AGGREGATE AGE-AT-MARRIAGE PATTERNS FROM INDIVIDUAL MATE-SEARCH HEURISTICS*

PETER M. TODD, FRANCESCO C. BILLARI, AND JORGE SIMÃO

The distribution of age at first marriage shows well-known strong regularities across many countries and recent historical periods. We accounted for these patterns by developing agent-based models that simulate the aggregate behavior of individuals who are searching for marriage partners. Past models assumed fully rational agents with complete knowledge of the marriage market; our simulated agents used psychologically plausible simple heuristic mate search rules that adjust aspiration levels on the basis of a sequence of encounters with potential partners. Substantial individual variation must be included in the models to account for the demographically observed age-at-marriage patterns.

In modern Western societies, deciding when to get married seems like a highly personal and individual choice. Individuals may believe that they are considering options and weighing possibilities that nobody else has ever had to think about in quite the same way. Yet, much research has pointed out the societal and economic constraints that influence even these personal decisions (e.g., Lloyd and South 1996). Indeed, when viewed from the aggregate level, the pattern of the age at which people first get married shows surprising regularity across populations (Coale 1971). Somehow, what people are doing in the mating game at the individual level seems to be following systematic rules that generate distinct patterns at the population level. But how? And how can we find out?

The scientific study of marriage has done little to answer these questions because of a strong division in focus among fields. A long tradition of sociological and demographic research has gathered and analyzed data on aggregate population-level patterns, such as age at marriage and proportion ever marrying, in cohorts from different historical and geographic settings. But this top-down macro perspective has typically obscured (or has not considered) how each individual makes a choice. Psychologists and economists, on the other hand, have studied and modeled the (often heterogeneous and culturally varying) individual-level processes that can end in the decision to cohabit or marry. But this bottom-up micro view has omitted the patterns that emerge in a group of such deciding individuals. Given that the two perspectives, individual and group level, have data and hypotheses that can help to constrain and explain the other, we should find a way to bring them together to speak to each other.

One common language that could connect both perspectives is that of mathematics. Building mathematical models has been done with some success (see, e.g., Coale and McNeil 1972; Diekmann 1989), but with a certain degree of violence done to the assumptions at both the micro and macro levels. In particular, allowing for significant variation in the strategies used by individuals quickly makes the mathematical models of their interactions intractable. As we argue in this article, it is exactly such individual-level variation that may underlie the emergence of the observed patterns at the population level. Thus, in

*Peter M. Todd, Center for Adaptive Behavior and Cognition, Max Planck Institute for Human Development, Lentzeallee 94, 14195 Berlin, Germany; E-mail:ptodd@mpib-berlin.mpg.de. Francesco C. Billari, Institute of Quantitative Methods and Innocenzo Gasparini Institute for Economic Research, Bocconi University, Milan. Jorge Simão, Computer Science Department, FCT/New University of Lisbon.

addition to being a challenging language in which to become fluent, mathematics may be inadequate for expressing the relationships that are crucial to understanding the micro/macro interactions in the marriage market.

Instead, we turn here to computer modeling as a lingua franca to foster communication between the top-down and bottom-up approaches. Agent-based simulation models that specify the mate-search and choice behavior of individual agents interacting in a group enabled us to capture and explore the impact of the vital variation that is often missing from mathematical models. We did so by controlling and monitoring the micro-level decision mechanisms of each agent and observing the patterns that emerged at the macro level as a consequence of their choices and interactions. This modeling approach is finding increasing application in the social sciences and beyond, enabling, as it does here, different previously separated research traditions to come together and illuminate each other (Epstein and Axtell 1996; Gilbert and Troitzsch 1999; Macy and Willer 2002). Although agent-based modeling has not yet become widespread in demography, the study of demographic behavior could benefit significantly from this approach (Billari and Prskawetz 2003).

In the rest of this article, we present our efforts to combine demographic and psychological approaches to marriage via agent-based modeling. Our aim is to explain the emergence of commonly observed patterns of age at first marriage (Coale 1971) as an outcome of the interaction of many instances of individual decision-making behavior. While marriage-age patterns have been explored from other perspectives in the past, our explanation here is novel in that it aims to do more than just account for the demographic data—we also require our models to meet the additional constraints of being psychologically plausible and fitting to other data on individual mate-choice behavior. We start with population-level empirical evidence on the distribution of ages at marriage and review existing explanations of the common invariant features of this distribution across cultures. We then take the bottom-up approach and simulate the behavior of a cohort of *satisficing* agents who are looking for (marriage) partners in situations of both one-sided and mutual-choice decision making. We find that plausible psychological mechanisms of choice that are suggested by the framework of *bounded rationality* need some refinements to be reconciled with the macro patterns of marriage choice. In particular, we show how population heterogeneity in strategies is compatible with observed macro patterns. As will become clear, the implications of our results open a wide space for future research developments, both on the side of empirical studies and on the side of agent-based modeling of social behavior.

THEORETICAL ACCOUNTS OF AGE PATTERNS OF MARRIAGE

The distribution of ages at which people first marry has been similar, at least in a qualitative way, across a broad range of geographic locations and historical periods (Coale 1971; Coale and McNeil 1972). After rising quickly from a minimum marriage age, this distribution follows a rough bell shape, with a long tail capturing people who marry late in life. Whereas Coale and McNeil studied this common pattern using the frequency distribution of age at marriage, here we use the more behaviorally relevant *hazard rate* of marriage. This rate, defined either in discrete or continuous time, is the probability of marriage (or density in the continuous-time case) conditional on the fact that an individual has not married before a certain exact age.

To illustrate the shape of these hazard rates for marriage, we show in Figure 1 the empirically observed functions[1] for men and women in three populations in the late

1. More specifically, we graphed the age-specific conditional probabilities of first marriage, that is, the number of first marriages of people who attained a given age x in a year divided by the number of still-unmarried individuals of age $x - 1$ at the beginning of the year.

