Two-Sided Estimation of Mate Preferences for Similarities in Age, Education, and Religion

John Allen Logan, Peter D. Hoff, and Michael A. Newton [1]

February 14, 2007

[1] John Allen Logan is Associate Professor, Department of Sociology, University of Wisconsin, Madison, WI 53706 (E-mail: logan@ssc.wisc.edu). Peter D. Hoff is Associate Professor, Department of Statistics, University of Washington, Seattle, WA 98195 (E-mail: hoff@stat.washington.edu). Michael A. Newton is Professor, Department of Statistics and Department of Biostatistics and Medical Informatics, University of Wisconsin, Madison, WI 53706 (E-mail: newton@biostat.wisc.edu). We are grateful to colleagues at the universities of Wisconsin and Washington, and members of ISA Research Committee 28, who provided helpful comments on earlier versions of this paper. We especially thank Eric Grodsky, Charles Halaby, Hsiang-Hui Daphne Kuo, and James E. Yocom for their specific suggestions. This research benefited from facilities provided by a grant from the National Institute of Child Health and Human Development (R24 HD047873) to the Center for Demography and Ecology at the University of Wisconsin-Madison.

Abstract

We propose a two-sided method to simultaneously estimate men's and women's preferences for relative age, education and religious characteristics of potential mates using cross-sectional data on married couples and single individuals, in conjunction with a behavioral model developed in game theory (Roth and Sotomayor 1990) and discrete choice estimation methods developed for simpler, one-sided choice situations (Train 2004). We use fixed effects to control for characteristics that are observed by the opposite sex but are missing from our data. Estimated mean preference coefficients determine the average degree to which measured characteristics of individuals affect others' evaluations of them as marital partners, while the model also accounts for variation of preferences around the means and for limitations in men's and women's information about members of the opposite sex. By assuming each individual chooses freely from the set of potential partners he or she finds available, we estimate preferences without having to observe these sets or to specify any details of the matching process. This makes our method robust to unknown features of the process. Application of the method to data from the first wave of the National Survey of Families and Households (Sweet, Bumpass, and Call 1988) indicates roughly symmetric or complementary preferences of men and women for age, education and religious affiliation characteristics of potential mates, and a much stronger preference for religious homogamy among conservative Protestants relative to mainline Protestants than was suggested by an earlier, retrospective study of religious differences in the temporal stability of marriages (Lehrer and Chiswick, 1993). Our method should be useful in many situations in which voluntary pairings have arisen through some complex process whose details have not been recorded. Besides marriage and cohabitation, data on employment, college attendence, and the coresidence of elderly parents with adult children often have this character, as do some biological data on non-human mating.

KEY WORDS: Two-sided matching; Mate preferences; Marriage model; Religious homogamy; Two-sided probit; Two-sided logit; Data augmentation; Markov chain Monte Carlo; Fixed effects; National Survey of Families and Households.

Many social surveys contain data on the end results of mutual choices made by individuals, but lack explicit information on the context or the constraints under which the decisions were made. For example, the National Survey of Families and Households (NSFH; Sweet, Bumpass, and Call 1988) contains characteristics of married couples and single individuals, but no information on the opportunities for marriage from which individuals chose to attain their marital states. Because of this, it is difficult to make direct inferences on the preferences that individuals have for the characteristics of potential mates. Knowing preferences for characteristics of spousal partners is important and of interest to social scientists, and goes beyond just measuring correlations of traits between spouses, or rates of marriage of men and women.

With knowledge of the preferences of women, for example, social scientists can understand and predict the implications of shortfalls in the supply of certain types of men, which might be due to changes in rates of imprisonment or to casualties of war. Or the effect of a reduction in school dropout rates among men could be predicted to reduce the rate of unwed motherhood by a specific range of values, if the distribution of women's preferences for men's educations (and other traits) were known. Policy makers could employ an increased range of potential interventions in favor of marriage by better knowledge of the motivations of both men and women to marry.

We propose a method for estimating preferences that men and women have for mates, using observations like those in the NSFH and many other surveys. Our specific goals are to infer mate preferences from observed matches, and to formulate the problem in a way that allows us to use the estimated preferences to make counterfactual predictions of marriage market outcomes in different circumstances of supply and demand. The particular preferences we will estimate are for the relative ages, educations, and religious affiliations of potential mates among non-Hispanic whites in the United States.

Other researchers have analyzed the NSFH and similar data sets with the goal of determining the characteristics of men and women that lead to or sustain marriages. The NSFH did ask unmarried respondents to directly state their preferences for certain partner characteristics; the replies were regressed on other characteristics of the respondents by South (1991) and Sprecher, Sullivan, and Hatfield (1994). Raley and Bratter (2004) found that unmarried individuals' stated preferences for hypothetical mate characteristics helped predict the likelihood they would be married roughly five years later. Johnson (1980) and Mare

(1991), among others, used log-linear models on similar data sets to evaluate the importance of personal characteristics on mate selection. Schoen and Wooldredge (1989) analyzed observed characteristics of men and women in marriages, using a "magnitude of marriage attraction" index for male-female pairs. Lehrer and Chiswick (1993) found that religious affiliations of husbands and wives were predictive of the temporal stability of the respondents' first marriages recorded retrospectively in the NSFH, controlling for husbands' and wive's ages, educations, and other characteristics. Bumpass and Sweet (1972) reported greater temporal stability of religious- and age-homogeneous marriages based on a regression analysis of a national sample of female respondents. None of these studies, however, attempted to infer preferences for spousal characteristics by observing the characteristics of matched couples. Kalmijn (1998) and Lehrer (2004) provide reviews of related research.

Potential applications of our estimation method include any situations in which voluntary pairings of actors are observed but information on the process leading to the pairings is not. Employment relationships, the mutual choices of colleges and would-be students, and more general household formation questions, such as the decisions of children and their aged parents to co-reside, are possible applications. The method may also have applications in non-human mating populations when only data on observed matches are available (see Bergstrom and Real (2000) on game-theoretic modeling of non-human mate choice).

1 Revealed Preferences in Typical Marriage Data

In contrast to previous analyses of marriage, we base our method on the idea of revealed preferences, which holds that a person's observed choice of one alternative over another indicates or reveals his or her preference for the characteristics of the alternative chosen. This simple idea, which underlies widely applied discrete choice models for situations where each person's available alternatives are observed by the researcher (e.g., Train 2003), is not immediately applicable to typical marriage data because the needed information about what alternatives were available to each person is not recorded. Typical marriage data are national samples of married and unmarried persons that record characteristics of men and women, including their spouses. That is, there are no lists of the individuals that would have agreed to marry each individual in such data sets, together with the values of their characteristics that would have made them more or less attractive as spouses. Without this information, the observed choices of mates do not directly reveal preferences for their characteristics, but instead reflect some unknown combination of preferences and opportunity constraints.

Complicating the fact that lists of available partners for each survey respondent cannot practically be recorded, it is usually very difficult even to define the particular market in which each person has sought a spouse. High rates of geographic mobility, travel, and other communication mean that spouses are often obtained from non-local sources. In addition, rates of mobility, travel and communication are not completely incidental with respect to marriage, but may reflect to some degree the conscious decisions of people to improve their pools of potential partners.

In spite of these difficulties, there would seem to be little doubt that the broad pattern of marital matches seen in typical marriage data reflects some common preferences and constraints, on average. Such average or typical preferences and constraints may be of interest for questions that transcend particular locales. In order to infer representative preferences of men and women, we suggest it may be unnecessary to observe the exact alternatives from which each man and woman in the sample made his or her choice. Instead, we approximate the alternatives available in each person's set by regarding the women in the broader sample as representative of the women whom the men might have known locally, and the men in the sample as representative of the women's acquaintances. We discuss some limitations of this approach and possible remedies for them subsequently.

