AGE-AT-MARRIAGE PATTERNS CAN EMERGE FROM INDIVIDUAL MATE-SEARCH HEURISTICS

PETER M. TODD
Center for Adaptive Behavior and Cognition
Max Planck Institute for Human Development
Lentzealle 94, 14195 Berlin, Germany
e-mail:ptodd@mpib-berlin.mpg.de

FRANCESCO C. BILLARI
Institute of Quantitative Methods, Università Commerciale Luigi Bocconi
and IGIER

JORGE SIMÃO
Computer Science Department, FCT / New University of Lisbon

Abstract

It is well known that the shape of the distribution of age at first marriage shows strong regularities across many countries and recent historical periods. We account for these patterns by developing agent-based models that simulate the aggregate behavior of individuals searching for marriage partners. Past models assumed fully rational agents with complete knowledge of the marriage market; our simulated agents use psychologically plausible simple heuristic 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.

INTRODUCTION

In modern Western societies, deciding when to get married seems like a highly personal and individual choice. We may feel that we are considering options and weighing possibilities that nobody else has ever had to think about in the same way. And yet much research has pointed out the societal and economic constraints that impact on 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 to date has done rather little to answer these questions, because of a strong division in focus between 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 typically obscures (or does not consider) 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 omits 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 in the middle to speak to each other. (Such multilevel understanding is also the aim of sociological rational choice theories, e.g. Hechter and Kanazawa 1997; however, many of these theories, particularly the “thin” ones, assume a psychologically implausible definition of rational behavior on the part of individuals, which we aim to avoid. As Petersen 1994, p. 499 says, “there is the concern that the analysis of a social system may go astray if one fails systematically in the description of the actors constituting the system.”)

One common language that we could attempt to get both perspectives to speak is that of mathematics. This has been done with some success (Billari 2000a; Coale and McNeil 1972; Diekmann 1989; Oppenheimer 1988), but with a certain degree of violence done to the assumptions at both the micro and macro levels. In particular, allowing for much variation in the strategies used by individuals quickly makes the mathematical models of their interactions intractable. As we will argue in this paper, it is exactly such individual-level variation that may underlie the emergence of the observed patterns at the population level. Thus, in addition to the difficulties of becoming fluent in the mathematical language, it may even be inappropriate 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 common tongue 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 enable us to capture and explore the impact of the vital variation that was missing from the mathematical models. We can do this by controlling and monitoring the micro-level decision mechanisms of each agent, and observing the patterns that emerge 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, it is clear that the study of demographic behavior can benefit significantly from this approach (Billari and Prskawetz 2003).

In the rest of this paper, we present our efforts to combine demographic and psychological approaches to marriage via agent-based modeling. Our aim is to explain the emergence of the commonly-found pattern of age at first marriage on the basis of individual decision-making behavior. While this pattern has already been explored from a few other perspectives, our explanation here is novel in that it must do more than just account for the demographic data—we also require the models to meet the additional constraints of psychological plausibility and fit 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 (Section 2). We then take the bottom-up approach and simulate the behavior of a cohort of *satisficing* agents looking for (marriage) partners in situations of both one-sided and mutual choice (Section 3). We find that plausible psychological mechanisms of choice suggested by the framework of *bounded rationality* need some refinements in order 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 (Section 4), 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 rather 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. While Coale and McNeil studied this common pattern in the frequency distribution of age at marriage, here we will 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. (The probability of being still unmarried at a certain age is also known as the survival function.)

To illustrate the shape of these hazard rates for marriage, we show in figure 1 the empirically observed functions (more specifically, the age-specific conditional probabilities of first marriage) for men and women in three populations of the late twentieth century: Romania, 1998, and Norway, 1978 and 1998. (These conditional probabilities are computed by dividing the number of first marriages of people who attain a given age *x* in a year by the number of still-unmarried individuals of age *x*-1 at the beginning of the year.) 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 rather different for Norway 1998, where non-marital cohabitation is widespread, it can still be described mathematically in a similar way. (If we included ages of transition to stable co-residential partnerships in this case, the curves would appear more similar.) 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.

**** FIGURE 1 ABOUT HERE *****

Hazard rates are preferable to frequency distributions for understanding behavior because they make clear the possibility that some individuals never show the behavior or event under study--here, marriage. For instance, the frequency distribution of age at death (after early adult years) is, like age at marriage, unimodal. Nevertheless, if one looks at hazard rates to death, they differ from the marriage hazard rates by being monotonically increasing with age (indicating the impact of the aging process), with the possible exception of extremely old ages (Vaupel et al. 1998). Clearly, while one can "survive" the marriage process by avoiding the transition to the married state, in the case of mortality there can be no long-term survivors. This important difference is clear in the hazard rate plots but obscured when one looks at the frequency distributions.

How can these common patterns in the ages at which people get married be explained? We can broadly classify existing demographic and sociological theories on the timing of marriage across two dimensions[1]. The first dimension can be considered as a continuum covering the ways in which people decide when to marry. At one extreme, the timing of marriage is a (more or less unthinking) response to social norms, while at the other extreme, marriages happen when rational individuals decide the time will maximize their utility function subject to constraints (Goldstein and Kenney 2001). Sociologists have traditionally emphasized the role of social norms, while economics has put the accent on rational choice (e.g., Becker 1981). Of course, there is a range of possibilities in between the two extremes, and some scholars have explicitly compared the weight of normative influence on the timing of marriage against the strength of rational choice Blossfeld and Huinink 1991).

The second dimension on which to classify theories concerns the heterogeneity of individuals, something more specific to modeling age patterns of marriage. On one end of the dimension, we can consider that all individuals belonging to the same cohort or population are subject to the same type of influence (e.g., they are all affected by a diffusion process as in Hernes 1972). On the other end, we can think of individuals as being heterogeneous in their choice of when to marry, according to some unobservable characteristics. This is the underlying hypothesis in split population compartment models like that of Coale and McNeil (1972; see also Bennet, Bloom, and Craig 1989; Bloom and Bennet 1990).

These two dimensions can help us to compare the existing models of age patterns of marriage. Three main types of formal behavioral models have been proposed (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 getting married, and that the length of time this process takes is governed by a stochastic process. More precisely, Coale and McNeil propose 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 very 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). This might be the consequence 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). For this reason, it is not clear how to classify the model along the normative versus rational choice continuum. In terms of population heterogeneity, the Coale-McNeil model assumes that the population is split between a group that is bound to get married and another group that is bound to remain single.

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 driving 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 they may become more eager to marry, they become less able to

---

[1] In what follows, we shall speak of marriage, but the reasoning and the data equally apply to other kinds of cohabiting union, including non-marital cohabitation.

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). In our two-dimensional classification, the Hernes model can be placed at the normative choice side, with an assumption of a homogeneous population.

Similar patterns fitting observed first marriage rates are produced by log-logistic diffusion models, where again the diffusion of behavior decreases with age. For instance, Brüderl and Diekmann (1995) proposed a log-logistic model with immunity, in which the diffusion process affects only a certain share of a cohort. Billari (2001) introduced a log-logistic model with a starting threshold, in which the diffusion process starts only after an individual reaches a certain age (which can vary for different sub-groups of a cohort). However, because this model does not allow for individuals who never marry, it is more useful for distinguishing the determinants (e.g., educational differences) of within-cohort differentials in the propensity to marry, than for representing the behavior of a cohort itself.