Figure 1. Hazard Functions for Marriage in European Populations

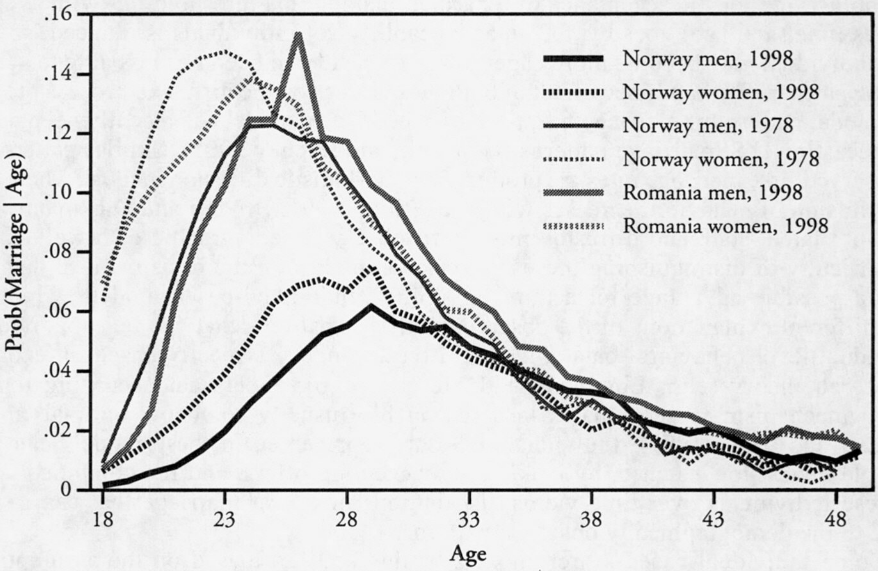

Source: Authors' elaborations on Eurostat, New Cronos database.

twentieth century: Romania in 1998, and Norway in 1978 and 1998. In all the cases shown in the figure, notice that the rise of age-specific probabilities is faster than its decrease. Although the shape of the curve looks somewhat different for Norway in 1998, where nonmarital cohabitation was widespread, it can still be described qualitatively in a similar way. In addition, hazard rates tend to converge to a level close to zero at later ages. This typical hazard-rate function can be observed for several other populations, and it is this overall pattern that we want to account for in our models.

Three main types of formal behavioral models have been proposed to explain these age patterns of marriage (Diekmann 1989): latent-state (or compartment) models, diffusion models, and search models. Such models are usually applied to analyze the behavior of a cohort of individuals as they age. Latent-state models of first marriage, like Coale and McNeil's (1972) model, hypothesize that individuals in a cohort pass through various stages in early adult life before they get married and that the length of time this process takes is governed by a stochastic process. More precisely, Coale and McNeil proposed that the age at entry into the marriageable state is normally distributed and that there are three subsequent exponentially distributed delays (corresponding to life stages) before marriage. Although the Coale–McNeil model fits observed data for a complete cohort or population well, it performs less well in the case of forecasting the behavior of a cohort by means of extrapolation (Goldstein and Kenney 2001; Henz and Huinink 1999), perhaps because of weakness in the behavioral assumptions of the model. The model has also been criticized for the absence of explicit assumptions regarding the workings of the search process (Burch 1993; Coale and Trussell 1996).

In diffusion models, mating happens by "contagion" from other people who are already mated. The model developed by Hernes (1972) is based on the idea that (first) marriages are influenced by two opposing forces that drive a cohort through a diffusion process. First, the pressure to marry increases with age because of the existence of social norms stating that "who marries late marries ill." Such norms are supposed to influence the threshold value for the acceptance of a partner, making this threshold fall with age. Second, as time (and age) goes by, the "marriageability" of individuals is reduced, so while each individual may become more eager to marry, he or she becomes less able to secure a willing partner. The combined effect of both these forces on the diffusion process produces a unimodal pattern for marriage hazard rates. The Hernes model has recently been applied to forecasting U.S. marriage patterns (Goldstein and Keeney 2001). Similar patterns that fit observed first-marriage rates are produced by log-logistic diffusion models, where again the diffusion of behavior decreases with age (Billari 2001; Brüderl and Diekmann 1995).

Both latent-state and diffusion models fit observed demographic data well. Indeed, the difficulty of distinguishing these macro-level models on the basis of fit to data indicates a possible advantage for a third possibility: individual-based models, which provide different explanations of the observed marriage patterns that can be supported with individual-level behavioral data as well. In these models, typically based on economic job-search theory (e.g., Lippmann and McCall 1976), agents act according to some search mechanism to seek mates in a reasonable (usually somehow optimal) manner (Burdett and Coles 1999). Individuals who are represented in these models can select possible mates, for instance, by making and accepting offers, and the combined actions of these individuals over time yields distributions of age at marriage that can be compared to the demographically observed patterns.

Some individual-based search models of this kind are based on the assumption of perfectly rational agents performing optimal searches; they also typically assume that the agents are homogeneous in terms of their rational behavior. Keeley (1979), for example, adopted an optimal model from job-search theory in which individuals set a threshold financial value for the minimal (monetary) benefit they seek in a marriage; if an individual finds a partner with whom their combined income can exceed this threshold, then they marry. The cost-benefit analysis that is necessary to set such an optimal search threshold assumes full knowledge of the environment of potential mates and full rationality on the part of the individual, and thus this model is subject to the common criticisms of such unrealistic assumptions (Chase, Hertwig, and Gigerenzer 1998; Oppenheimer 1988): real human decision makers have only limited knowledge of the situation they face (here, the distribution of values of possible available mates and their own value on the marriage market), limited ability to process whatever amount of information they do have, and limited time within which to make a decision. The presence of these limitations implies that we should build specific models of individual marriage-search processes starting from the assumption that individuals act according to *bounded* rationality, as we describe in the next section.[2]

MODELING SEQUENTIAL SEARCH PROCESSES

To construct an agent-based model to account for population-level demographic phenomena that are related to age at first marriage, we can create a set of simulated individuals who go about trying to marry (or mate) and monitor their success (or lack thereof) over time. Essentially, we want these agents to live out a life composed of the following steps:

2. It is important to keep clear the distinction between the nature of the decision mechanisms that humans use and the way that these mechanisms came about. Individuals commonly rely on choice mechanisms with limited information use and processing, rather than employ complex optimizing processes to reach decisions. However, the mechanisms that humans use are themselves likely to have arisen through processes that are more akin to (constrained) optimization, namely, biological or cultural evolution (Macy and Flache 1995). This evolutionary origin can result in our limited psychological mechanisms being nonetheless well fit to the situations and environments in which they are used, a match between mind and world termed *ecological rationality* (Gigerenzer and Todd 1999).

first, grow up until they reach the minimum marriageable age, possibly learning something along the way that will aid in their later marriage process; second, start to look for a marriage partner; third, if an acceptable (and agreeing) marriage partner is found, marry and leave the still-unmarried population, otherwise get a bit older, possibly learn something from the failed experience, and (if not too old) return to Step 2 to look again. We will record the age at which each individual first gets married (note that there are only first marriages in this version of the model) and the overall number of individuals who ever get married and then compare these data with the empirically observed facts to see how well this model fits. To be concrete, this model requires a specification of the way in which potential marriage partners are met and of how an individual searches through the potential partners. What are the possibilities that we should consider?