2 Interdependent Outcomes

Some clarification of the distinctive approach represented in our statistical procedures is needed. Figure 1 presents scatterplots of men's and women's ages and educations from the NSFH data that we analyze below. Each plotted point in the interior of the top panel shows the husband's and wife's ages in a particular married couple (points have been jittered to make multiple observations with identical values discernable). The bottom panel shows husbands' and wives' years of education. Along the axes of the two scatterplots we have also plotted the ages and educations of the unmarried men and women in the data set. Previous analyses of marriage have focused on the joint distributions of characteristics of husbands and wives, data corresponding to the interior points in the scatterplots of Figure 1, omitting the unmarried persons plotted along the

Figure 1: Observed Age and Education Distributions of Men and Women, NSFH Data.

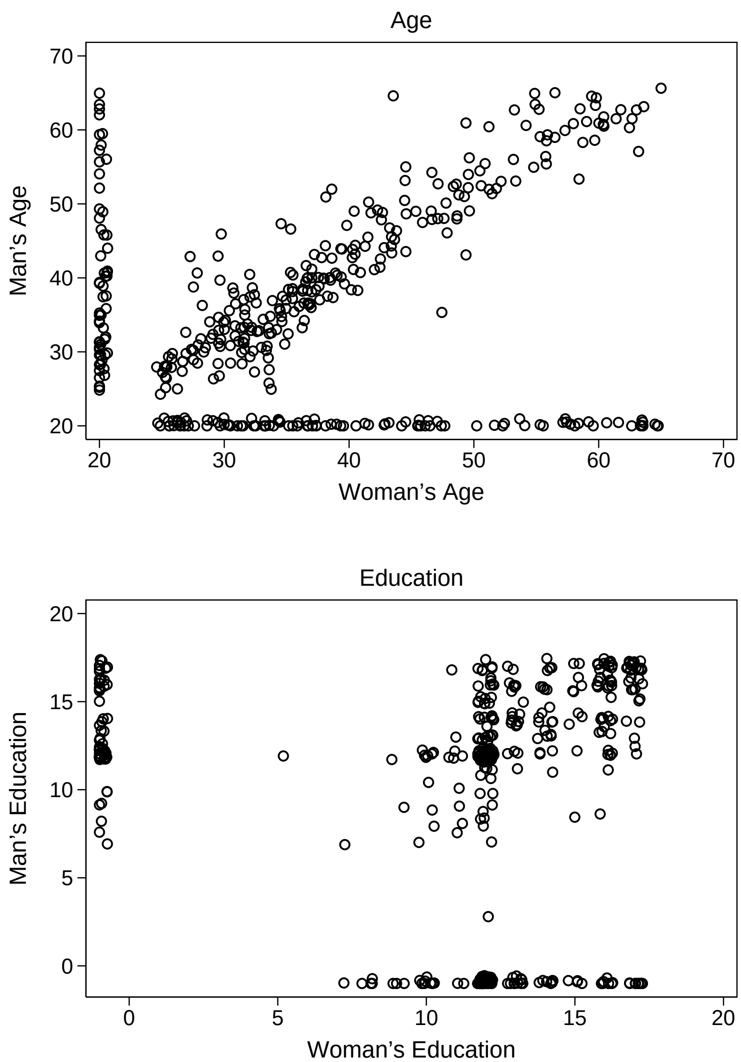

axes. Standard regression models have been used to estimate the distributions of husbands' characteristics conditional on wives', and vice versa (e.g., Stevens, Owens, and Schaefer, 1990). Such regressions are of limited use for the goals we have established. The regression coefficients cannot differentiate the effects of wives' and husbands' preferences for spousal characteristics from the opportunity-constraining effects of the distributions of desired characteristics in the population. Nor can they be used to make plausible predictions of the effects of changes in any of the components. Our method will allow us to do both.

Scatterplots like those in Figure 1 suggest an interpretation different from those underlying regression analyses. In our interpretation the data do not represent draws from a joint probability distribution of men's and women's characteristics with fixed mean and covariance, nor conditional draws from one sex's characteristics given the other's. Instead, the data in the scatterplots — here including the characteristics of the unmarried individuals along the axes — satisfy fixed marginals on both axes. These fixed marginals are the pre-existing age or education distributions of the men and women making up the population. The reasonable requirement that the marginal distributions of characteristics should not be changed by the matching process means the plotted points are not independent observations. Figure 2 illustrates this implicit interdependence using a small, hypothetical scatterplot of men's and women's education levels. The plotted circles labeled 1|1, 2|2, and 3|3 represent marriages of men 1 through 3 with women 1 through 3, respectively. Any different arrangment of points in the scatterplot would entail coordinated shifts among men and women. For example, if woman 1 were to be matched with man 2 instead of man 1, a new point would form in the location that preserves their two levels of education, as indicated by the shaded circle marked 2|1. The circles labeled 1|1 and 2|2 would disappear. This matching of woman 1 with man 2 would leave man 1 and woman 2 unattached. The figure shows one possible resolution of this situation in which woman 2 joins with man 3 to form circle 3|2, and man 1 and woman 3 move to the distributions of unattached individuals along the axes, indicated by circles 1|- and -|3. None of the three original labeled circles remains, while the four new circles preserve the appropriate education levels of their men and women. The set of possible scatterplots is evidently finite and systematically constrained, once the marginal distributions of characteristics are given. Our statistical analysis will take the marginal distributions of quantitative characteristics of men and women as fixed, and will consider the observed joint distribution of characteristics, given the marginals, to reflect

Figure 2: Interdependencies of Scatterplot Positions

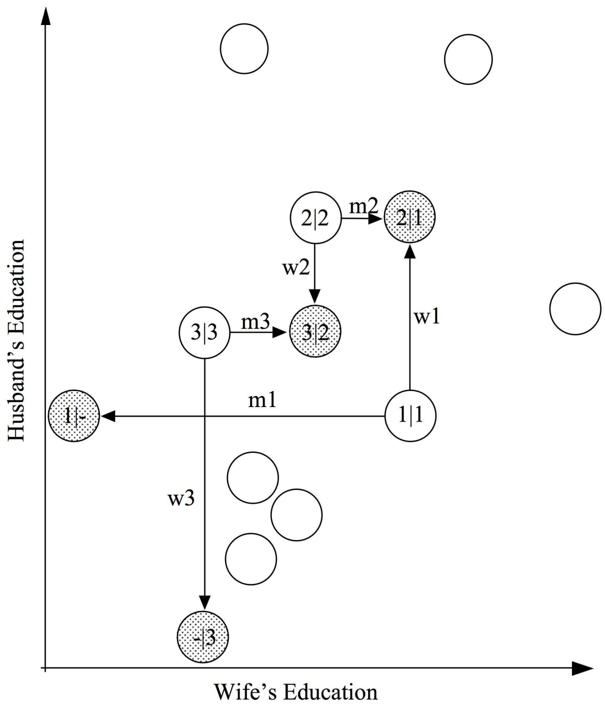

the free mutual choices of men and women according to their preferences for those characteristics. This is a new way of looking at scatterplots of matching data. Rather than considering the data points to be independent draws from suitable distributions, as is done in standard regressions, we view them as highly interdependent observations and base our estimations on a model of the interdependence.

