to a first-contact. The horizontal axis is the photo rating of the person sending the first-contact. The regressions were estimated separately for different groups of responders. The first group comprises users who fall within the 1-20th percentile of the photo ratings distribution within their gender, etc. The estimates shown are from a sample of users in the 30-39 age range.

First Contacts – Reason for Joining Site




Figure 5.1

First Contacts – Looks Ratings




Figure 5.2

First Contacts – Self Described Looks




Figure 5.3

First Contacts – Height




Figure 5.4

First Contacts – Body Mass Index




Figure 5.5

First Contacts – Income




Figure 5.6

First Contacts – Educational Achievement

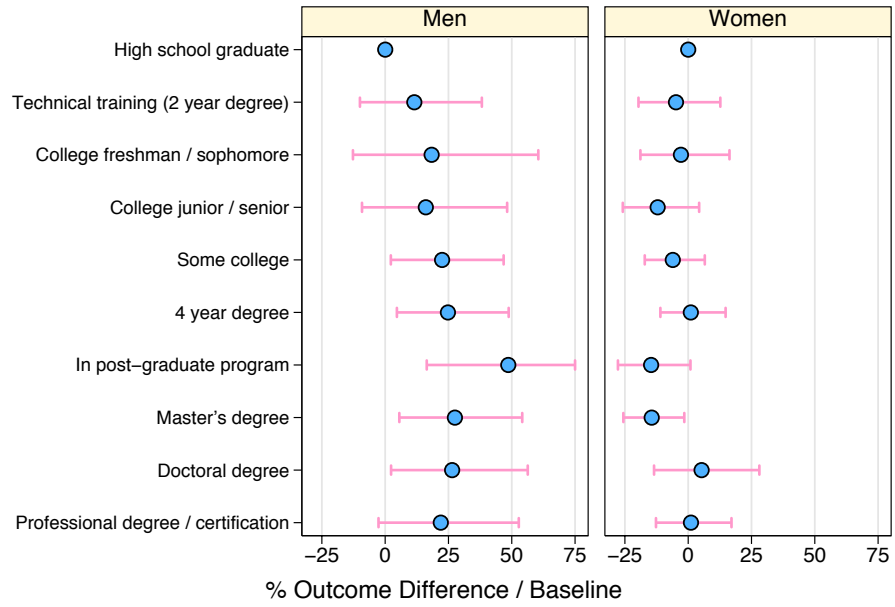


Figure 5.7

First Contacts – Educational Achievement

Segmentation: Education




Figure 5.8

Stated Preference for Ethnicity

Which Ethnicity Do Users Prefer?




Figure 5.9

First Contacts – Ethnicity

Outcomes w.r.t. Ethnic Groups




Figure 5.10 — Note: Each graph displays how own ethnicity affects own outcomes (number of first-contact e-mails received) with respect to potential partners belonging to a certain ethnic group. The bottom right graph, for example, shows how the number of unsolicited e-mails received from Asian men varies by the ethnicity of the receiving woman. In each graph, the ethnicity of the sender is chosen as the baseline ethnicity with respect to which outcomes are measured.

First Contacts – Ethnicity

Outcomes w.r.t. Whites with or without a specific ethnicity preference




Figure 5.11