We start by specifying the nature of the environment in which the marriage process takes place—that is, how potential partners are encountered. One approach would be to say that all the potential partners are simultaneously available to an individual who is seeking to marry, and the individual must just compare them and choose the one who most closely matches some preference. This is the view of the marriage market proposed by some economists who are interested in how stable matchings can be made between men and women who have complete knowledge of all available partners (Bergstrom and Real 2000). Although this full-knowledge assumption may apply to some small societies, it does not seem to match most of the cases of large populations for which demographers have collected age-at-marriage data. Instead, people who are seeking mates (or other things, such as houses, jobs, and even consumer products) often must choose between a set of options that they see not all at once, but one after another, sequentially. These situations are typically characterized by low (or zero) probability of being able to recall, or return to and choose, previously seen options once they have been passed by (e.g., individuals one has dated and broken up with in the past are probably not still interested in rekindling the relationship later). The problem then becomes one of deciding when to stop searching and to go with the currently available option.

Given this environment for marriage decisions, what kind of search mechanisms can people use to make their choices and stop their hunt? Again there are two main types of approaches. Inspired by the optimizing perspective of unbounded rationality mentioned in the last section, one could attempt to gather as much relevant information as possible about the distribution of available partners and then choose in a way that maximizes the chance of getting the best mate. For instance, one could attempt to compute the optimal point at which to stop searching, given the trade-off between time and other costs that accumulate with each alternative seen, and the chance that the next alternative that is checked will be better than those that were encountered previously. This approach could involve extensive calculations, such as Bayesian updating of probability estimates or assessments of the costs of forgone opportunities, and thus require considerable time and computational resources.

But to make choices in a useful amount of time, real agents must use a limited search across options because real decision makers have only a finite amount of time, knowledge, attention, or money to spend on a particular decision (Todd 2000). And indeed, there is considerable evidence that people who are faced with sequential search tasks use simple rules to make their choices (Dudey and Todd 2002; Hey 1982; Moon and Martin 1990; Seale and Rapoport 1997). As such, people are acting in accordance with what Simon (1990) called bounded, rather than unbounded, rationality—making decisions within the bounds of time, information, and computational ability that the task environment and human cognitive capacities impose on them. The notion of unbounded rationality, following the tenets of logic and probability theory, is a convenient fiction for constructing mathematical models of economic behavior, but to understand real human behavior, one should construct models of the actual bounded psychological processes that guide decision making. Moreover, the simple heuristics that people often use to make decisions with limited time and information not only are easier and faster to employ, they also can be surprisingly successful when applied in the proper task environments (Gigerenzer, Todd, and the ABC Research Group 1999).

Simple search mechanisms require a quick and easy way to decide when to stop looking for options, that is, a stopping rule. What kinds of simple stopping rules are reasonable for a marriage model? For realistic search situations in which the distribution of available options (here, potential mates) is not known or well characterized and the costs of a search (here, the loss of all other opportunities) cannot be accurately assessed, traditional rational models cannot be readily applied and optimal stopping points cannot be calculated. Instead, for such decision problems, Simon (1990) proposed a *satisficing* approach to searches, in which individuals check successive alternatives until they find one that is good enough (rather than optimal) for their goals. This approach can be implemented by means of an aspiration level that individuals somehow set and then use in searching further, stopping that search as soon as an option is encountered that exceeds the aspiration level. Here, we assume that all potential marriage partners can be assessed on some unidimensional quality scale, so that the searchers can set a quality (or mate value) aspiration level for stopping a search upon finding a suitable aspiration-exceeding partner. The exact way in which the aspiration level is set depends on further details about the search situation that is encountered. (Of course, cultural norms and individual emotions exert strong influences on the search for and choice of a mate. Although we do not explore their roles here, both could operate in the search process as we present it, for instance, by affecting the aspiration level that is set and by indicating when an aspiration level has been met, as in falling in love.)

One way to conceive of the search for a marriage partner is as a shopping expedition in which potential mates are encountered one by one, choices must be made on the spot (no recall), and the final choice is made in a unilateral fashion by the searcher—the partner has no say in the marriage decision. Here, mate search can be characterized as one-sided, which is clearly a somewhat unrealistic simplification of the mate-search process for most (if not all) cultures, but which has proved useful as a starting point for mathematical modeling (cf. the *secretary* or *dowry problem* in probability theory; see Ferguson 1989; Seale and Rapoport 1997). Todd (1997; see also Todd and Miller 1999) showed that a simple satisficing heuristic would do well at finding good partners in this situation: set an aspiration level at the highest-quality mate that one has seen during the first dozen or so potential partners, continue searching until a new partner is seen who exceeds that level, and select (marry) that person. Billari (2000) tested whether such a simple rule would produce the expected age-at-marriage patterns by recording how long the search would go on until a suitable partner was found and mapping that search length onto age.[3] He found that if all individuals used the same length of "learning time" to set their aspiration level (e.g., after seeing exactly 12 potential mates), age at first marriage peaked unnaturally immediately after the initial learning phase; only when variation in learning times was introduced did the typical right-skewed bell distribution appear.

Two-Sided (Mutual) Search Processes

The reason that one-sided mate searches, with one sex doing the searching and making the decisions, is an unrealistic model is that in many modern cultures, mate search is

3. Note that here choosing a mate is equated with marriage, and the number of potential partners seen is equated with age. The second mapping, from partners to age, is an assumption of this modeling work that needs to be tested against demographic data; alternatively, the simulations could be compared with data on the number of partners before, rather than age at, first marriage. But for now, given the difficulty of obtaining such data, we make the reasonable assumption that a linear relationship exists between age and number of partners.

mutual: at the same time one sex is evaluating members of the other sex as prospective mates, they are themselves being evaluated in turn. If a particular man does not meet the standards of a particular woman in whom he is interested, for instance, then his courtship attempts are doomed to failure. Furthermore, in contrast to the solipsistic lone-searcher model (e.g., Todd 1997), searching individuals interact in the real world, at a minimum because they are vying for the same set of potential mating partners. How can we model this more realistic two-sided mutual-choice situation?

To explore how different mate-search rules can work in a two-sided setting,[4] we start with a population containing two sets of searchers, 100 simulated men and 100 simulated women, each with a distinct mate value between 0.0 and 100.0 and all in competition with one another (within each sex) for the same set of possible mates. Each individual has the ability to assess accurately the mate values of members of the opposite sex, but (initially) lacks any knowledge of his or her own mate value. Each individual begins his or her simulated life by assessing and making (or not) practice marriage or mating offers to some specific number of members of the opposite sex during an "adolescence period" (akin to the learning phase in Billari 2000). That is, for each potential partner an individual sees during adolescence, the individual judges whether the other's mate value is above his or her own aspiration level and, if so, makes an offer (which, however, cannot result in actual marriage during this initial period). Over this time, individuals can also adjust their aspiration level on the basis of whom they encounter and what happens during each encounter (e.g., offers or rejections).