Another key to our analysis of typical marriage data such as the NSFH is the recognition that the data contain no direct information on the processes that led to the observed marriages, but only on the resulting state, the observed marriages themselves. We will therefore model the state rather than the process. On the assumption that marital partners remain together voluntarily we will estimate the preferences that their existing unions reveal. These preferences could also have influenced the unobserved and unmodeled processes by which the marriages formed, as we will discuss.

The strategy of modeling the state of observed matches rather than the process of match formation was also used in an employment context by Logan (1996, 1998). Logan called his statistical model two-sided logit, since the choices made on either side of the market implied a standard logit model once the preferences of the other side were taken into account. However, his method required that alternatives on one side of the market (jobs) be represented as categories (occupations) for computational reasons, while we will treat both sides individually, avoiding a substantial loss of information.

3 The Marriage Model

We now present our behavioral model of marriage, in which men and women voluntarily choose their preferred partners from among the set of those who would be willing to marry them. This model is a parametric version of an abstract model developed in game theory, described by Roth and Sotomayor (1990). We will take advantage of game theoretic results about the model in developing our estimation strategy.

A marriage market involves a population of men and women in which each individual has a preference ordering over the members of the opposite sex. In this paper, we assume the preference orderings are derived from subjective utilities that individuals have for forming pairs with one another and for being single: Man i's utility for pairing with woman j is denoted as $U_{i,j}$, and man i's utility for being single is $U_{i,0}$. Similarly, $V_{j,i}$ and $V_{j,0}$ denote woman j's utilities for man i and for being single, respectively.

In the context of the marriage market, a matching m is a set of pairs of men and women, such that each individual is a member of at most one pair. A matching m determines a husband function h, where $h(j)$ gives the index of the man married to woman j, and $h(j) = 0$ if woman j is single. Similarly, the matching can be written in terms of a wife function w, mapping the indices of the men to those of the women.

We assume each actor in the system chooses a partner or remains single so as to maximize his or her utility over the set of available opportunities: Each man forms a pair with the highest-utility woman among those available to him, or remains single if his utility for singleness exceeds that for his highest-rated available woman. Women do the same, with the roles reversed. Of course, not every actor has the same opportunities: The set of available partners varies from actor to actor, and depends in a complicated way on the utilities of the other actors in the system. For a given set of utilities and a given matching, the available pairing states for man i and woman j are given by their opportunity (or choice) sets

$$O(i) = \{j : V_{j,i} \geq V_{j,h(j)}\} \cup \{0\}$$

$$O(j) = \{i : U_{i,j} \geq U_{i,w(i)}\} \cup \{0\},$$

where the state of being single (denoted "0") is available to each actor. Thus each man has the opportunity to be single or to pair with any woman who prefers him to her match. Opportunities of the women are similar, with the roles reversed. We assume there are no ties among the utilities, so that preferences are strict.

For the purposes of estimation we assume that men and women have selected their matches (whether married or self-matched) voluntarily so as to maximize their utilities given their available alternatives, and that they have utilities for each member of the opposite sex. This means that no man and woman believe they can improve their matches by dissolving their current partnerships and forming a new one with each other, a condition that is called stability in the game theory of marriage.

Stability in this technical sense has no direct temporal implication and does not mean marriages are destined to persist. The willingness of particular men and women to continue in their matches might change in the future for exogenous reasons such as entries or exits of potential partners in the market or changes in the characteristics of its present members. In addition, the preferences of men and women for characteristics of partners can be expected to change over time, possibly in part because their own characteristics are changing. However, considering the population as a whole we expect that overall rates of change in the characteristics of matched partners will be slow, and for this reason we compare our estimates of preferences to estimates of characteristics associated with temporally stable marriages later in this paper.

The assumed voluntary nature of pairings does not mean that marriages are assumed to be uniformly happy. Partners' decisions to stay married take into account the potential costs as well as gains associated with leaving, and large potential costs can outweigh the gains to be expected from decamping to an otherwise more desirable alternative mate. Finally, the assumption that individuals have utilities over all opposite sex members does not mean that they have detailed or accurate knowledge of potential mates' characteristics. Rather, we will model the utilities as having a mean value determined by mean preference coefficients, with a distribution around the mean that reflects differences of opinion as well as imperfect information about both measured and unmeasured characteristics.

Using the above notation, stability implies

$$U_{i,w(i)} \geq U_{i,j} \ \forall j \in O(i), \forall i$$
$$V_{j,h(j)} \geq V_{j,i} \ \forall i \in O(j), \forall j. \tag{1}$$

Stable matchings, in the technical sense, are of central importance in the application of two-sided matching models. Since a matching that is not stable could be improved on by the voluntary actions of system members, it is reasonable to expect that real systems will approximate stable matchings. In addition it has been shown that there is at least one stable matching (and typically many) for any given configuration of utilities in a market, so that a stable matching is always achievable (Roth and Sotomayor 1990).

The marriage model is composed of individuals but has a systemic quality. Each actor observes the set of partners available to him or her, and then chooses the one that offers the highest utility. The choices available to actors on either side of the market are constrained by the preferences of members of the opposite sex, but also by the willingness (or not) of more desirable members of their own sex to match with the partners they themselves would prefer. The marriage model is a system model of the interdependent effects of the preferences and characteristics of all men and women.

4 Inference

Given a stable matching for a population of men and women, our goal is to estimate the extent to which certain observed characteristics of an individual influence the unobserved utilities that other actors have for pairing with that individual. We assume that a set of unknown mean preference parameters θ along with observed female characteristics X, male characteristics Y, and other components give rise to utilities U, V that men and women have for each other, and that the utilities in turn determine preference rankings that men and women have for the members of the opposite sex. The particular rankings vary across the rankers due to unobserved components of utility that will be represented with additional parameters. Individuals form pairs with one another according to their individual preference rankings of those particular members of the opposite sex available to each of them. This in turn gives rise to the data we observe.

Under the assumption that the system is stable, the data tell us of the occurence of two events:

- $M = m$: The matching m was realized by the system of actors.

- $m \in \mathcal{S}$: The matching m is in the set $\mathcal{S}$ of stable matchings for the system.

A full probability model for the data, representing the outcome of the process leading to the stable matching, can therefore be written compactly as $p(M = m, m \in \mathcal{S}|\theta, \psi)$, where ψ represents parameters other than preferences that also govern the process, and the dependence on observed characteristics is left implicit. We factor the full probability model of the data as follows:

$$p(M = m, m \in \mathcal{S}|\theta, \psi) = p(M = m|m \in \mathcal{S}, \theta, \psi)p(m \in \mathcal{S}|\theta)$$

The first component of the probability, $p(M = m|m \in \mathcal{S}, \theta, \psi)$, depends on the particular mechanism or process by which actors formed pairs with one another, which is likely to have involved some strategic behavior that was not a simple reflection of their preferences. Roth and Sotomayor (1990) show that the use of appropriate strategies can affect which particular stable matching is attained, among all those possible for a given configuration of utilities. Since this can benefit certain actors, strategic behavior should be expected. Without knowledge of the matching mechanism, including these strategies, specification of the first component is difficult. However, the difficulty can be avoided if we base our inference only on the second component, the marginal probability of the stability of the observed matching. Roth and Sotomayor (1990) prove that stability, once attained, is a function only of the preferences of the actors, so its maintenance does not depend on strategic considerations. We denote this by distinguishing the preference parameters θ that determine stability from the parameters ψ that involve strategies and other factors besides preferences influencing the matching process. The second component of the probability model provides information on the preferences no matter which stable matching has been attained. Although some information may be lost by basing inference only on this marginal likelihood, we emphasize that as a result, our method of estimation does not depend on the particular process that generated the matching.