To compare how these latent-state and diffusion models apply to particular marriage patterns, we show in figure 2 their application to continuous-time data on age at first union (including marriage and non-marital cohabitation) from a ten-year birth cohort of Italian women. (The data and the procedures used to estimate the parameters of the models are described in detail in Billari 2000b.) We compare the fitted Coale-McNeil, Hernes, and log-logistic with immunity models against a parametric estimate of the observed marriage hazard rates (a three-year piecewise-constant hazard rate model). In all cases, it is possible to observe the characteristic unimodal shape of the curve, with a long tail to the right converging to zero. Although from a statistical point of view (according to the Bayesian Information Criterion--see Raftery 1995) the Hernes model performs better than the others with this particular data, it is striking how closely all three models fit the cohort patterns (at least from a qualitative perspective). The difficulty of distinguishing these macro-level models on the basis of their fit to the data indicates a possible advantage for individual-based models: They provide different explanations of the observed marriage patterns that can be supported with individual-level behavioral data as well.

**** FIGURE 2 ABOUT HERE *****

Finally, in the third type of model, 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 instantiated in these models can select possible mates, for instance by making and accepting offers, and the combined actions of these individuals over time yields the patterns of marriage age that we are interested in. Some models of this kind are based on the assumption of perfectly rational individuals performing optimal search; they also typically assume the agents are homogeneous in terms of their rational behavior. Keeley (1979), for example, adopts an optimal model from job search in which individuals set a threshold financial value for the minimal (monetary) benefit they seek in a marriage; if they find a partner with whom their income can exceed this threshold, then they marry. The cost-benefit analysis 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; Reich 2000): 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.

Countering the idea that marriage search is driven by complete knowledge, Oppenheimer (1988) developed a more plausible theory of how marriage timing is influenced by the uncertainty surrounding a potential partner’s future value. This uncertainty can arise particularly through lack of knowledge about the partner's (and one's own) future job prospects. Hence this theory predicts that highly uncertain job market conditions will lead to widespread delays in the timing of first marriages. But while such a theory can explain qualitative changes in the initial slope of the age-at-first-marriage curve, it does not account for the overall shape of the curve that we are concerned with here, which remains common across different economic and cultural situations. (The same holds true for Lloyd and South's 1996 application of this theory to situations with different ratios of men to women on the marriage market—fewer available women leads to later marriage for men, but the shape of the function across age is unspecified.) This is, at least in part, because Oppenheimer's theory did not specify the mechanisms of individual search behavior; at its core, it still assumed an unpsychological optimizing job search model, though one that had to take into account environmental constraints (here, employment uncertainty, and in Lloyd and South, varying sex ratios). Expecting that individuals can perform optimization with additional constraints remains a psychologically unrealistic view of rationality (Gigerenzer and Todd 1999). Instead, 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] Thus, in terms of the two theory classification dimensions, our models will not be fully rational in the traditional sense (nor norm-following, but rather somewhere in between), and as we will see, they will need to incorporate some measure of individual heterogeneity in order to produce the empirically observed age at marriage pattern.

MODELING SEQUENTIAL SEARCH PROCESSES

To construct an agent-based model to account for population-level demographic phenomena relating to age at first marriage, we can create a set of simulated individuals that 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: 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 looking 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?

---

[2] It is important to keep clear that, while individual humans commonly use search and decision mechanisms with limited information-processing requirements rather than complex optimizing processes, those mechanisms themselves are likely to have arisen through a process more akin to optimization with constraints, namely biological or cultural evolution (Macy and Flache 1995), which can lead them to be well-fit to the situations in which they are used.

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 of the potential partners are simultaneously available to an individual 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 interested in how stable matchings can be made between men and women who have complete knowledge of all available partners (Bergstrom and Real 2000). While this may apply to some very small societies, it does not seem to match most of the cases of large populations where demographers have collected age-at-marriage data. Instead, people seeking mates (or other things, such as houses, jobs, even consumer products) often must choose between a set of options that we do not see 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., people we have dated and broken up with in the past are probably not still interested in rekindling the relationship again later). The problem then becomes one of deciding when to stop searching and go with the currently-available option.

This problem of how to decide when to stop searching applies to whatever set of potential mates an individual may encounter (assuming the set is larger than just one or two possibilities). That is, even after the set of potential partners has been constrained by all the social factors commonly studied by sociologists, such as age, ethnicity, employment (current and future possibilities), intelligence, religion, status, etc. (Kalick and Hamilton 1986; Kalmijn 1998; Lloyd and South 1996; Oppenheimer 1988), there will still be a (presumably fairly large) number of people that one could encounter and appraise as a marriage partner. Given that this remaining constrained set is encountered over time and varies along some dimensions of interest to the individual seeking marriage, a sequential search must still take place and hence a process for deciding when to stop the search must be used (but cf. Reich 2000, for an opposing view).

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 search, given the tradeoff between time and other costs that accumulate with each alternative seen, and the chance that the next alternative checked will be better than those encountered previously. This could involve extensive calculations, such as Bayesian updating of probability estimates or assessments of the costs of foregone opportunities, and thus require considerable time and computational resources.

But to make choices in a useful amount of time, real agents must employ limited search for options to choose between, 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 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 Herbert Simon 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 (Simon 1990). 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, we should construct models of the actual bounded psychological processes that guide our decision making.

Leaving behind unbounded rationality for a more plausible view of human behavior does not imply, however, entering the realm of irrationality and error. The simple psychologically realistic decision mechanisms that people use can also perform very well, provided that they are used in the proper settings—that is, provided that they are employed in environments that are structured in ways they can exploit. This combination of bounded mental mechanisms operating in environments they are attuned to yields *ecological rationality* (Gigerenzer, Todd, and the ABC Group 1999), the idea that the limited human mind can exploit the rich structure available in the environment to reach good decisions without extensive time, information, or computation. For instance, simple mechanisms for sequential search can perform nearly as well as the best approaches known, provided the search environment is properly structured (e.g., stable over time—see Dudey and Todd 2002). The study of ecological rationality involves the exploration of such decision mechanisms and the nature of their fit to appropriate environments (both rules for sequential search and "fast and frugal" heuristics for other forms of inference and choice—Gigerenzer et al. 1999). As we will show, demographic data can provide another source of evidence, beyond individual-level data, for discovering what decision heuristics people may be using in different contexts.