After this adolescence period, the simulated men and women meet up in a further set of randomly assigned pairs, and they can either make a real proposal (an offer to mate) to their paired partner or decline to do so. If both individuals in a pair make an offer to each other, then this pair is deemed married and the two individuals are removed from the population. Otherwise, both individuals remain in the marriage pool to try again with someone else. This pairing-offering-marrying cycle is repeated until every individual is married or until every individual has had the opportunity to assess and propose to every member of the opposite sex exactly once.

With this simulation framework, we can test and compare different search and stopping mechanisms that the individuals use according to how many individuals in the population get married, how well matched the pairs end up being, and when the marriages occur. We can compare search rules along these dimensions not only against each other but also against sociological, demographic, and psychological data. On the first dimension, a worldwide effort to study (first) marriage patterns has shown that in most societies, 80% to 100% of adults marry (United Nations 1990). (The exception is the emerging pattern of nonmarital cohabitation in some countries, but for our purposes, it can be considered equivalent to marriage.) Second, as we indicated earlier, a large body of research in sociology and psychology has demonstrated the high degree of homogamy that is evident in marriage patterns, along such dimensions as ethnicity, religion, socioeconomic status, attractiveness, intelligence, and height (Coltrane and Collins 2001; Kalmijn 1998). This degree of homogamy has been quantified in some cases in a way that provides useful data for vetting our models, such as the high correlation, between .4 and .6, of the physical attractiveness of people in married couples (Kalick and Hamilton 1986). By taking attractiveness as a rough proxy for mate value (when other dimensions are held constant), we have a plausible numeric target (about .5) for the within-pair mate-value correlations that come out of our simulations. Finally, we have the age-at-first-marriage curves discussed earlier.

---

4. See Todd and Miller (1999) and Dudey and Todd (2002) for further results regarding other search heuristics, and contact the authors for the Lisp code used.

How do different search rules fare on these dimensions? First, trying to use a one-sided mate-search rule in the two-sided (and competitive) setting has disastrous results for most of the population. For instance, if everyone checks a dozen members of the opposite sex and sets an aspiration level equal to the highest mate value seen, then only 7% of the population will end up in mutually agreeing pairs (Todd and Miller 1999). Furthermore, only the highest-valued individuals end up mated with this rule (mostly in the top 10% of the population). This situation is certainly counter to human experience (as well as to that of other species in which mates select each other mutually, as in some monogamous animals), where the majority of individuals, across a wide range of relative mate values, are able to find mates. Clearly, a different kind of search rule must be used for mutual search.

An individual can achieve a much more successful two-sided mate search simply by using his or her own mate value (or slightly less) as the aspiration level for deciding which members of the opposite sex to propose to—assuming now that this mate value is known. With this approach, most of the population can succeed in finding and pairing up with mates of a similar value to their own (Miller and Todd 1998; Todd and Miller 1999). When we look at the hazard function for marriage, however, we see an unrealistic exponentially decreasing function (Figure 2a), similar to what appeared in Billari's (2000) one-sided search case. Thus, merely changing the search setting to two-sided choice does not, by itself, lead to a realistic distribution of marriage times. However, as Billari found, introducing variation in learning times (here, letting the adolescence period vary normally) proves to be a crucial factor that is sufficient to create the familiar unimodal hazard curve, as shown in Figure 2b.

But there is also a problem with this strategy: the accurate knowledge of one's own (relative) mate value that this strategy requires is not necessarily easy to come by. Individuals cannot be born with it because it is context sensitive (it depends on the others in one's social circle) and changes with age. Without this initial knowledge, then, people must somehow estimate their own mate value if they are to use it to form an aspiration level. What learning mechanisms could individuals use to arrive at aspirations that are in line with their own quality?

The one-sided learning rule presented earlier used only the information about the mate values of individuals who were encountered during the adolescent learning period, which does not reflect anything about the searcher's own mate value in our random-meeting environment. But there is more information available that can be used to infer one's own value: whether or not each encountered individual made a mating offer. In this case, the simulation model works as follows.

First, we set initial aspiration levels to an intermediate value of 50 for everyone, under a "no-knowledge" assumption (in these simulations, it does not make much difference if all individuals have the same initial aspiration level, whether 50 or otherwise, or if initial aspiration levels are randomly normally distributed). Next, each male A encounters a randomly chosen female B as the first step in each one's adolescent learning period. Male A can accurately see female B's mate value, and vice versa. If B's mate value is higher than A's aspiration level (currently 50), then A will make an offer to B; otherwise A will reject B. Similarly, B checks whether A's mate value is higher than her current aspiration level (also initially 50) and makes an offer to A or rejects him accordingly. A may then alter his aspiration level, for instance (depending on the learning rule he is using), adjusting it upward if he receives an offer from B and downward if he receives a rejection. B will do the same, depending on what she receives from A. So if A's mate value was 25 and B's mate value was 70, for example, then after this encounter, A may end up with a new aspiration level of 45 (following rejection by B), and B may end up with a new aspiration level of 55 (following an offer from A). Then A and B each go on to a second random encounter, using their new aspiration levels to determine whether to

Figure 2. **Hazard Functions for Marriage in a Population of Simulated Agents Who Are Searching for Mates Using Aspiration Levels Close to Their Own Mate Value**