Explicitly representing the unobserved utilities in the marginal probability of stability gives

$$p(m \in \mathcal{S}|\theta) = \int p(m \in \mathcal{S}|U, V)p(U, V|\theta)\ dUdV, \tag{2}$$

where dependence on the characteristics X, Y remains implicit. Note that the probability that matching m is stable is either 0 or 1, given U, V, so that the first term in the integral is simply a zero-one indicator of the event that the utilities lie in the constrained space required by the stability of the observed matching. In this paper, we focus primarily on models for $p(U, V|\theta)$ that are linear in the preference parameters (although more general parametric models are discussed in Section 5.3). For example, we might model man i's and woman j's utilities for each other and for being single as

$$\begin{aligned} U_{i,j} &= \alpha' x_{i,j} + \epsilon_{i,j} \\ V_{j,i} &= \beta' y_{j,i} + \gamma_{j,i}, \end{aligned} \tag{3}$$

where α is a vector of mean male preference coefficients for the characteristics $x_{i,j}$ each man i perceives in each woman j, and for the characteristics $x_{i,0}$ he perceives in being single, and β, $y_{j,i}$, and $y_{j,0}$ are defined similarly. The error terms ϵ, γ in the linear models represent the individual deviations from mean utilities discussed earlier. A convenient, initial model for error is that the ϵ's and γ's are i.i.d. from a known distribution centered around zero. With the parametrization above, the estimand of interest is the set of mean preference parameters $\theta = (\alpha, \beta)$.

The linear model of equation (3) provides for pair-specific disturbances representing unmeasured characteristics of each potential match, as evaluated by both the man and woman involved. However, it seems likely that other unmeasured characteristics that are perceived by all evaluators, but are not recorded in the data, would also affect matchings. The net effect of such characteristics can be represented by sets of fixed effects, one for each man and each woman. These are introduced as coefficients of dummy indicators for each man and woman, appearing in the utility functions of the opposite sex. When such fixed effects are introduced for all men and women, a separate constant term for each side's utility function can no longer be separately identified, and so is dropped.

5 Estimation Methods

Our goal is to estimate the preference coefficients θ using the information that the observed matching is stable. Maximum likelihood estimation amounts to finding those values of θ which maximize (2). The maximum likelihood estimation methods of Logan (1998) would be applicable to the problem with minor alterations, except the estimation time associated with those methods makes them impractical here. The difficulty lies in the calculation of $p(m \in \mathcal{S}|\theta)$, which is a complicated integral over a high-dimensional space.

Instead of trying to maximize (2) in θ, we take a Bayesian approach to the estimation: Our inference is based on the posterior distribution $\pi(\theta|m \in \mathcal{S}) \propto \pi(\theta)p(m \in \mathcal{S}|\theta)$, where $\pi(\theta)$ is a probability distribution representing our prior uncertainty about θ. Although the posterior involves the same intractable integral $p(m \in \mathcal{S}|\theta)$, inference is made feasible by approximating the posterior via samples of θ from a Markov chain.

5.1 MCMC

In contrast to many applications, in our model the contributions from different individuals to the likelihood $p(m \in \mathcal{S}|\theta)$ are not multiplicative. Constraints induced by the assumption of stability connect the matches in a complicated way: Given the values of U, V, the probability of stability, $p(m \in \mathcal{S}|U, V)$, is an indicator function equal to unity if condition (1) is satisfied, and otherwise equal to zero. As none of the utilities is observed, $p(m \in \mathcal{S}|U, V)$ must be integrated over U and V in order to calculate $\pi(\theta|m \in \mathcal{S})$, the posterior of θ. Markov chain Monte Carlo provides a method for approximating such integrals: While calculation of $\pi(\theta|m \in \mathcal{S})$ or direct sampling from $\pi(\theta|m \in \mathcal{S})$ is impractical, sampling values of θ, U, V from a distribution approximating $\pi(\theta, U, V|m \in \mathcal{S})$ is feasible. The sampled θ values can be examined marginally to estimate $\pi(\theta|m \in \mathcal{S})$. Such data augmentation is often used to simplify posterior sampling (Tanner and Wong, 1987).

Let $\phi = (\theta, U, V)$ represent a state of the unknowns in our model, which includes both the unobserved utilities as well as the unknown parameters. We use the Metropolis-Hastings algorithm to iteratively construct sequences of ϕ from which we can approximate $\pi(\theta|m \in \mathcal{S})$. Given the most recent state ϕ^b, a new value ϕ^{b+1} is constructed by separately updating each unobserved utility, as well as separately updating each component of the preference parameters $\theta = (\alpha, \beta)$. Each update consists of sampling a proposal value for the utility or preference parameter being updated, then accepting or rejecting the proposal with the appropriate probability. The generation of ϕ^{b+1} from ϕ^b constitutes one scan of the algorithm. As $b \to \infty$, the sampling distribution of ϕ^b approaches the desired posterior $\pi(\phi|m \in \mathcal{S})$, so the empirical distribution of the samples of ϕ from the Markov chain can be used to approximate $\pi(\phi|m \in \mathcal{S})$ (Tierney 1994).

Specific methods of constructing such Markov chains are given below. We first address the simple case involving conjugate linear models for utilities, and then discuss a method for constructing chains for some non-conjugate models.

5.2 Conjugate Models: Two-Sided Probit

In this section, we assume the utilities follow the linear model given by (3), the error terms are independent standard normal deviates, and the preference coefficients α, β are a priori independent and normally distributed: $\alpha \sim \mathcal{N}(\mu_\alpha, \Sigma_\alpha)$, $\beta \sim \mathcal{N}(\mu_\beta, \Sigma_\beta)$. This specification constitutes a two-sided probit model: The probability of each man's match, given his unobserved opportunity set, has the form of a probit model, as does the probability of each woman's match, given her opportunity set.

Utilities: We first consider updating the matching utilities of males. For a particular male i and female j, we sample a proposed utility $U^*_{i,j}$ from its full conditional distribution, which is the distribution of the utility conditional upon the data and the current state of all other utilities and parameters. Sampling from a full conditional distribution is Gibbs sampling, and the acceptance probability of such a sample is always unity.

For normal error terms, the full conditional distribution of $U^*_{i,j}$ given all other unknowns is simply a normal distribution constrained by the other utilities and the stability of the observed matching (cf. Albert and Chib 1993). For example, if man i is matched to woman $w(i)$, then he must have a higher utility for matching with $w(i)$ than for being single or matching with any other woman available to him. Similarly, if man i is single then $U_{i,0}$ must be higher than the utility of being matched to any woman available to him. More precisely, the constrained sampling is as follows:

- if $j \neq w(i)$
    1. if $j = 0$ or $V_{j,i} > V_{j,h(j)}$ (i.e., j is available to i), sample $U^*_{i,j}$ from the normal distribution with mean $\alpha' x_{i,j}$, conditional upon $U^*_{i,j} < U_{i,w(i)}$.
    2. if $V_{j,i} < V_{j,h(j)}$ (i.e., j is not available to i), sample $U^*_{i,j}$ from the normal distribution with mean $\alpha' x_{i,j}$ (unconstrained).
- if $j = w(i)$, sample $U^*_{i,j}$ from the normal distribution with mean $\alpha' x_{i,j}$, conditional upon $U^*_{i,j} > U_{i,0}$ (if $j \neq 0$) and $U^*_{i,j} > U_{i,k}$ for any $k : V_{k,i} > V_{k,h(k)}$.