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 our 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 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 has proposed a *satisficing* approach to search, in which individuals check successive alternatives until one is found that is good enough (rather than optimal) for their goals (Simon 1990). This approach can be implemented by means of an aspiration level that individuals set somehow and then use in further search, stopping that search as soon as an option is encountered that exceeds the aspiration level. Here we will assume that all potential marriage partners can be assessed on some unidimensional quality scale, so that searchers can set a quality (or mate value) aspiration level for stopping search on a suitable partner. The exact way in which the aspiration level is set depends on further details about the search situation encountered. We will now consider two specific situations: one-sided search, in which a searcher considers a sequence of potential partners who have no say in the decision, and two-sided or mutual search, in which two populations of searchers (males and females) are assessing each other simultaneously and must both agree to any marriage. (Of course, cultural norms and individual emotions exert very strong influences on mate search and mate choice. While 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-Sided Search Processes

One way to conceive of the search for a marriage partner is as a shopping expedition in an open-air market, where the searcher wanders past a series of, say, mango stands, checking the wares of each until a suitably ripe mango is found, which the searcher then procures and takes home. From the searcher's perspective then, the available options are encountered one by one, the distribution of quality is unknown, previously passed-by options are no longer available (assume that someone else at the market buys the previously seen mangos), and a final choice can be made in a unilateral fashion--no mango ever argues that the purchaser is not good enough for it. Hence search here can be characterized as one-sided (and mostly non-competitive). Clearly this is a somewhat unrealistic simplification of the mate search process for most (if not all) cultures, but it has proven a useful starting point for our modeling efforts, and we can test how much the simplifications affect the results we are interested in. Now we must ask, how can searchers operate effectively in this situation? More specifically, what mechanisms are appropriate for setting aspiration levels to guide search?

We have help in answering this question because the one-sided mate search problem just laid out is close to a widely studied problem in probability theory known as the *dowry problem*[3]. In the dowry problem it is the searcher's aim to find the one woman out of a set of *N* with the highest dowry (e.g., money for marriage; alternatively, one can think of a woman searching for the one man in some set with the highest income). . The women are assessed by the searcher sequentially in a random order and the searcher has no knowledge about the distribution of dowry values. With each new woman seen, the searcher learns her dowry (or the current rank of her dowry, in the strictest version of the problem), and then must choose between stopping the search and thus marrying the current woman, or continuing the search to look for a higher dowry. If the searcher continues he cannot go back and choose an earlier woman -- that is, there is no ability to “recall” past alternatives in this search.

To maximize the chance of selecting the highest dowry, the searcher should look through the first 37% of the women, set an aspiration level equal to the highest dowry seen in that 37% sample, and then select the first woman seen thereafter who has a dowry above the aspiration level. (See Ferguson 1989, for a review of the literature on this problem and its optimal solution.) But this method of setting the aspiration level requires searching through 74% of the available options on average before a choice is made, and results in only a 37% success rate for the goal of picking the single highest dowry. Such an optimizing goal along with the long search and low success it leads to do not match the typical human mate search process. Nor in fact do people in experimental versions of the dowry problem search as long as the optimal 37% rule dictates (Kahan, Rapoport, and Jones 1967; Seale and Rapoport 1997).

Instead, Todd (1997; see also Todd and Miller 1999) proposed that more realistic goals for mate search would include maximizing the expected mean mate value selected (or mean dowry received in the dowry problem) or the probability of finding a mate in the top quarter of the population quality distribution. To perform well given such goals, and without knowing the exact number of available alternatives ahead of time, a searcher only needs to check a much smaller set of potential partners before setting an aspiration level – Todd and Miller (1999) found via simulation that assessing on the order of 12 partners works well (when the total number of available partners is in the hundreds or thousands). That is, the searcher can make good choices by checking the (unidimensional) quality levels of the first 12 potential partners seen, remembering the highest quality among those 12 and setting an aspiration level at that value, and then from the 13th potential partner on, stopping search on the first person seen whose quality exceeds the aspiration level. Such a rule

---

[3] It is also known in the statistical literature as the *secretary problem* (Ferguson 1989; Seale and Rapoport 1997). One difference between this problem and the mate search scenario laid out earlier is that the number of available options here is known beforehand. Note also that the solutions to this problem can apply to any set of options, including one that has been restricted beforehand to a limited range of values; for instance, even if one’s mate search has been restricted through one’s upbringing to a particular social strata, these search mechanisms can still be applied to finding an individual with a high relative value or rank within that restricted set.

results in finding a good partner (according to those goals) in over 90% of the times it is applied, after searching through about 30 potential partners on average.

However, empirical support for the use of such quick aspiration-setting rules in one-sided search situations remains to be gathered (along with determining whether or not mate search conforms to such situations in the first place). Seale and Rapoport (1997) have found experimental evidence that people do search in settings akin to the dowry problem by setting aspiration levels after a short initial search. But Dudey and Todd (2002) have observed that people change their search behavior appreciably when directed to search with different goals such as selecting a high value or selecting one in the top 10% of a distribution. This indicates that no single heuristic is likely to be employed across the different variations of the one-sided search task, and hence we must gather more evidence to determine just what psychological mechanism is being used in a particular setting.

The need for further evidence to constrain our individual-level psychological models is exactly where demographic research comes in. Billari (2000a) recognized that the simple search rules proposed by Todd (1997) would have testable population-level implications when used by a group of individuals all looking for mates--implications of the sort that demographic data could empirically assess. In particular, Billari wanted to model a group of simulated individuals using a simple aspiration-level search rule to see whether their distribution of ages at first marriage would match the nuptiality patterns described in section 2. If not, then what changes would be necessary to make the group-level behavior conform to the observed regularities?

Billari (2000a) simulated agents searching individually through a population of 100 potential mates with varying (uniformly distributed) qualities. Each agent first checked 12 potential mates and set an aspiration level equal to the highest quality seen, then looked through the remaining 88 possibilities, stopping on (selecting) the first individual encountered with a quality exceeding the aspiration level. Billari plotted the distribution of the "age" at which agents chose a mate, where each individual seen corresponds to a "year" of life.[4] But rather than producing the right-skewed unimodal pattern seen in human societies, Billari found (and proved mathematically) that this simple rule creates a distribution that starts instantaneously (at age 13 in this case) and falls rapidly thereafter--see Figure 3a, where we have rerun his simulation and plotted the results as a hazard function for the probability of entering marriage at a given age. Thus, the simple aspiration-level search rule, at least when used by a group of homogeneous non-interacting searchers, does not produce the expected population-level behavior.

**** FIGURES 3a, 3b ABOUT HERE *****

---

[4] 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 itself 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.

What is the minimal alteration to this model that could make it match the human marriage data? Billari struck upon introducing heterogeneity to the searching agents, simply by giving them a range of "learning times" to use to set their initial aspiration level. That is, rather than every agent first checking (say) 12 potential mates, each agent has its own specific number of individuals to check before setting its aspiration level. Billari investigated the effect of various distributions of these learning times, and found that when these distributions were strongly age-skewed--that is, when more agents used longer learning times--an age-at-first-marriage curve approximating the traditional unimodal shape was produced. But how realistic is such a distribution? This is very difficult to assess (we can find out how many partners someone has before marriage, but how can we know the number of partners they had before setting an aspiration level for marriage?), but a normal or uniform distribution of lengths of learning times would seem a more natural default assumption. And in fact, when we give the population of agents a normal distribution of lengths of learning periods (with mean 12 and standard deviation 3), a right-skewed roughly unimodal curve for the marriage hazard rate does again emerge (Figure 3b; the same effect occurs for uniform distributions, but is not shown). We have also investigated the effects of introducing or modifying other forms of variation in the population, for instance using a normal distribution of mate values rather than a uniform distribution, but have not found any that produce the same skewed-unimodal distribution of marriage ages. (Other forms of individual variation, such as in risk attitudes, could also be explored given the proper instantiation in the model.)