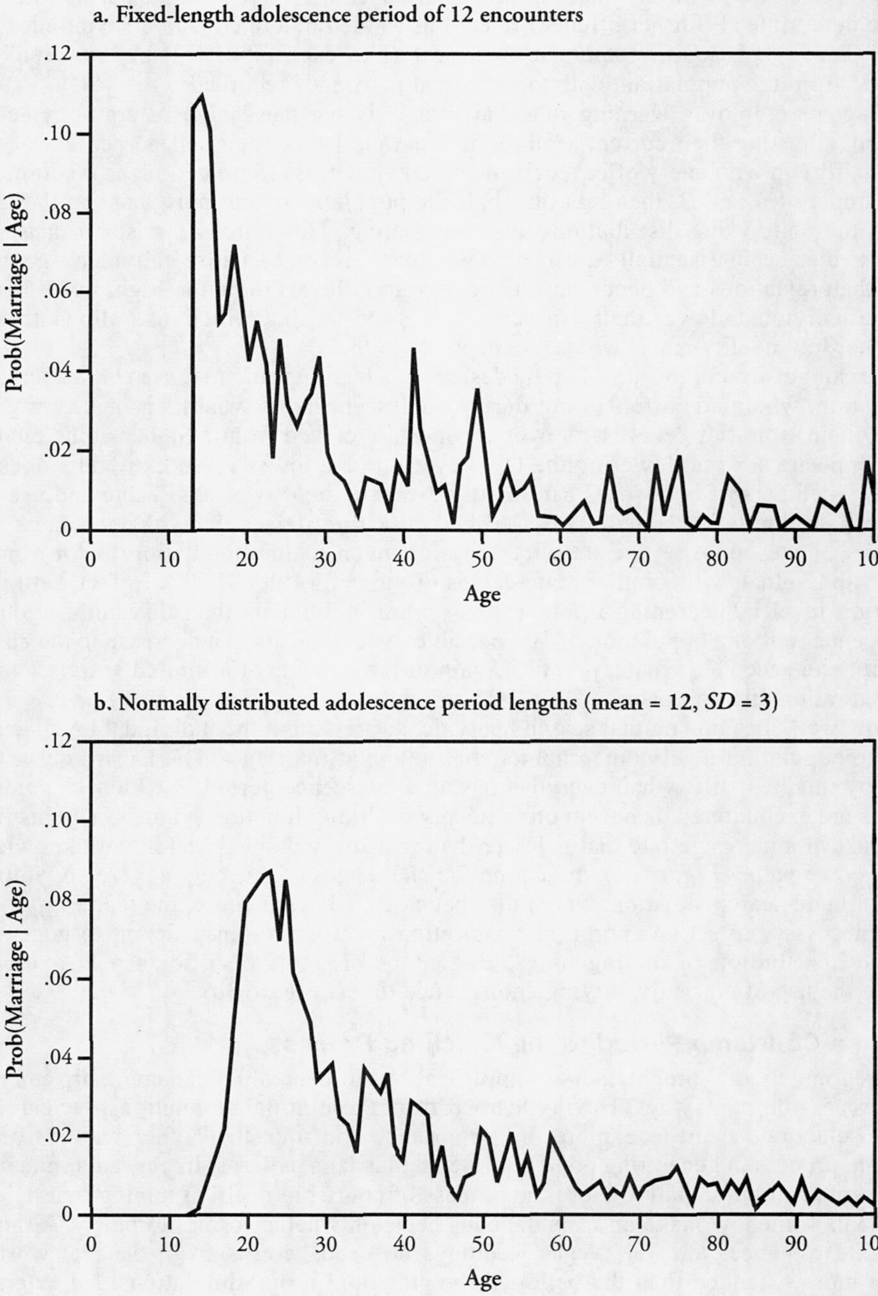

make an offer to or reject the next individual they encounter, and adjusting their aspiration levels further, depending on what they then receive in turn. This process of random encounters and aspiration-level adjustment continues through the adolescent learning phase for some predetermined number of encounters. Then the true mating phase begins, as described earlier, with the final aspiration levels from the adolescence period fixed and used to determine all further offers or rejections as the random encounters continue. During this mating phase, any pair who make offers to each other will be "married" and removed from the population, with their "age at marriage" recorded.

If agents employ a learning rule that uses only the data about offers or rejections received, adjusting their current aspiration level (and hence their self-perception of their own quality) up with every offer received and down with every rejection, as mentioned in the example of A and B, then less than half the population—and only those in the lower half of the mate-value distribution—ends up mating. This outcome arises because this learning rule acts in essentially a vain manner: above-average-quality individuals get more offers than rejections and hence raise their aspiration levels to be too high, while below-average individuals lower their aspiration levels too far, but which also allows them to find other low-quality mates who are acceptable to them.

We can get around this problem by designing a learning rule that uses both sources of information: who made offers or not during adolescence, and what their quality was. By raising their aspiration levels with every proposal received from a higher-value member of the opposite sex and lowering the level every time a lower-value individual does not propose, members of both sexes can rapidly estimate their own mate value and use it to pair up with similarly valued mates. With such a rule, fewer than 20 encounters with members of the opposite sex are necessary for much of the population to form mated pairs of individuals with similar mate values (Todd and Miller 1999).[5] In fact, setting an aspiration level by searching through many more individuals than this number during adolescence (out of a population of 100 possible mates) results in a decrease in the chance of finding an acceptable mate, pointing again to the benefits of a limited search within a bounded rationality approach.

How well does this mutual search heuristic, successful at the individual level, accord with the population-level demographic data on age at marriage? The hazard curve produced by this heuristic, when everyone has an adolescence period in which 12 potential partners are encountered, is once more a steeply declining function (Figure 3a). This function peaks at a marriage rate that is lower than that for two-sided search with knowledge of one's own value (Figure 2a), indicating the challenges of this competitive and initially ignorant mate-search situation. Given the shape of the hazard curve, the learning process that is necessary to set an appropriate aspiration level here is insufficient to generate a realistic distribution of mating times. But, as before, this insufficiency is overcome through the use of normally varying adolescence times (Figure 3b).

Adding a Courtship Period to the Matching Process

The foregoing models are, of course, simplifications of the real human courtship and marriage process in many ways. Having learned from these initial simulations, we can continue to elaborate them to explore the importance and impact of other features of the courtship process in generating population-level marriage patterns. In particular, the models presented so far actually ignore the process of courtship itself, assuming that pairs of individuals somehow make an instantaneous decision whether or not to marry. Although whirlwind romances and Las Vegas weddings do occur, courtship in the real world is usually more extended than the hello–yes/no situation in our simulations. An extensive

5. However, the number of mated individuals hovered around an unrealistically low 50%, a problem that is addressed by the mating models with courtship presented in the following section.

Figure 3. Hazard Functions for Marriage in a Population of Simulated Agents Who Are Searching for Mates Using Aspiration Levels Learned Based on Offers and Rejections Received During an Adolescence Period