The utilities $V_{i,j}$ of women for men are sampled the same way, except with V, h interchanged with U, w. Note that sampling the utilities is equivalent to sampling the error terms: For example, sampling $U_{i,j}$ constrained to be less than $U_{i,w(i)}$ is equivalent to sampling $\epsilon_{i,j}$ constrained to be less than $U_{i,w(i)} - \alpha' x_{i,j}$. See Robert (1995) for efficient methods of sampling from constrained normal distributions.

Preferences: Normal priors for regression coefficients are conjugate for linear models with normal errors, so the conditional distributions of α, β given the utilities are normal as well. A straightforward calculation shows:

- $\alpha|_{\beta,U,V,m\in\mathcal{S}} \sim \mathcal{N}(\tilde{\mu}_\alpha, \tilde{\Sigma}_\alpha)$, where

  $\tilde{\Sigma}_\alpha^{-1} = \sum_{i=1}^{n_m} X_i X_i' + \Sigma_\alpha^{-1}$, $\tilde{\mu}_\alpha = \tilde{\Sigma}_\alpha (\sum_{i=1}^{n_m} X_i U_i + \Sigma_\alpha^{-1} \mu_\alpha)$

- $\beta|_{\alpha,U,V,m\in\mathcal{S}} \sim \mathcal{N}(\tilde{\mu}_\beta, \tilde{\Sigma}_\beta)$, where

  $\tilde{\Sigma}_\beta^{-1} = \sum_{j=1}^{n_f} Y_j Y_j' + \Sigma_\beta^{-1}$, $\tilde{\mu}_\beta = \tilde{\Sigma}_\beta (\sum_{j=1}^{n_f} Y_j V_j + \Sigma_\beta^{-1} \mu_\beta)$

in which

- n_m and n_f are the male and female population sizes,
- X_i is a matrix with man i's self-matching characteristics in the first column and the characteristics he perceives in each woman in columns 2 through $n_f + 1$,
- Y_j is the corresponding matrix of women's perceived matching characteristics, with $n_m + 1$ columns,
- U_i is an $(n_f + 1) \times 1$ vector of the utilities of man i for self-matching and for each woman,
- V_j is the corresponding vector of utilities of women.

Such normal distributions are easy to sample from, and so we use them to generate our proposal values α^* and β^*. Since these are Gibbs samples, the acceptance probability of each proposal is unity.

5.3 Non-conjugate Models: Two-Sided Logit

We now consider models of the form

$$
\begin{aligned}
U_{i,j} &= u(x_j, y_i|\alpha) + \epsilon_{i,j}, \\
V_{j,i} &= v(y_i, x_j|\beta) + \gamma_{j,i},
\end{aligned}
\tag{4}
$$

where the functions u and v are determined by the preference parameters α and β, respectively, and the error terms are independent and distributed according to known distributions. One convenient choice for the error distribution is the type 1 extreme value (Gumbel) distribution, having a closed-form CDF, making constrained sampling relatively easy. When specification (3) is substituted for (4), the non-conjugate model with Gumbel disturbances is a completely individual-level, two-sided logit model.

Utilities: Conditional on the preference parameters, the distribution of utilities can be resampled by sampling the errors ϵ and γ, constrained by stability. This proceeds in a manner similar to that in the case of normally distributed errors, as described above.

Preferences: Generally there will be no simple conditional distribution for a preference parameter given the utilities. Therefore, instead of generating parameter updates from a full conditional distribution, we can use a random walk proposal distribution, that is, we sample a proposal value uniformly from a box around the current value. For example, a possible proposal distribution q for α is

$$\begin{aligned} q(\alpha^*|\alpha) &= A^{-1} \qquad \text{for } \alpha^* \in \alpha \pm \delta_1 \\ &= 0 \qquad \text{otherwise} \end{aligned}$$

where δ is a vector of the same dimension as α, and A is the volume of the box spanned by $\alpha \pm \delta_\alpha$. Updates for β can be generated similarly.

Because such a proposal distribution is not the full conditional, the acceptance probability is not necessarily unity, and needs to be calculated. In what follows, note that the errors ϵ, γ together with the preference parameters θ determine the utilities, and so we can base our Metropolis algorithm and acceptance probabilities on these quantities. For example, when updating the α parameter the Metropolis-Hastings ratio, or update probability, is

$$\begin{aligned} \frac{\pi(\theta^*, \epsilon, \gamma | m \in \mathcal{S})}{\pi(\theta, \epsilon, \gamma | m \in \mathcal{S})} \frac{q(\alpha|\alpha^*)}{q(\alpha^*|\alpha)} &= \frac{p(m \in \mathcal{S}|\theta^*, \epsilon, \gamma)\pi(\alpha^*)p(\epsilon)p(\gamma)}{p(m \in \mathcal{S}|\theta, \epsilon, \gamma)\pi(\alpha)p(\epsilon)p(\gamma)} \frac{q(\alpha|\alpha^*)}{q(\alpha^*|\alpha)} \\ &= \frac{p(m \in \mathcal{S}|\theta^*, \epsilon, \gamma)\pi(\alpha^*)}{p(m \in \mathcal{S}|\theta, \epsilon, \gamma)\pi(\alpha)} \qquad (5) \\ &= p(m \in \mathcal{S}|\theta^*, \epsilon, \gamma)\frac{\pi(\alpha^*)}{\pi(\alpha)}, \qquad (6) \end{aligned}$$

where $\theta^* = (\alpha^*, \beta)$. Equality (5) holds because the proposal distribution is symmetric, i.e. $q(\alpha^*|\alpha) = q(\alpha|\alpha^*)$. Equality (6) holds because $p(m \in \mathcal{S}|\theta, \epsilon, \gamma)$ is simply the zero-one indicator variable of whether or not m is a stable matching under the values of θ, ϵ, γ, which must be one since θ was accepted at a previous step of the chain. Therefore, the probability of accepting α^* is $p(m \in \mathcal{S}|\theta^*, \epsilon, \gamma)\pi(\alpha^*)/\pi(\alpha)$, which is $\pi(\alpha^*)/\pi(\alpha)$ if m is a stable matching under θ^*, and zero otherwise. The acceptance probabilities for β are similar.

6 Parametrizations and Identifiability

One formulation of the parametric marriage model described above involves a linear model for the latent, unobserved utilities. The linear model is composed of unknown preference parameters $\theta \in \Theta$ and observed covariates X, Y. We say a linear model is identifiable for the marriage model if for each $\theta_1, \theta_2 \in \Theta$, $\theta_1 \neq \theta_2$, there exists a state of the covariates X, Y such that $p(m \in \mathcal{S}|\theta_1) \neq p(m \in \mathcal{S}|\theta_2)$, where conditional dependence on the covariates is again implicit (this is similar to Manski's (1988) definition of identifiability for logistic regression models). Even though this notion of identifiability is fairly non-restrictive, the complexity of $p(m \in \mathcal{S}|\theta)$ makes it difficult to determine which linear models are identifiable. However, the standard requirements for identification in discrete choice models (see Train 2003) must be met on both sides of our model. As in those models, the scale of the preference coefficients is not identified, and only differences of characteristics across alternatives have identifiable effects.

In addition to the standard identifiability issues of discrete choice models, there may be multiple local maxima of the posterior distribution. These maxima might give rise to a variety of interpretations of the data: For example, consider a population in which women with a relatively high value for some characteristic x tend to be single. Using a univariate version of the linear model (3), this trend in the population could be explained by a male preference for low values of x ($\alpha < 0$), a female self-preference for high values of x ($\beta > 0$), or both.