Thus, a population of agents following a simple rule for searching for suitable mates--checking a few and setting an aspiration level to guide further search--can not only exhibit bounded (or ecological) rationality by making good choices with a small amount of effort and information, but can also be characterized by a distribution of ages at mate choice (first marriage) that approximates empirically observed human values. This is provided that the population is suitably heterogeneous, for instance in the amount of time that each agent takes to learn about the available set of mates and set an aspiration level. In this way, the real-world demographic age-at-first-marriage data has helped to inform and constrain our individual-based psychological modeling. However, this simplified model is of course just a first step, lacking some of the important features of many human mating decisions. In particular, the search process we have presented so far is just one-sided, with the searcher making the final decision without any input from those being searched. For many human cultures (and many other species), mate choice is instead a (more) mutual affair, with both sexes taking part in the decision process. We amend our models to address this issue in the next section.

Two-Sided (Mutual) Search Processes

One-sided mate search, in which the members of one sex do the searching and make the decisions, is clearly not a realistic model of the adaptive problem of mate choice actually faced by people in many cultures (including modern Western ones). The problem is that 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 male does not meet the standards of a particular female he is interested in, for instance, then his courtship attempts are doomed to failure. Furthermore, in contrast to the model presented in the previous section, searching individuals do interact, at a minimum because they are vying for the same set of potential mating partners. The way this two-sided matching process can be seen in a microcosm in the "pairing game" classroom demonstration (Ellis and Kelley 1999). In this game, individuals in a group each wear a card on their forehead with a number (e.g. from 1 to 30) that they do not know but which others can see; everyone is then told to pair up with the highest-numbered person they can (without speaking). After a flurry of evaluation of others and extending of hands in attempted offers, individuals with high numbers soon learn that they are in demand and select another person with a high number, sealing the pairing with a handshake. After the high values leave the population, individuals with lower numbers have to then settle for the best similarly lower number they can get.

What mechanisms might underlie this matching process? In this competitive setting, the kind of one-sided search rules explored above perform poorly: If everyone in the population has been setting their aspiration levels based on the highest mate value seen in the first dozen potential mates they encounter (or first few numbers they see in the pairing game), then everyone will end up with a rather high aspiration for whom they will agree to mate with. Given these high aspiration levels, the trouble is that only those rare individuals with high mate values (relative to the rest of the population) will be chosen as mates, and most individuals will end up alone.

Instead, as can be seen in both the pairing game and real life, most individuals do not end up alone—which indicates that individuals seem to be able to adjust their courtship targets on the basis of experience. Either they could be aiming for the best available, and lowering their aim as the quality of the individuals at the top falls (because the best pair up early and leave the mating or pairing game), or they could be using their own value as a target to match, which implies that they must be learning about their own value somehow, even though they begin without knowing it. Kalick and Hamilton (1986), in an early discipline-bridging individual-based simulation aimed at bringing insights from the field of biology into the field of psychology, modeled these two types of mutual mate search strategies to find out whether both would end up producing empirically-observed patterns of assortative mating or homogamy (similar individuals pairing up, e.g. like marrying like). They started with the fact that observations of many human populations show that individuals in couples are highly correlated in attractiveness (correlations between 0.4 and 0.6 in different studies). This finding had led social scientists in the 1960's to propose the “matching hypothesis” that people actively seek a mate matched to them in attractiveness. But this seems to contradict experimental data indicating that people in general tend to prefer more physically attractive individuals as prospective partners (i.e., not taking their own attractiveness into account). To explain this apparent contradiction, Kalick and Hamilton set up individual-based simulations to study the relationship between individual-level preferences and population-level patterns. In their simulations, randomly selected individuals with particular attractiveness values are paired up sequentially in “dates.” Both individuals in a date then use a probabilistic acceptance criterion to decide whether or not they accept each other, and, if both agree, they mate and leave the population. A discounting factor was introduced to make individuals less choosy with time.

Kalick and Hamilton's results demonstrated that universal preferences for high attractiveness, as was already obvious for preferences for similar attractiveness (matching), can produce realistic degrees of intra-couple correlation of attractiveness (.55). This is because, just as is seen in the pairing game demonstration described above, higher attractiveness individuals tend to pair (and leave the mating pool) earlier than lower quality individuals, leaving the lower quality individuals with no option but to lower their aspirations (if not their preferences) and pair amongst themselves. However, Kalick and Hamilton's support for a mate search mechanism based on preferences for high attractiveness is open to criticism. First, an unrealistically high number of dates (evaluations of members of the opposite sex) was required in the model for a realistic intra-couple attractiveness correlation to be obtained and a significant percentage of the population to mate (e.g., it took 40 “dates” for the correlation to reach .43 and for 86% of the individuals to mate—see Aron 1988). Second, Kalick and Hamilton let in a similar-attractiveness-based preference adjustment through the back door: By including a choosiness-discounting factor that increased with time (which was necessary to keep the simulation from taking even longer), this adjustment applied more to individuals who were left longer in the simulation—that is, those who were less attractive. This indicates that high-attractiveness-based mechanisms that rely on a (possibly slow) process of marriage-market clearing from the top down may be less realistic than mechanisms seeking similarly attractive mates. But seeking mates similar in attractiveness assumes that one knows one’s own attractiveness in the first place, which as we will see can be a problematic assumption.

To explore how different search rules can work in a two-sided setting, we must amend our earlier model (see Todd and Miller 1999, and Dudey and Todd 2002, for more details). This time, similar to Kalick and Hamilton’s models, we start with a population containing two sets of searchers, 100 males and 100 females, 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”. 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 up or down from an initial value. Here we set initial aspiration levels to an intermediate value of 50 for everyone, under a “no-knowledge” assumption (it turns out not to make much difference in these simulations if all individuals have the same initial aspiration level, whether 50 or otherwise, or if initial aspiration levels are randomly normally distributed).

After this adolescence period, males and females meet up in random 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.

With this simulation framework, we can test and compare different search and stopping mechanisms used by the individuals. In comparing the effectiveness of different search rules in this setting, we are interested more in population-level measures across all agents than in just the individual success that was natural to consider in the one-sided search case. Now we want to see how many individuals in the population get married, how well matched the pairs end up being, and when the marriages occur[5]. 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 between 80% and 100% of adults marry (United Nations 1990). (The exception is the emerging pattern of non-marital cohabitation in some countries, but for our purposes this can be considered equivalent to marriage.) Second, as indicated earlier, a large body of research in sociology and psychology has demonstrated the high degree of homogamy evident in marriage patterns, along dimensions including ethnicity, religion, socioeconomic status, attractiveness, intelligence, height, and the like (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 physical attractiveness of people in married couples (Kalick and Hamilton 1986). By taking attractiveness here as a rough proxy for mate value (when other dimensions are held constant), we

---

[5] These last two measures, the goodness and timing of matches, are combined in Oppenheimer’s (1988) notion of search efficiency, which she defines as “the likelihood of finding a good match for each unit cost of investment—that is, the direct and indirect (e.g., time) costs invested per period of search” (p. 569), but her definition leaves unspecified exactly how to quantify the concept.

have a plausible numeric target (around .5) for the within-pair mate value correlations that come out of our simulations. And third, we have the age-at-first-marriage curves discussed earlier.