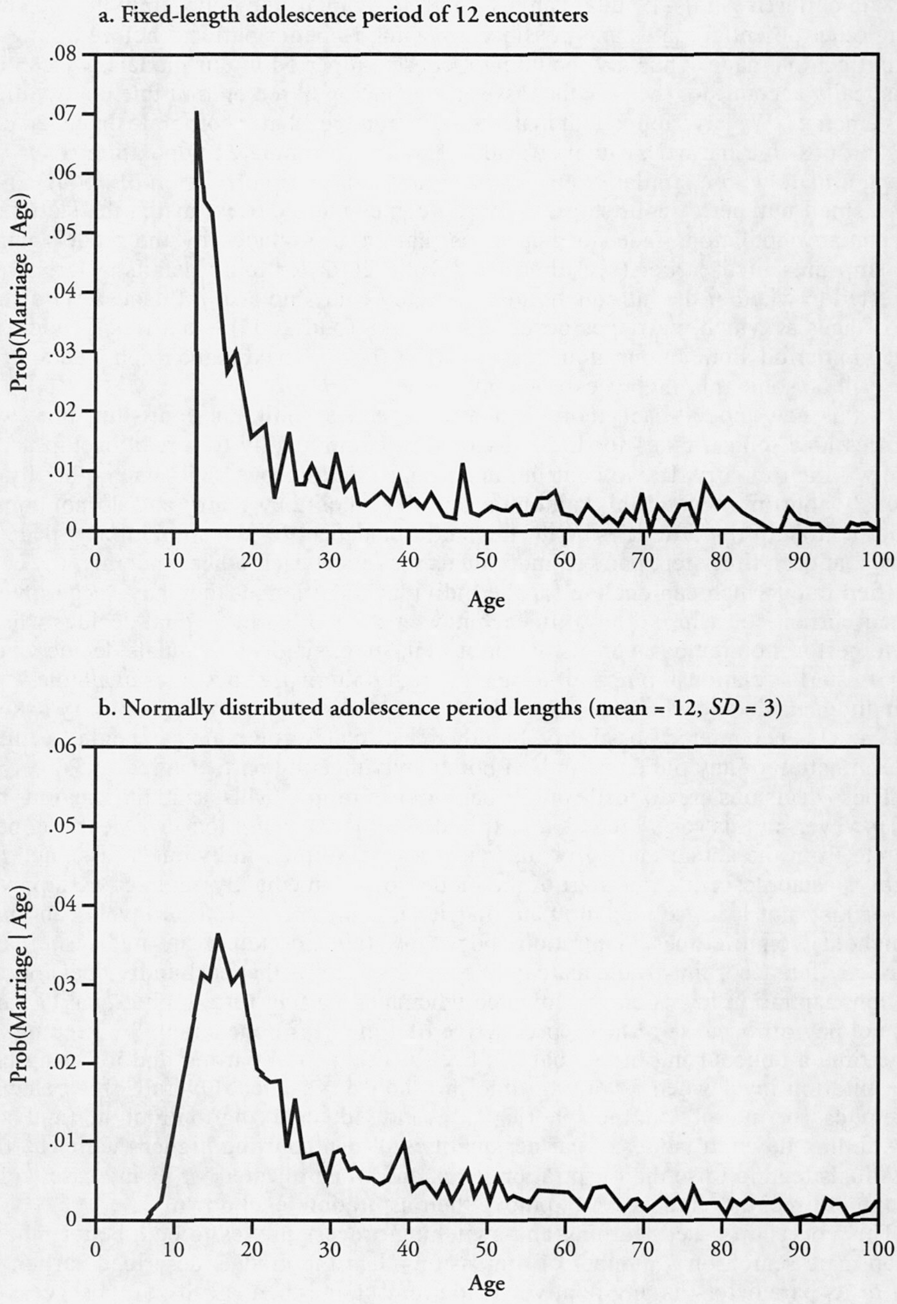

courtship period can serve a number of functions (Simão and Todd 2002): it can allow more information to be gathered about a potential partner, resulting in a better decision about his or her quality and potential match; it can enable an assessment of the potential partner's willingness to commit to a longer-term relationship (important in helping individuals to avoid the risk of abandonment, which can particularly affect women who are left with children to raise); and it can give both individuals the opportunity to keep monitoring other potential mates and possibly to switch to better partners before a long-term commitment is made. Thus, by including a courtship period in our models, we can more realistically account for the way that assortative mating based on multiple quality dimensions emerges. We have found that mate-search strategies that incorporate extended courtship and possible partner switching lead to most of a simulated population (over 95%) finding a mate with a similar overall quality (a within-pair correlation of about .5) after only a small number of courtships. These outcomes more closely match the statistics of real human populations than do the values that were produced by the models without courtship presented earlier (see Simão and Todd 2002 for more details). Here, we are interested in whether the introduction of a lengthy courtship also influences the distribution of ages at which marriages occur (Simão and Todd 2003)—that is, can adding the courtship period alone to our models account for the observed demographic data, or will there still be some missing necessary component?

In this new model, each individual has a specific minimum courtship time, which specifies how long it takes for the individual to commit fully to a relationship and become willing to marry. If two courting individuals "fall in love" with each other by continuously courting beyond this minimum duration, then they marry and do not consider further courtship opportunities. With this additional feature, the simulation proceeds as follows: at each time step, pairs of individuals may meet each other randomly at a certain specified rate (which can decline for the individuals the longer they have been involved in their current courtship, if any). In each new encounter, an individual decides what action to perform on the basis of his or her current state: single individuals decide whether to try to start a relationship or wait to see if a better alternative becomes available. Courting individuals decide whether to continue to court their current partner or to try to switch to the newly encountered possibility. In either case, a new courtship can begin (leading to the termination of any old ones) only if both newly meeting parties agree.

These decisions are all made on the basis of aspiration levels, as in the previous models. However, in this case, there is no explicit adolescent aspiration-level learning period separate from the actual mating or marrying period. Rather, individuals (assumed to be roughly postadolescent at the start of the model) can immediately begin courtship, which may or may not lead to their ultimate marriage, and their aspiration levels can change throughout their lifetimes. Aspirations begin low (all individuals are undiscriminating) and rise or fall according to the mate quality of the partners that each individual courts. In addition, aspiration levels can be lowered whenever waiting for a higher-quality partner does not pay off because of lost reproductive lifetime. It is done simply by keeping track of the time a noncourting individual has been waiting for a partner and lowering his or her aspiration level when a waiting-time threshold is reached. Overall, this behavioral strategy can be interpreted, metaphorically, as individuals trying to climb up (and sometimes falling down) a ladder of partner qualities. When courting higher-quality partners, individuals tend to raise their aspiration level, and when rejected, or in any case with the passage of time, individuals tend to move their aspiration level down.