The relative strengths of men's and women's coefficients on dummy-variable indicators of common group memberships will also be poorly identified, since weaker preferences of either gender can be approximately offset by stronger preferences of the other. Such identifiability problems may be ameliorated somewhat by including other characteristics in the model, or by making a priori assumptions that certain characteristics are desirable (i.e. making assumptions about the signs of the preference parameters). Even though simulation studies, not reported here, give us some confidence that our methods will assign higher probability to the true values of the preference parameters, we note the importance of a careful parametrization of the model.

7 Results

We now analyze data from the 1988 National Survey of Families and Households (Sweet, et al., 1988). The NSFH was a national probability sample of English- and Spanish-speaking men and women living in U.S. households and, like all national surveys known to us, does not contain detailed population information on the persons making up the local marriage market that any of its sample members could have considered. Therefore we use the distribution of characteristics in the sample to approximate the distribution of characteristics in the local marriage markets that individuals were exposed to. We disregard any local or regional differences in marriage markets. We model the current matches of sample members to estimate the preferences of men and women for the age, education levels, and religious affiliations of mates.

Observations are of single householders and married householders whose spouses had completed the spouse questionnaire. The race of both the householder and any spouse was restricted to non-Hispanic White, and the ages of both were restricted to the range 25 through 65. Age and education were measured in years, with education top-coded in the data at 17 years. Cases with missing values on any analysis variables were dropped. To reduce the computational burden, we drew a simple ten percent random subsample of 255 married couples and 164 single persons, a total of 314 men and 360 women.

The NSFH recorded the religious affiliations of respondents and spouses using 65 detailed denominations, which we recoded into the five categories No Religion, Mainline Protestant, Conservative Protestant, Catholic, and Other Religion, adapting categories defined by Lehrer and Chiswick (1993). The Other category contained religions with small numbers of observations in the final data set, including Jewish, Mormon, Orthodox, and East Asian. Five dummy variables were created to indicate whether each possible pairing of men and women would constitute a match within the same religion. Following Lehrer (1998), who based her procedure on intermarriage patterns observed by Lehrer and Chiswick (1993), we used different rules across the religions to define what would be considered a marriage within the same religion. For conservative

Protestant denominations and members of Other Religions, we required marriage within the same, detailed denomination. For mainline Protestants, a marriage between two persons of different denominations, but still within the mainline category, counted as the same religion. The No Religion and Catholic labels were single denominations in the data, and needed no special rule. To clarify by an example, our rules counted marriages between Episcopalians and Lutherans, two mainline denominations, as homogamous, but counted marriages between members of the conservative Baptist and Pentacostal denominations, or between Jews and Mormons, as heterogamous. Using NSFH retrospective data on first marriages, Lehrer and Chiswick (1993) found homogamous conservative Protestant marriages to be temporally more stable than homogamous mainline Protestant and Catholic marriages; we are interested in whether our estimates of preferences for same-religion partners, based on the NSFH cross-sectional data, will show similar differentials.

Age and education distributions in the NSFH data were presented graphically in Figure 1 above. The scatterplots suggest a tendency of actors to form matches with others having similar age and education levels. This could reflect preferences of actors to choose mates like themselves. For this reason, we consider a variant of the model presented in equation (3), using the symmetric specification for men's and women's utilities shown in Table 1.

Age and education differences are defined as the potential mate's values minus the evaluator's. Corresponding squared differences are also specified. There are five separate dummy indicators for matches within the same religions. All variables are defined to take the value zero for self-matching, so the self-matching utilities are equivalent to their disturbances as defined in equation (3). All disturbances are independent standard normal random variables, so the model is two-sided probit.

Two variants of the model were estimated, as described below. In each case, two Markov chains of length 2×10^6 scans were constructed using the methods of Section 3.2, having different random seeds, but each having starting values of zero for all regression coefficients. The prior distributions for the regression parameters were taken as independent normal distributions, each with a mean of zero and variance of 100. The state of each Markov chain was recorded every 200 scans. The chains indicate slow mixing, but approximate stationarity by about 500,000 scans.

Means and standard deviations of preference parameters were calculated using the pooled scans after the

Table 1: Posterior means and standard deviations of men's and women's preference parameters.

Parameter	No fixed effects Men	No fixed effects Women	Fixed effects Men	Fixed effects Women
Constant	-1.491*	-1.413*	—	—
	(0.097)	(0.099)		
Age difference	-0.043*	0.031*	-0.031	0.069*
	(0.015)	(0.016)	(0.021)	(0.021)
(Age difference)^2	-0.007*	-0.006*	-0.008*	-0.010*
	(0.001)	(0.001)	(0.002)	(0.002)
Education difference	0.004	-0.008	0.024	0.010
	(0.023)	(0.023)	(0.033)	(0.039)
(Education difference)^2	-0.013*	-0.013*	-0.016*	-0.032*
	(0.006)	(0.006)	(0.008)	(0.008)
No religion	0.448	0.420	0.299	0.406
	(0.388)	(0.401)	(0.547)	(0.538)
Mainline Protestant	0.449*	0.396*	0.421*	0.572*
	(0.128)	(0.133)	(0.198)	(0.191)
Conservative Protestant	0.855*	0.764*	1.212*	1.530*
	(0.174)	(0.178)	(0.277)	(0.287)
Catholic	0.560*	0.501*	0.697*	0.849*
	(0.151)	(0.155)	(0.225)	(0.231)
Other religion	1.592*	1.478*	2.601*	2.090*
	(0.362)	(0.359)	(0.735)	(0.600)

Table 2: Rough comparison of averaged estimates with Lehrer and Chiswick (1993).

Study	No Religion	Mainline	Conservative	Catholic
Present study	.353	.497	1.366	.773
Lehrer & Chiswick	.257	.822	1.004	.911

first 500,000 for each Markov chain, and are presented in Table 1. Asterisks indicate estimates greater than 1.96 standard deviations in absolute value. Because of the pooling, standard deviations reflect both within- and between-chain variability. The value of $\sqrt{\hat{R}}$, a measure of the ratio of between-chain to within-chain variance (Gelman 1996), was calculated for each parameter. Values close to 1.0 indicate that between-chain variability is not high, based on what would be expected from the variability within chains. The values for all coefficients in the table are all below Gelman's suggested cutoff of 1.2.

The first two columns of Table 1 report estimates from the complete sample of 255 married couples and 164 single persons described above, omitting fixed effects. The last two columns are estimates from a fixed effects specification that also accounts for characteristics of men and women that are assumed to be perceived by all members of the opposite sex but are not included in the observed data. This is our preferred model because we believe that some unmeasured variables such as wealth, earning ability and beauty are important for marriage. The fixed effects estimates are roughly similar to those in the first two columns. Men prefer a negative age difference, meaning a younger wife, while women prefer an older husband; the point estimate of the absolute strength of the women's preference is more than two times the men's. Both men and women have significant and similar preferences for smaller squared differences in age and education with their spouses. Neither gender prefers the other to have a level of education higher or lower than their own.