How do different search rules fare on these dimensions? As we indicated above, trying to use a one-sided mate search rule in the two-sided (and competitive) setting has rather 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 very highest-valued individuals end up mated with this rule (mostly in the top 10% of the population). This is certainly counter to human experience, as well as that of other mutually-selecting (e.g., some monogamous) species, where the majority of individuals, across a wide range of relative mate values, are able to find mates. (As we saw in the previous section, this search rule also generates an unrealistic distribution of marriage ages unless it is modified to include variable learning periods.) Clearly, a different kind of search rule must be used for mutual search.

Much more successful two-sided mate search can be achieved by an individual simply 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 again see the unrealistic exponentially-decreasing function (Figure 4a) that appeared in the 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. As before, though, introducing variation in learning times (here, letting the adolescence period vary normally) does suffice to create the familiar unimodal curve, as shown in Figure 4b.

**** FIGURES 4a, 4b ABOUT HERE *****

But there is also a problem facing this strategy: The accurate knowledge of one's own mate value that it requires is not necessarily an easy thing to come by. Individuals cannot be born with it, because it is both context-sensitive (it depends on the others around) and changes over age with development. 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 in line with their own quality?

The one-sided learning rule presented above only used the information about the mate values of individuals encountered during the adolescent learning period. But there is more information available: whether or not each encountered individual made a mating offer. If agents use just this latter data in a learning rule, adjusting their aspiration level (and hence self-perception of their own quality) up with every offer received and down with every rejection, less than half of the population ends up mating--and only those in the lower half of the mate value distribution. This is because this learning rule is too vain: Above-average-quality individuals get more offers than rejections and hence raise their aspiration levels to be too high, while below-average individuals conversely lower their aspiration levels too far, but which also allows them to find other low-quality mates that 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 one's aspiration level 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 that to pair up with similarly-valued mates. With such a rule, less than twenty encounters with members of the opposite sex are necessary for much of the population to form mated pairs of individuals with similar mate value (Todd and Miller 1999)[6]. In fact, setting an aspiration level by searching through many more than this during adolescence (out of a population of 100 possible mates) results in a decrease in chance of finding an acceptable mate, pointing again to the benefits of 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 adolescent period in which 12 potential partners are encountered, is once again a steeply declining function (Figure 5a). This function peaks at a marriage rate below that for one-sided non-competitive search (Fig. 3a) and two-sided search given knowledge of one's own value (Fig. 4a), indicating the challenges of this competitive and initially-ignorant mate search situation. Given the shape of the hazard curve, we see that the learning process necessary to set an appropriate aspiration level here is insufficient to generate a realistic distribution of mating times. But as before, this is overcome through the use of normally-varying adolescence times (Figure 5b).

**** FIGURES 5a, 5b ABOUT HERE *****

Adding a Courtship Period to the Matching Process

The above 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. While whirlwind romances and Las Vegas weddings do occur, courtship in the real world is usually rather more extended than the hello-yes/no situation in our simulations. An extensive courtship period can serve a number of functions (Simão and Todd 2002): It can allow more information to be gathered about a potential partner, allowing a better decision to be made 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 (particularly important in helping women to avoid the risks of abandonment); and it can give both individuals the opportunity to keep monitoring other potential mates, and possibly 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 incorporating extended courtship and possible partner switching lead to most of a simulated population (over 95%) finding a mate with a similar overall quality (within-pair correlation around .5) after only a small number of courtships. These outcomes more closely match the statistics of real human populations than do those values produced by the models presented earlier without courtship (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 do we still need to include some further modification?

---

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

In this new model, each individual has a specific minimum courtship time, which specifies how long it takes for the individual to fully commit to a relationship and become willing to marry. (This courtship time can fall with age, reflecting a faster decision process later in life.) 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. Individuals encounter one another stochastically at a rate which can decline with increasing involvement in the current courtship process. With each encounter, individuals decide what action to perform based on 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 try to switch to the new possibility. In either case, a new courtship is begun (and any old ones are terminated) if both 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 post-adolescent 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 lifetime. Aspirations begin very 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. This is done simply by keeping track of the time a non-courting individual has been waiting for a partner and lowering his or her aspiration level when a waiting time threshold is reached. Overall, this behavior 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 will tend to raise their aspiration level, and when rejected, or in any case with the passage of time, individuals will tend to move their aspiration level down.

This courtship-based learning-and-switching process results in almost all individuals quickly finding and consequently marrying mates of similar levels of quality. But how quickly? In Figure 6 (solid line), we see that the distribution of ages at which individuals marry is an unrealistically tight spike, generated by most of the high-quality individuals marrying very quickly, coupled with a long tail produced by the low-quality individuals marrying over a much greater period of time. While 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 sufficient to account for the population-level demographic patterns. Following our finding earlier 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 of time within which to appraise their own quality and switch to a better partner if feasible. If we introduce variation into the minimum courtship time, making it normally distributed across individuals instead of fixed for everyone, we find that the ages at marriage do now more closely follow the demographic patterns (Figure 6, dashed line), again showing the importance of this simple manipulation of our models.

**** FIGURE 6 ABOUT HERE *****

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 are alone insufficient to generate age-at-marriage distributions reflecting human patterns. Going from non-competitive one-sided search to 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 non-mating 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 we have explored so far are still not entirely satisfying. Our initial models have allowed us to expand the range of reasonable explanations to consider for the demographic data, but there are of course many ways in which these simple mechanisms of individual search are unrealistic. While this simplicity is a necessary first step in modeling, “introducing complications, and incorporating results from other lines of inquiry, only as simpler, more parochial specifications are well-understood” (Macy and Flache 1995:89), more psychologically plausible models should enhance our understanding of the processes involved.

One unrealistic aspect of our current model is the assumption that a decision is made instantaneously 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. One possibility could be that potential partners are assessed on a succession of cue dimensions, each one taking longer to evaluate than the previous (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, reject any below that) 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. Furthermore, whatever aspiration levels that individual searchers use, whether hard or soft, could be modifiable not only through interactions with potential partners, but also through social comparison with one’s competitors (Mussweiler, 2003).

The range of demographic data 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 is much more difficult to come by (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—Oppenheimer 1988). Individual search models make predictions about the relationships between age at marriage and the quality of those individuals getting married. Being able to test these predictions empirically would also be very 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, or 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 this against 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 (such as delaying onset, if not changing the entire skewed-bell pattern), such as financial uncertainty or sex ratio, in our individual-based models? And can this approach be extended further, to cultures with different marriage traditions where a sequential search model (at least at the individual level, as opposed to perhaps the parental or family level) may well not apply?

All of these future research directions will be enhanced, if not 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, sociology, and psychology simultaneously. Further, 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 remaining in the efforts of the other. In the end, these simulations of masses of interacting individuals serve to 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

Aron, A. 1988. “The Matching Hypothesis Reconsidered Again—Comment on Kalick and Hamilton.” *Journal of Personality and Social Psychology*, 54(3), 441–46.

Becker, G.S. 1981. *A Treatise on the Family*. Cambridge, MA: Harvard University Press.

Bennet, N.G., D.E Bloom, and P.H. Craig .1989. “The Divergence of Black and White Marriage Patterns.” *American Journal of Sociology* 95:692–722.

Bergstrom, C.T. and L.A. Real. 2000. “Towards a Theory of Mutual Mate Choice: Lessons From Two-Sided Matching. *Evolutionary Ecology Research* 2(4):493–508.

