Population-Wide Marriage Patterns Produced by Individual Mate-Search Heuristics

Peter M. Todd[1] and Francesco C. Billari[2] *

[1] Center for Adaptive Behavior and Cognition, Max Planck Institute for Human Development, Lentzeallee 94, D-14195 Berlin, Germany; email: ptodd@mpib-berlin.mpg.de

[2] Institute of Quantitative Methods, Bocconi University, viale Isonzo 25, I-20135 Milano, Italy; email: francesco.billari@uni-bocconi.it

Abstract. The psychological problem of uncovering the mechanisms by which people select mates or marriage partners and the demographic problem of understanding the emergence of population-level patterns of marriage can be brought together to illuminate each other. In this paper we combine a top-down demographic approach with a bottom-up psychological approach to study marriage and mate search via agent-based simulations. We model a group of agents searching for mates using individual aspiration-based *satisficing* heuristics, which are simple, psychologically plausible mechanisms suggested by the framework of *bounded rationality*. By comparing the resulting population-level patterns with the general features of demographically observed age-at-marriage distributions, we find that the psychological mechanisms must be refined to account for the empirical data. In particular, we show the importance of heterogeneity across individuals for producing realistic marriage times.

1 Introduction

The choice of a partner for marriage or cohabitation is one of the key events in the course of our lives, and as such is a topic of keen personal interest for most adults. We observe it and talk about it at both the aggregate level ("Why are all my friends getting married before me?") and the individual ("How did she ever find a man like that?"). But, probably owing to the complexity of jumping between individual and aggregate perspectives, we seldom seem to put the two levels together ("Maybe everyone is getting married by this age because they are all looking for a partner in the same way she did").

The scientific study of marriage has been similarly divided. A long tradition of 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

* We would like to thank Tom Burch, Henriette Engelhardt, and Padma Rao Sahib for their valuable comments and suggestions. This paper was written while Francesco C. Billari was working at the Max Planck Institute for Demographic Research.

to cohabit or marry. But this bottom-up micro view omits the patterns that emerge in a group of such deciding individuals. It seems apparent that the two perspectives, individual and group level, have data and hypotheses that can help to constrain and explain the other. But how can the bottom-up and top-down approaches be brought together in the middle to speak to each other?

To bridge the gap in the study of the marriage market, we have developed agent-based models that simulate the mate search and choice behavior of individual agents interacting in a group. These simulations allow us to control and monitor the micro-level decision mechanisms of each agent, and observe 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 [17,14].

In the rest of this paper, we present our efforts to combine demographic and psychological approaches to marriage via agent-based modeling. 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 demographic behavior.

2 The Demographic Patterns of Age at Marriage

For over a century, demographers have been concerned with the careful analysis of empirical patterns related to the important events of human life: mortality, fertility, migration, and marriage. At times this empirical investigation has been independent from behavioral theories, while in other cases such theories (usually built by economists or sociologists, owing to demography's historic connection with those fields) have been used to derive formal behavioral models of the observed population data. The study of the demographic patterns of the age at marriage is no exception. In this section, starting from some facts about the distributions of age at first marriage that are well-known to demographers, we briefly review the main attempts at modeling this age-dependent process from a formal point of view. [1]

The idea that the age pattern of first marriage has been roughly stable (at least in a qualitative way) across places and historical periods became popular among

[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.

population scholars following the work of Coale [9,10]. Coale started from the observation that the frequency distribution of the age at which people marry for the first time is often a unimodal curve, with a long tail to the right. This empirical observation led Coale and McNeil to develop a demographic model able to fit distributions of the age at first marriage in different populations, including contemporary and historical populations in various geographical areas. (Their "standard" distribution for reference is that of Sweden in 1865-69.) Coale and McNeil's model predicts the frequency distribution of the age at first marriage by using only three parameters: 1) the proportion of individuals eventually getting married; 2) the mean age at first marriage; and 3) the variance of the distribution of ages at first marriage. Coale and McNeil (1972) introduced a behavioral interpretation of the model, in which the process of first marriage is seen as stochastically determined by the transition through a sequence of latent states in young adulthood. People start out being not marriageable, and they then enter into a marriageable state according to some given waiting-time distribution. Coale and McNeil's formula has subsequently been applied to a vast amount of demographic data, and the equation underlying their demographic model has been found to describe a vast amount of demographic data.