This courtship-based learning-and-switching process proves to be a better model of human mate search on a number of dimensions than the models described earlier, with most of its parameters having relatively little impact and the majority of its success arising from the introduction of courtship and switching. In particular, this model results in almost all individuals quickly finding and consequently marrying mates of similar levels

Figure 4. Hazard Functions for Marriage in a Population of Simulated Agents Using Extended Courtship

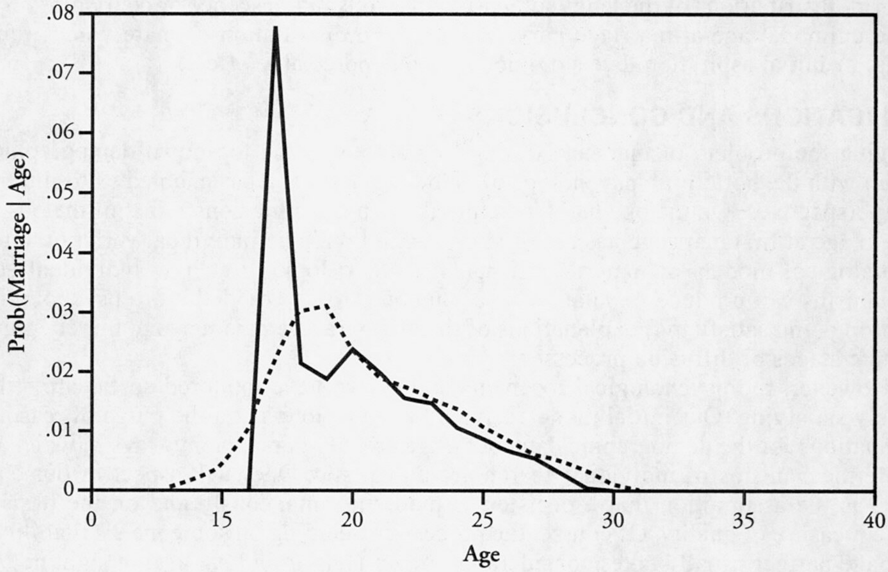

Notes: Solid line: individuals with a fixed five-year courtship duration before marriage. Dashed line: individuals with a normally distributed courtship duration before marriage (mean = 5 years, *SD* = 3).

of quality. But how quickly? In Figure 4 (the solid line), the distribution of ages at which individuals marry is an unrealistically tight spike, generated by most of the high-quality individuals marrying quickly, coupled with a long tail produced by the low-quality individuals marrying over a much longer period. Although this outcome produces a testable prediction regarding the relation between mate quality and age at marriage that accords qualitatively with some observed data (cf. Kalick and Hamilton 1986, whose model makes similar predictions), it differs from the expected age-at-first-marriage distribution much as did our first models. Thus, the courtship process alone does not appear to be sufficient to account for the population-level demographic patterns. Following our earlier finding that individual variation in learning time (adolescence) will lead to more-realistic marriage-age distributions, we can test for a similar outcome in this case. Here, we do not have a separate learning phase to alter, but the courtship period serves a related function, allowing individuals a period within which to appraise their own quality and, if feasible, to switch to a better partner. If we introduce variation into the minimum courtship time, making it normally distributed across individuals instead of fixed for everyone, the ages at marriage more closely follow the demographic patterns (Figure 4, the dashed line), again showing the importance of this simple manipulation of our models.

To summarize, what we have found so far in our explorations of mate-search mechanisms through a demographic lens is that various aspects of the individual search mechanism and task setting alone are insufficient to generate age-at-marriage distributions that reflect human patterns. Going from a noncompetitive one-sided search to a competitive two-sided (mutual) search did not create the expected skewed unimodal distribution or hazard function, nor did adding learning processes to the two-sided search, whether a nonmating adolescent trial period or an extended adult courtship period. Instead, we found that the introduction of variation across individuals in the population could lead to the appropriate patterns—but not just any type of variation. Only the inclusion of normal (or uniform) distributions of the length of learning periods (adolescence or courtship) resulted in the unimodal age-at-marriage curve; varying the distribution of mate values (quality levels) or initial aspiration levels did not have an appreciable effect.

IMPLICATIONS AND CONCLUSIONS

Studying the problem of marriage timing by combining the top-down demographic approach with the bottom-up psychological modeling approach has enabled us to illuminate both perspectives. On the one hand, meeting the demographic constraints of the observed data on age at first marriage has required us to build realistic individual variation into our psychological models of mate search. On the other, looking at how individual search mechanisms can produce population-level outcomes has provided more psychologically satisfying (and satisficing) explanations of the demographic data, beyond merely pointing to latent stages or diffusion processes.

However, the psychological mechanisms that we have explored so far are still not entirely satisfying. Our initial models have allowed us to expand the range of reasonable explanations for the demographic data, but there are, of course, many ways in which these simple mechanisms of individual search are unrealistic. One such aspect of our current model is the assumption that a decision is made instantaneously and on the basis of a single measure of quality. Of course, the process of deciding on someone's suitability as a marriage partner usually takes considerably longer (falling in love at first sight not withstanding) and involves many dimensions. By incorporating some form of this extended appraisal process into our model, we may be able to make more-accurate aggregate temporal predictions. For instance, potential partners could be assessed on a succession of cue dimensions, each one taking longer to evaluate than the previous one (e.g., status may take longer to assess than physical attractiveness, and personality may take longer still). Then the courtship process could be stopped at any point that one of these dimensions does not reach some desired level (Miller and Todd 1998), or it could be extended to gather more information if the uncertainty on a given dimension is too high (Oppenheimer 1988). Another criticism is that the use of strict cutoff aspiration levels in our model (accept any individual above a certain quality level and reject any below it) implies unrealistically deterministic behavior. Instead, a graded acceptance function (or adding noise into the quality-appraisal process) should result in more reasonable probabilistic behavior and more individuals finding partners more quickly.

The range of demographic data that we used to constrain and assess our psychological models should also be expanded. We need better data on the number of partners before marriage, not just on the age at marriage, but this information is much more difficult to obtain (both because the definition of "partner," as used in our models, is not specified, and because it is not clear how to determine when someone is engaged in the mate-search process at all; see Oppenheimer 1988). Individual search models make predictions about the relationships between age at marriage and the quality of those individuals who get married. Being able to test these predictions empirically would also be useful for distinguishing our models, but again there are the problems of defining and collecting demographic data about mate quality, for instance, whether to use income, as is common in some sociological studies; attractiveness, as psychologists have explored; or some combination of these and other dimensions. If these problems can be overcome, then we can also look at within-married-couple correlations of mate quality as a function of age at marriage and compare them to the predictions of our models.

Finally, the issue of cultural and historical differences must be addressed—how can we account for the factors that clearly make a difference in some aspect of marriage timing (e.g., delaying onset, if not changing the entire skewed-bell pattern), such as financial uncertainty or sex ratio? And can this approach be extended further, to cultures with different marriage traditions for which a sequential search model (at least at the individual level, as opposed to perhaps the parental or family level) may well not apply?

All these future research directions could be enhanced or enabled by the use of agent-based computer simulations that bring together top-down and bottom-up approaches to the same questions. At the same time, these simulation models bring together the disciplines themselves from which the disparate approaches arise. We have shown here how computer models can address questions from demography, economics, sociology, and psychology simultaneously. Furthermore, they can open up new directions for research: by creating new demands on the data or models of the other field, each side can also point out some of the interesting questions that remain in the other's efforts. In the end, these simulations of masses of interacting individuals remind us that the patterns we see at various social levels come down to the richly varying, yet at least partly rule-following, behavior of single (and married) people.