Billari, F.C. 2000a. “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: Hermes Science Publications.

———. 2000b. L'analisi delle biografie e la transizione allo stato adulto. Aspetti metodologici e applicazioni ai dati della Seconda Indagine sulla Fecondità in Italia. Padova, Italy: Cleup editrice.

———. 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: Physica Verlag.

Bloom, D.E. and N.G. Bennet. 1990. “Modeling American Marriage Patterns.” *Journal of the American Statistical Association* 85:1009–17.

Blossfeld, H.-P. and J. Huinink. 1991. “Human Capital Investments or Norms of Role Transition? How Women’s Schooling and Career Affect the Process of Family Formation.” *American Journal of Sociology* 97:143–68.

Brüderl, J. and A. Diekmann. 1995. “The Log-Logistic Rate Model. Two Generalizations With an Application to Demographic Data.” *Sociological Methods & 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: International Union for the Scientific Study of Population.

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

Chase, V.M., R. Hertwig and G. Gigerenzer. 1998. “Visions of Rationality.” *Trends in Cognitive Science* 2(6):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.

Ellis, B.J. and H.H. Kelley. 1999. “The Pairing Game: A Classroom Demonstration of the Matching Phenomenon. *Teaching of Psychology* 26(2):118–21.

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(3):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: 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(2):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 & Organization* 3:65–81.

Kahan, J.A., A. Rapoport and L.V. Jones. 1967. “Decision Making in a Sequential Search Task.” *Perception & Psychophysics* 2:374–76.

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* 74:1097–1119.

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.

Mussweiler, T. (2003). “Comparison Processes in Social Judgment: Mechanisms and Consequences.” *Psychological Review* 110:xxx-xxx.

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

Petersen, T. 1994. “On the Promise of Game Theory in Sociology.” *Contemporary Sociology* 23:498–502.

Raftery, A.E. 1995. “Bayesian Model Selection in Social Research.” Pp. 111-96 in *Sociological Methodology 1995*, edited by P.V. Marsden. Oxford: Basil Blackwell.

Reich, W. 2000. “Heuristics as Plausible Models of Rationality?” *Acta Sociologica* 43:251–58.

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.

Simão, J., and P.M. Todd. 2003. “Emergent Patterns of Mate Choice in Human Populations.” *Artificial Life* 9:403-417.

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.

Todd, P.M. 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.

Vaupel, J.W., J.R. Carey, K. Christensen, T.E. Johnson, A.I. Yashin, N.V. Holm, I.A. Iachine, A.A. Khazaeli, P. Liedo, V.D. Longo, Y. Zeng, K.G. Manton and J.W. Curtsinger. 1998. “Biodemographic Trajectories of Longevity.” *Science* 280(5365):855–60.

Figure Captions

FIGURE 1. Hazard function for marriage (Number of first marriages of people who attain a given age *x* in a year by the number of still-unmarried individuals of age *x-1* at the beginning of the year). Source: own elaborations on Eurostat, New Cronos database.

FIGURE 2. Hazard function for marriage according to formal demographic models fitted to cohort data (Italian women born 1946-55, FFS individual-level data). Source: Billari (2000).

FIGURE 3. Hazard function for marriage in a population of simulated agents. Each agent searches in a noncompetitive one-sided setting, using an aspiration level to guide search after an initial “adolescence” learning period. **a**. The aspiration level is set to the highest value seen in the first 12 possible mates. **b**. The aspiration level is set to the highest value seen in the first *k* mates, where *k* is drawn from a normal distribution with mean 12, standard deviation 3 (i.e., adolescence length is normally distributed).

FIGURE 4. Hazard function for marriage in a population of simulated agents. Each agent searches in a competitive two-sided setting, using an aspiration level slightly below its own mate value to guide search after an initial “adolescence” period (which in this case involves no learning). **a**. The adolescence period covers the first 12 possible mates. **b**. The adolescence period is normally distributed with mean 12 and standard deviation 3.

FIGURE 5. Hazard function for marriage in a population of simulated agents. Each agent searches in a competitive two-sided setting, using an aspiration level to guide search after an initial “adolescence” learning period. **a**. The aspiration level is adjusted according to acceptances and rejections (see text) received during the first 12 possible mates. **b**. The aspiration level is adjusted according to acceptances and rejections received during the first *k* mates, where *k* is drawn from a normal distribution with mean 12, standard deviation 3 (i.e., adolescence length is normally distributed).

FIGURE 6. Hazard function for marriage in a population of simulated agents using extended courtship. Each agent begins searching for a marriage partner starting at age 12, using an aspiration level that varies with experience to decide when to initiate a courtship period. During courtship, individuals can switch to another better potential partner if one comes along; after courting one individual for a specific required period of time, marriage occurs. The solid line shows the hazard rates for marriage with age when all individuals have a 5-year required courtship duration; the dashed line shows the hazard rates when each individual has a required courtship period drawn from a normal distribution with mean 5 years and standard deviation of 3 years.

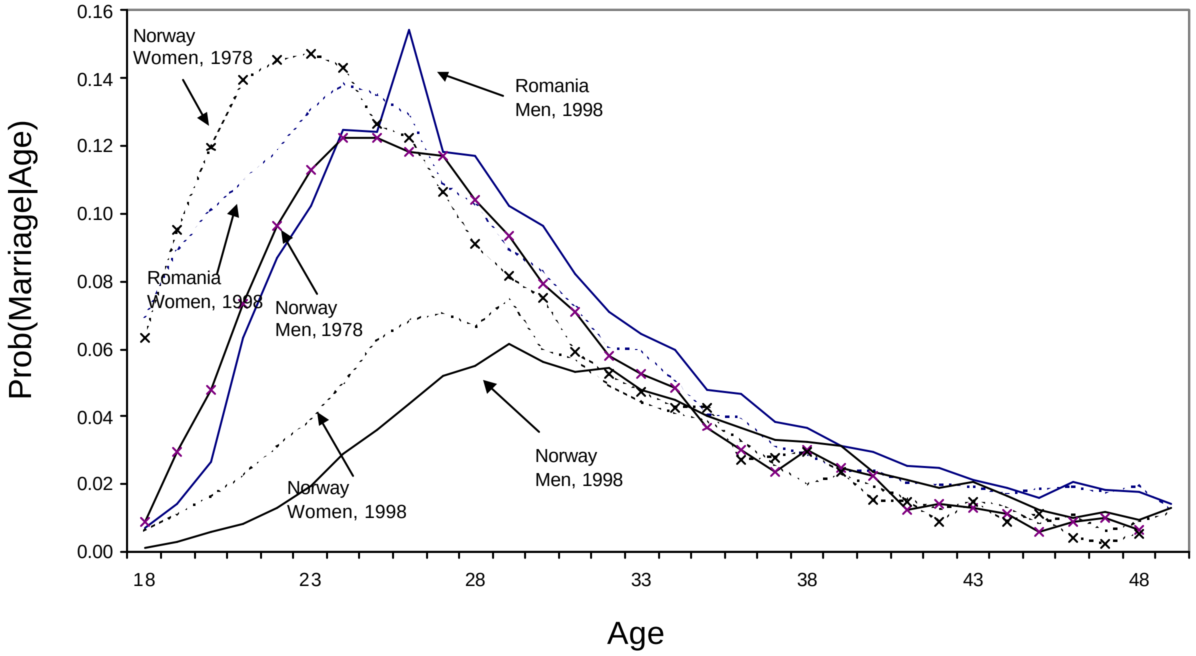
0.16
0.14
0.12
0.10
0.08
0.06
0.04
0.02
0.00
Prob(Marriage|Age)
18
23
28
33
38
43
48
Age
Norway
Women, 1978
Romania
Men, 1998
Romania
Women, 1998
Norway
Men, 1978
Norway
Women, 1998
Norway
Men, 1998