The estimates of the men's and women's preferences for a spouse of the same religion are mostly strong, but their relative sizes across genders are poorly identified for the reason mentioned in section 6. Taking their average values across genders as summaries, the categories fall in this order of increasing preference for homogamy: No religion (0.353), mainline Protestant (0.497), Catholic (0.773), conservative Protestant (1.371), Other religion (2.346). A direct comparison to Lehrer and Chiswick's (1993) estimates of the effects of homogamy on *temporal* stability in the NSFH data is made difficult because of their use of a different method (proportional hazards), different data (retrospective first marriages), and a different categorization. However, after dropping the incompatible Other Religion category, and averaging Lehrer and Chiswick's two separate coefficients for mainline Protestant homogamous matches, as defined for our analysis, we regressed our four estimates on theirs, and used the regression coefficients to rescale their estimates to rough comparability with our own, as seen in Table 2. The order of their estimated effects of homogamy on the temporal stability of marriages, using retrospective data, agrees with our estimates of preferences for same-religion matches, using cross-sectional data. We note a larger difference between the conservative and mainline Protestants in our results, with our estimate for conservative Protestant homogamy 2.76 times that for mainline Protestants, compared with Lehrer and Chiswick's factor of 1.22. Though these are different types of estimates than Lehrer and Chiswick's, we are reassured by their rough similarity. As discussed earlier, we believe that a slow overall rate of change in the determinants of stable matches in the game theoretic sense should also tend to produce a degree of temporal stability, as studied by Lehrer and Cheswick.

Leaving aside the coefficients for Other Religion, which apply to a relative handful of cases, the estimated preferences of women tend to be stronger than men's. However, preference coefficients are implicitly measured against the assumed common scale of the disturbances in the men's and women's utility functions. Since the model provides no empirical basis for asserting that the disturbances have the same scale across the genders, it is inappropriate to make direct cross-gender comparisons of the absolute magnitudes of the effects (cf. Train 2003 on comparing coefficients of one-sided models across markets). One interpretation of the overall magnitude differences is that men and women may have similar relative preferences for age, education, and religion, but that the actual scales of the disturbances may differ. That is, the effects of the measured variables compared with the disturbances appear larger for the women; there is less to be explained by the disturbances for the women, once men's education, age, religion and fixed effects are taken into account. Notice that this pattern is present only in the fixed effects estimates. Since the women's estimates increase in absolute size relative to the disturbances when fixed effects are introduced, it seems reasonable to conclude that fixed effects — characteristics observable by all members of the opposite sex, but not in the data — are more important in explaining women's preferences. These may reflect observable differences in men's earning power, wealth, or status, for example.

The combined linear and squared effects of age and education preferences of men and women are shown graphically in Figure 3, with age preferences shown by solid lines and education by dotted. The net effect of the linear and squared age terms is that men prefer women about 1.9 years younger, and women prefer men about 3.5 years older than themselves. The combined effects of absolute and relative preferences for

Figure 3: Estimated Age (——), Education (– –), and Religious Preferences of Men and Women.

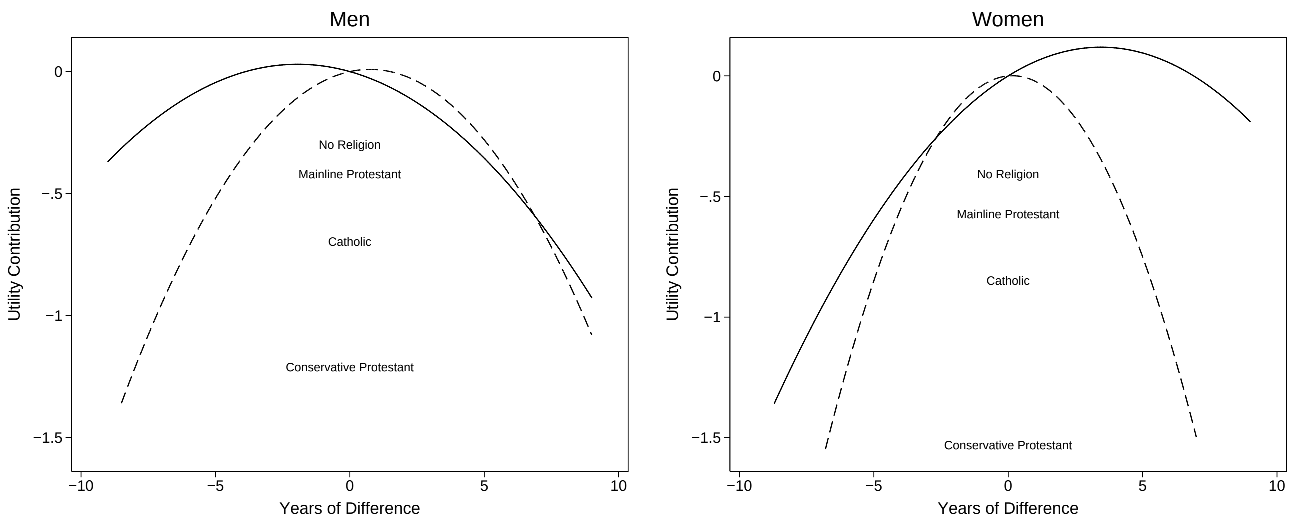

education are similar across genders: Men and women prefer partners with roughly the same education levels as themselves, but the point estimate of the men's preference is weaker.

Figure 3 also compares religious homogamy preferences with the age and education curves by plotting estimates for the religions at vertical locations equal to minus one times their Table 1 values. Plotted this way, the locations indicate the age- and education-difference equivalents of a potential mate *not* sharing the same religion as the evaluator. For example, consider a mainline Protestant man evaluating a woman of the same age and education as his own. The left panel of the figure shows that the contribution of her age and education to his utility for her is zero. If she is of a different religion, however, her overall contribution is reduced by the vertical distance between zero on the axis and the location of the mainline Protestant marker; this is a reduction of 0.421 utility points (taking the exact value from Table 1). That is the same reduction in utility that would pertain if the woman were about 7.3 years older or younger than the ideal age, or 5.1 years different from the most preferred education level. For mainline Protestant women, the equivalent age and education differences for heterogamous mates are 7.6 and 4.2 years respectively. The effects are much stronger for conservative Protestants, for whom finding a mate in the same, detailed denomination is important enough to offset an age difference from the ideal of about 12.3 (12.4) years, or an education difference of about 8.7 (6.9) years for men (women). Such a strong difference between estimates for conservative and mainline Protestants did not appear in Lehrer and Chiswick's retrospective temporal stability estimates (Table 2). Whether this reflects a stronger differential in contemporary rather than historical preferences might be investigated in a subsequent study, but space limitations preclude it here.

8 Discussion

The two-sided method we propose for estimating mate preferences from typical marriage data represents reality in a simplified form. We use the distribution of characteristics in the sample to approximate the distributions of characteristics in the marriage markets to which different individuals are exposed, and then estimate separate mean preference coefficients for men and women, allowing for random variability of utilities at the individual level and providing fixed effects for characteristics observed by the opposite sex but not included in our data. We now consider under what circumstances this approach might be justified, and what changes in the model could be considered when those circumstances do not apply.

Our intent is not to model the full process that led individuals in our sample to marry, but instead to estimate their mate preference coefficients using the stability assumption, that all men and women are matched voluntarily. We use the other sample members as stand-ins for the unknown potential mates that each person might have had available for choosing. Doing this misrepresents the actual choices available to each person in two possibly important ways. First, the size of the modeled choice set may be different in size from typical real choice sets. We have done the estimation so that each man (woman) has a choice set composed of as many of the 360 women (314 men) as prefer him (her) to their current matches. (The actual number of choices each faces depends on his (her) desirability.) Does the assumption that there are somewhat more than three hundred known members of the opposite sex influence the results? Simulations, not presented here, show little dependence of the parameter estimates on differences between the assumed size of the known population and the actual size, within broad limits. It is also reassuring to note a recent estimate that the number of acquaintances for both men and women in the general population has a mode of 610 and mean of 750, meaning that each gender knows roughly 350 members of the opposite sex, which approximates our sample size (Zheng, Salganik and Gelman 2005).