REFERENCES

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

Billari, F.C. 2000. "Searching for Mates Using 'Fast and Frugal' Heuristics: A Demographic Perspective." Pp. 53–65 in *Applications of Simulation to Social Sciences,* edited by G. Ballot and G. Weisbuch. Oxford, England: Hermes Science.

———. 2001. "A Log-Logistic Regression Model for a Transition Rate With a Starting Threshold." *Population Studies* 55:15–24.

Billari, F.C. and A. Prskawetz, eds. 2003. *Agent-Based Computational Demography.* Heidelberg, Germany: Physica Verlag.

Brüderl, J. and A. Diekmann. 1995. "The Log-Logistic Rate Model. Two Generalizations With an Application to Demographic Data." *Sociological Methods and Research* 24:158–86.

Burch, T.K. 1993. "Theory, Computers and the Parameterization of Demographic Behavior." Pp. 377–88 in *Proceedings of the International Population Conference, Montreal, Canada,* Vol. 3. Liege, Belgium: International Union for the Scientific Study of Population.

Burdett, K. and M.G. Coles. 1999. "Long-Term Partnership Formation: Marriage and Employment." *Economic Journal* 109:307–34.

Chase, V.M., R. Hertwig, and G. Gigerenzer. 1998. "Visions of Rationality." *Trends in Cognitive Science* 2:206–14.

Coale, A.J. 1971. "Age Patterns of Marriage." *Population Studies* 25:193–214.

Coale, A.J. and D.R. McNeil. 1972. "The Distribution by Age of the Frequency of First Marriage in a Female Cohort." *Journal of the American Statistical Association* 67:743–49.

Coale, A.J. and J. Trussell. 1996. "The Development and Use of Demographic Models." *Population Studies* 50:469–84.

Coltrane, S.L. and R. Collins. 2001. *Sociology of Marriage and the Family—Gender, Love, and Property,* 5th ed. Belmont, CA: Wadsworth.

Diekmann, A. 1989. "Diffusion and Survival Models for the Process of Entry Into Marriage." *Journal of Mathematical Sociology* 14:31–44.

Dudey, T. and P.M. Todd. 2002. "Making Good Decisions With Minimal Information: Simultaneous and Sequential Choice." *Journal of Bioeconomics* 3:195–215.

Epstein, J.M. and R. Axtell. 1996. *Growing Artificial Societies: Social Science From the Bottom Up.* Cambridge, MA: MIT Press.

Ferguson, T.S. 1989. "Who Solved the Secretary Problem?" *Statistical Science* 4:282–96.

Gigerenzer, G. and P.M. Todd. 1999. "Fast and Frugal Heuristics: The Adaptive Toolbox." Pp. 3–34 in *Simple Heuristics That Make Us Smart,* by G. Gigerenzer, P.M. Todd, and the ABC Research Group. New York: Oxford University Press.

Gigerenzer, G., P.M. Todd, and the ABC Research Group. 1999. *Simple Heuristics That Make Us Smart.* New York: Oxford University Press.

Gilbert, N. and K.G. Troitzsch. 1999. *Simulation for the Social Scientist.* Buckingham, England: Open University Press.

Goldstein, J.R. and C.T. Kenney. 2001. "Marriage Delayed or Marriage Forgone? New Cohort Forecasts of First Marriage for U.S. Women." *American Sociological Review* 66:506–19.

Henz, U. and J. Huinink. 1999. "Problems Concerning the Parametric Analysis of the Age at First Birth." *Mathematical Population Studies* 7:131–45.

Hernes, G. 1972. "The Process of Entry Into First Marriage." *American Sociological Review* 37:173–82.

Hey, J.D. 1982. "Search for Rules for Search." *Journal of Economic Behavior and Organization* 3:65–81.

Kalick, S.M. and T.E. Hamilton. 1986. "The Matching Hypothesis Reexamined." *Journal of Personality and Social Psychology* 51:673–82.

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

Keeley, M.C. 1979. "An Analysis of the Age Pattern of First Marriage." *International Economic Review* 20:527–44.

Lippmann, S.A. and J.J. McCall. 1976. "The Economics of Job Search: A Survey (Part I)." *Economic Inquiry* 14:155–89.

Lloyd, K.M. and S.J. South. 1996. "Contextual Influences on Young Men's Transition to First Marriage." *Social Forces* 74:1097–119.

Macy, M.W. and A. Flache. 1995. "Beyond Rationality in Models of Choice." *Annual Review of Sociology* 21:73–91.

Macy, M.W. and R. Willer. 2002. "From Factors to Actors: Computational Sociology and Agent-Based Modeling." *Annual Review of Sociology* 28:143–66.

Miller, G.F. and P.M. Todd. 1998. "Mate Choice Turns Cognitive." *Trends in Cognitive Sciences* 2:190–98.

Moon, P. and A. Martin. 1990. "Better Heuristics for Economic Search—Experimental and Simulation Evidence." *Journal of Behavioral Decision Making* 3:175–93.

Oppenheimer, V.K. 1988. "A Theory of Marriage Timing." *American Journal of Sociology* 94:563–91.

Seale, D.A. and A. Rapoport. 1997. "Sequential Decision Making With Relative Ranks: An Experimental Investigation of the 'Secretary Problem.'" *Organizational Behavior and Human Decision Processes* 69:221–36.

Simão, J. and P.M. Todd. 2002. "Modeling Mate Choice in Monogamous Mating Systems With Courtship." *Adaptive Behavior* 10:113–36.

———. 2003. "Emergent Patterns of Mate Choice in Human Populations." *Artificial Life* 9:403–17.

Simon, H.A. 1990. "Invariants of Human Behavior." *Annual Review of Psychology* 41:1–19.

Todd, P.M. 1997. "Searching for the Next Best Mate." Pp. 419–36 in *Simulating Social Phenomena,* edited by R. Conte, R. Hegselmann, and P. Terna. Berlin: Springer.

———. 2000. "Fast and Frugal Heuristics for Environmentally Bounded Minds." Pp. 51–70 in *Bounded Rationality: The Adaptive Toolbox,* edited by G. Gigerenzer and R. Selten. Dahlem Workshop Report. Cambridge, MA: MIT Press.

Todd, P.M. and G.F. Miller. 1999. "From Pride and Prejudice to Persuasion: Satisficing in Mate Search." Pp. 287–308 in *Simple Heuristics That Make Us Smart,* by G. Gigerenzer, P.M. Todd, and the ABC Research Group. New York: Oxford University Press.

United Nations. 1990. *Patterns of First Marriage: Timing and Prevalence.* New York: United Nations.