Figure 1

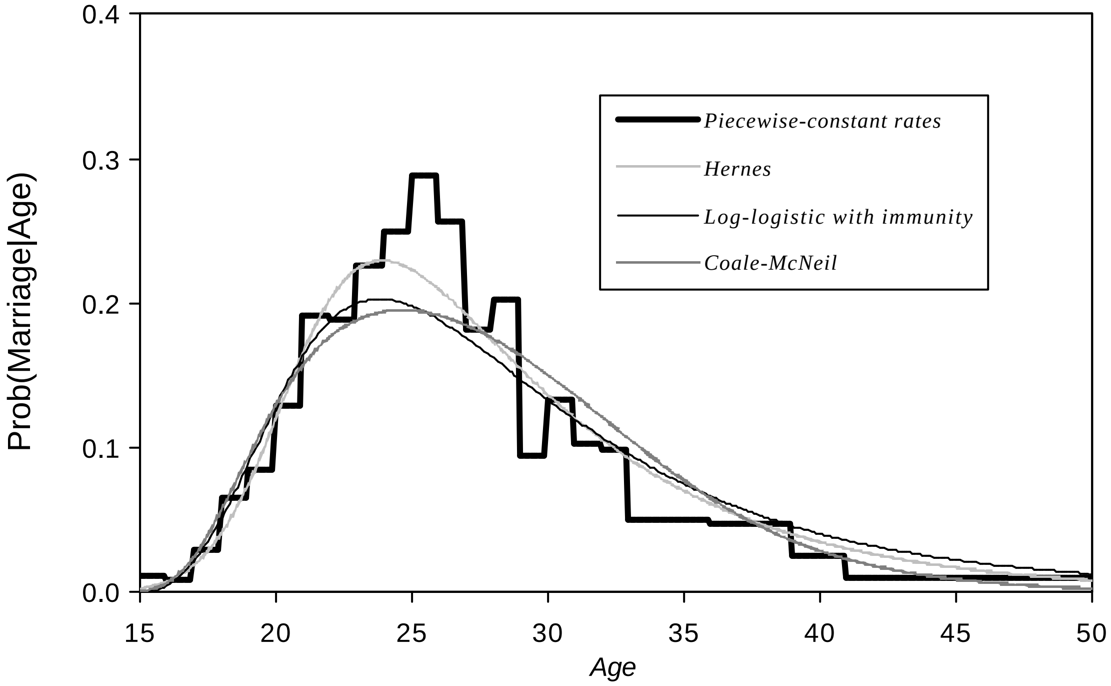
0.4
0.3
0.2
0.1
0.0
Prob(Marriage|Age)
15
20
25
30
35
40
45
50
Age
Piecewise-constant rates
Hernes
Log-logistic with immunity
Coale-McNeil