The second way in which our use of sample members as stand-ins misrepresents the actual choices available to respondents is that the locales in which actual respondents choose mates are not random samples of the larger population. A main difference between real locales and the modeled ones is in the greater social homogeneity of the former. It is also true that mate preferences presumably differ across locales and population subgroups, which we do not model. However, it is possible to accommodate the latter kind of variability by interacting indicators of locale or subgroup with the evaluated characteristics as in the usual regression style of analysis. More detailed behavioral models could be constructed by combining the present estimation methods with more specific ideas of choice set formation and information flow. The present model's assumption that the whole set of sample members can be taken to represent local marriage markets should generally not be replaced, however, by deterministic and equally false assumptions, such as that marriages can only be made within the confines of particular communities, groups or networks. Instead, the distribution of information about potential partners should be modeled probabilistically.

The goodness of fit of our model can be assessed by comparing characteristics of the observed matching to characteristics of matchings obtained from the posterior predictive distribution. To generate a random matching from the posterior predictive distribution, preference parameters α_{pred} and β_{pred} are sampled from the Markov chain and then combined with the observed characteristics X and Y to generate a set of utilities U_{pred} and V_{pred} according to model (3). Generating a matching M_{pred} from U_{pred} and V_{pred} requires specification of a particular matching process, a specification we have been able to avoid in our estimation procedure. Here, however, game theory provides help in bracketing the types of matchings that can occur: It is known that there are two extreme types of matchings for any given set of utilities. One is optimal for men in the sense that no man would prefer any other stable matching to it, and the other is optimal for women in the same sense, roles reversed. Any other stable matching must lie between these two with respect to the preferences of all men and women. If these two stable matchings are identical, then it is the only matching that is stable for the given set of utilities. A simple algorithm due to Gale and Shapley (1962) will generate either the male- or female-optimal matching for any set of utilities depending on which gender is accorded a certain role. Using this algorithm, we can generate both of these extreme matchings $\{M^m_{\text{pred}}, M^f_{\text{pred}}\}$ from U_{pred} and V_{pred}, and compare features of these matchings, such as the joint distributions of spousal attributes, to the same features of the observed matching.

A number of extensions to the proposed approach seem promising for future work. One benefit to modeling the interplay of preferences and characteristics of mates explicitly is the possibility of using the preference estimates for counterfactual projections of marriage rates that would result from exogenous changes in distributions of potential mates. This possibility is still to be explored, but remains a strong motivation for the type of analysis undertaken. We note that the proposed model might also be used to adjudicate between different theories of mate selection, using alternative specifications of the preferences on the two sides of the market. For example, Buss (1989) and colleagues conducted reported-preference studies in 37 cultures, finding evidence that men and women value traits of partners asymmetrically, consistent with differentials predicted by an evolutionary biological theory. Another extension of the model would relax the cross-sex mating requirement, allowing both hetero- and homosexual matches, with a same-sex indicator variable similar to the same-religion indicator used here. Our general model and modifications to it could prove useful for a variety of matching or choice problems in which the options available to each actor are limited by the preferences of others.

References

Albert, J. H., and Chib, S. (1993), "Bayesian Analysis of Binary and Polychotomous Response Data," *Journal of the American Statistical Association*, 88, 669–679.

Bergstrom, C. T., and Real, L.A. (2000), "Towards a Theory of Mutual Mate Choice: Lessons from Two-Sided Matching," *Evolutionary Ecology Research*, 2, 493–508.

Bumpass, L.L., and Sweet, J.A. (1972), "Differentials in Marital Instability: 1970," *American Sociological Review*, 37: 754-766.

Buss, D.M. (1989), "Sex differences in human mate preferences: Evolutionary hypotheses tested in 37 cultures," *Behavioral and Brain Sciences*, 12, 1–49.

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

Gelman, A. (1996), "Inference and monitoring convergence," in W.R. Gilks, S. Richardson and D.J. Spiegelhalter, eds., *Markov Chain Monte Carlo in Practice*, New York: Chapman and Hall, 131–143.

Johnson, R.A. *Religious Assortative Mating in the United States*, New York: Academic.

Kalmijn, M. (1998), "Intermarriage and Homogamy: Causes, Patterns, Trends," *Annual Review of Sociology*, 24, 395–421.

Lehrer, E.L. (1998), "Religious Intermarriage in the United States: Determinants and Trends," *Social Science Research*, 27, 245–263.

Lehrer, E.L. (2004), "Religion as a Determinant of Economic and Demographic Behavior in the United States," *Population and Development Review*, 30, 707-726.

Lehrer E. L., and Chiswick, C. U. (1993), "Religion as a Determinant of Marital Stability," *Demography*, 30(3), 385-404.

Logan, J. A. (1996), "Opportunity and Choice in Socially Structured Labor Markets," *American Journal of Sociology*, 102 ,114–160.

Logan, J. A. (1998), "Estimating Two-Sided Logit Models," *Sociological Methodology*, 28, 139–173.

Manski, C. (1988), "Identification of Binary Response Models," *Journal of the American Statistical Association*, 83, 729–738.

Mare, R.D. (1991), "Five Decades of Educational Assortative Mating," *American Sociological Review*, 56, 15–32.

Raley, R. K., and Bratter, J. (2004), "Not even if you were the last person on earth! How marital search constraints affect the likelihood of marriage." *Journal of Family Issues*, 25, 167–181.

Robert, C.P. (1995) "Simulation of Truncated Normal Variables," *Statistics and Computing*, 5, 121-125.

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

Roth, A. E., and Vande Vate, J.H. (1991), “Incentives in Two-Sided Matching With Random Stable Mechanisms,” *Economic Theory*, 1, 31–44.

Schoen, R., and Wooldredge, J. (1989), “Marriage Choices in North Carolina and Virginia, 1969-71 and 1979-81,” *Journal of Marriage and the Family*, 51, 465-481.

South, S.J. (1991), “Sociodemographic Differentials in Mate Selection Preferences,” *Journal of Marriage and the Family*, 53, 928–940. [NSFH.

Sprecher, S., Sullivan, Q., and Hatfield, E. (1994), “Mate Selection Preferences: Gender Differences Examined in a National Sample,” *Journal of Personality and Social Psychology*, 66(6), 1074–1080.

Stevens, G., Owens, D., and Schaefer, E.C. (1990), “Education and Attractiveness in Marriage Choices,” *Social Psychology Quarterly*, 53, 62-70.

Sweet, J. ,Bumpass, L., and Call, V. (1988), ”The Design and Content of The National Survey of Families and Households.” Center for Demography and Ecology, University of Wisconsin-Madison, NSFH Working Paper #1.

Tanner, M. A., and Wong, W.H. (1987), “The Calculation of Posterior Distributions by Data Augmentation,” *Journal of the American Statistical Association* 82,528–540.

Tierney, L. (1994), “Markov Chains for Exploring Posterior Distributions” (with discussion), *The Annals of Statistics*, 22, 1701–1762.

Train, K. (2003), *Discrete Choice Methods With Simulation*, Cambridge: Cambridge University Press.

Zheng, T., Salganik, M.J. and Gelman, A. (2005), “‘How many people do you know in prison?’: Using overdispersion in count data to estimate social structure in networks,” unpublished manuscript, Department of Statistics, Columbia University, June 13, 2005.