Figure 2

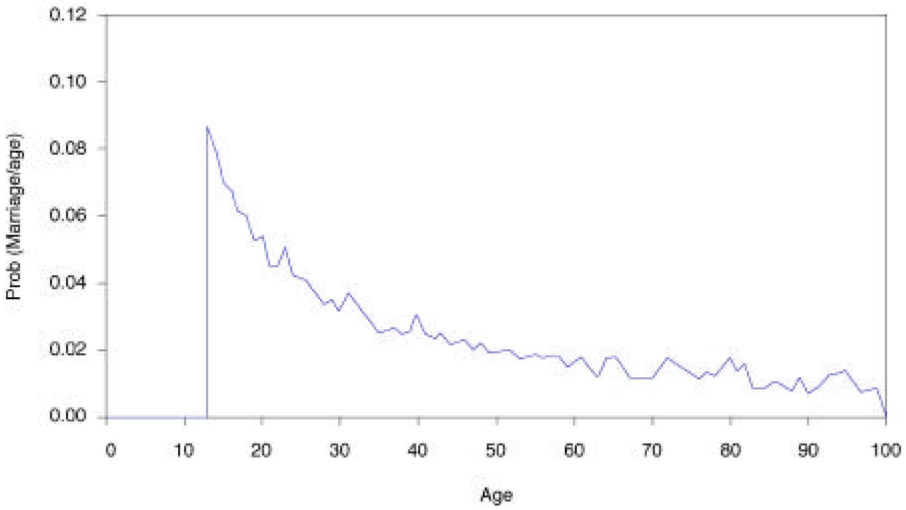

Figure 3a

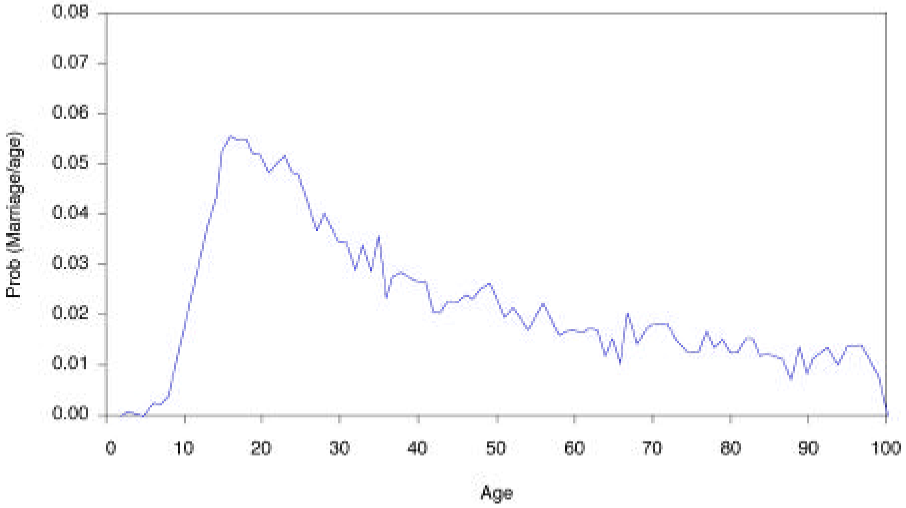

Figure 3b

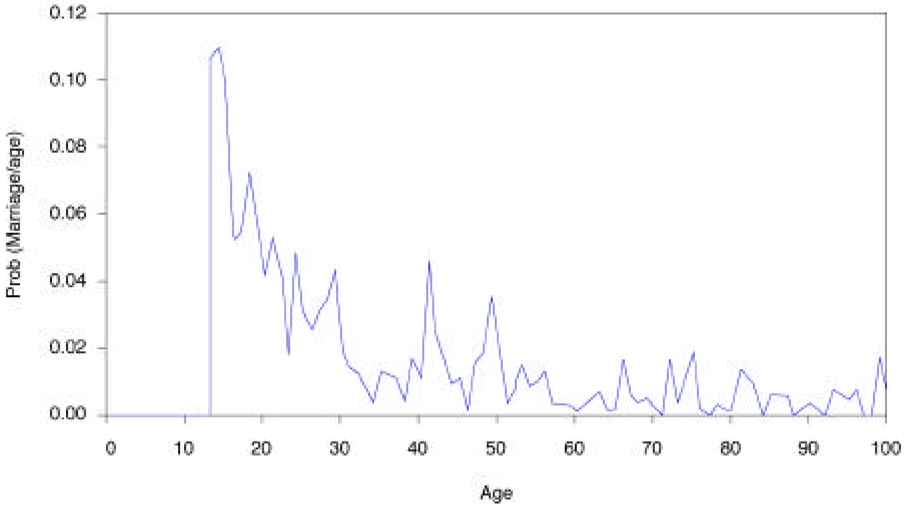

Figure 4a

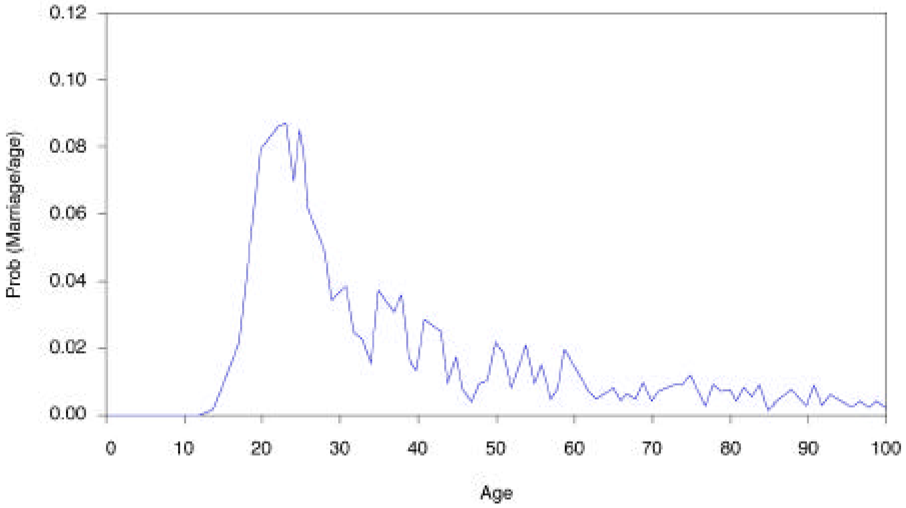

Figure 4b

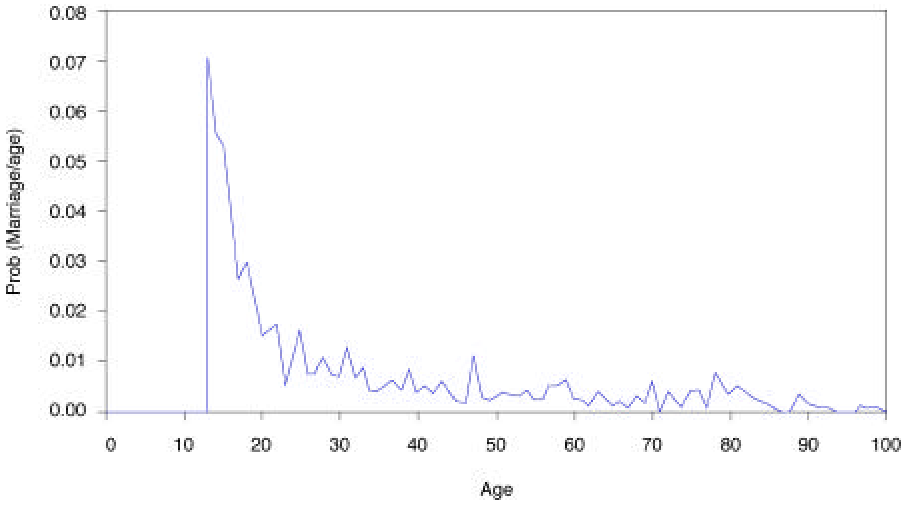

Figure 5a

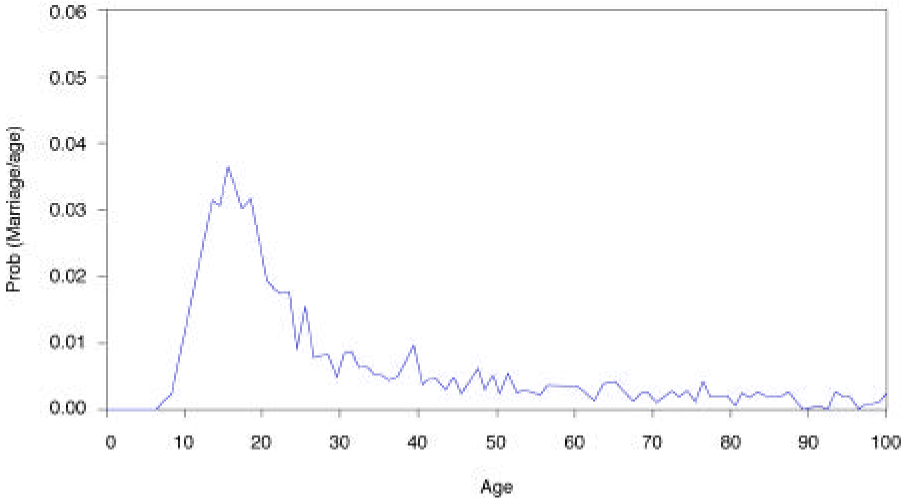

Figure 5b

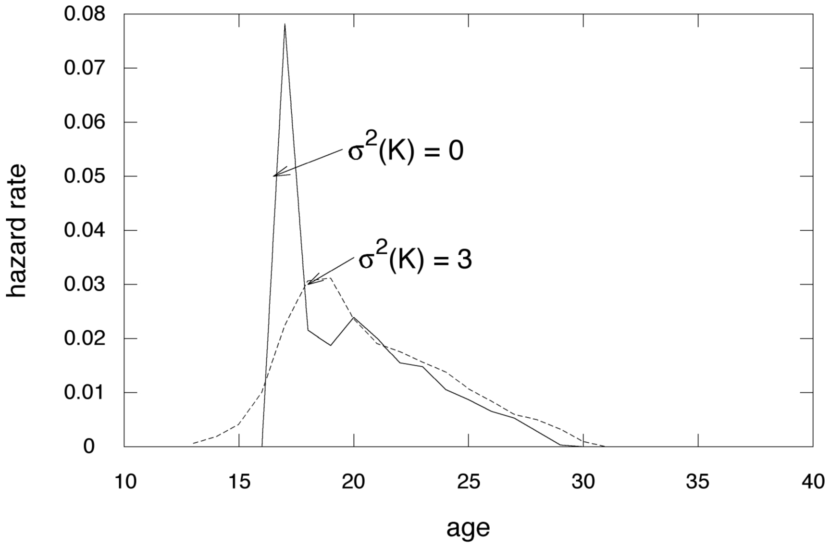
0.08
0.07
0.06
0.05
0.04
0.03
0.02
0.01
0
hazard rate
10
15
20
25
30
35
40
age
σ²(K) = 0
σ²(K) = 3

Figure 6