Instead of looking at the frequency distribution of ages at marriage as Coale and McNeil did, it is more suitable from the behavioral point of view to look at the *hazard rate* of marriage. This rate, defined either in discrete or continuous time, is the probability of marriage during the next time interval (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.) Coale and McNeil's mathematical formulation of the unimodal frequency distribution with a long tail to the right implies a skewed-bell shape also for this hazard rate. That is, the probability of marrying within an age interval for those who are not married at the beginning of the interval first increases with age, and then slowly decreases. In addition, the fact that hazard rates monotonically decrease after a certain age implies, if the speed of decrease is high enough, that not every member of a cohort will marry. (This proportion of the population that never marries is captured explicitly in the first parameter of the model as listed in the previous paragraph.)

Hazard rates are preferable for understanding behavior because they are not influenced by the probability of surviving in a certain state – here, the unmarried state – up to the starting point of the age interval. Hazard rates give then a better representation of the age-dependence of the "risk" of experiencing a certain event. 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 [36]. 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 implies that the survival function in the case of marriage, as opposed to death, never reaches the zero level.)

To illustrate the shape of these hazard rates for marriage, we show in Fig. 1 the empirically observed functions (in the form of 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. [2] 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 very different for Norway 1998, where non-marital cohabitation is widespread (and thus many people never make the transition to marriage although they make the transition to a stable co-residential partnership). In addition, hazard rates tend to converge to a level close to zero at later ages. Such a pattern can be observed for several other populations.

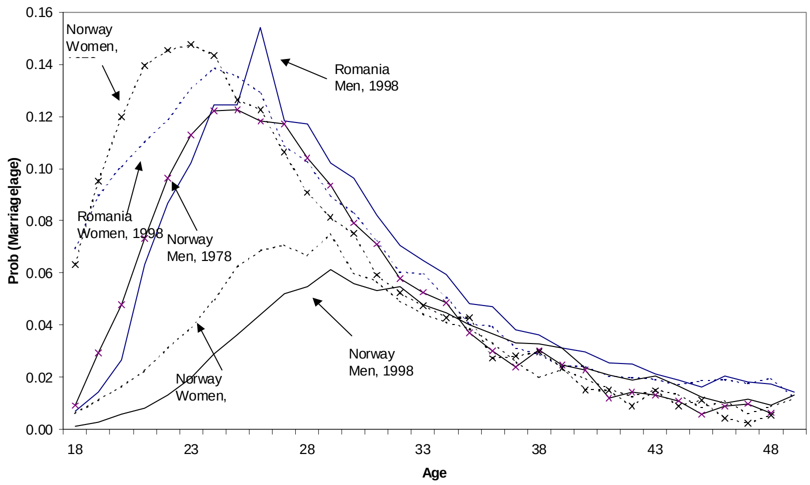

Fig. 1. Hazard function for marriage (number of first marriages of people who attain a given age x in a year divided by the number of still-unmarried individuals of age $x-1$ at the beginning of the year. *Source: own elaborations on Eurostat, New Cronos database*

To explain the skewed unimodal hazard rate to first marriage, three main types of formal behavioral models have been proposed [12] : latent-state (or compartment) models (as first introduced by Coale and McNeil [10]), diffusion models, and search models. Such models are usually applied to analyze the behavior of a cohort of individuals as they age. While these models can assume heterogeneity of characteristics among individuals, in most cases they are based on a top-down approach that can-

[2] 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.

not explicitly accommodate interactions between individuals or complex individual behavioral patterns.

Latent-state models of first marriage, like Coale and McNeil's 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 and 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 [19,18]. 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 sociological aspects of the search process (e.g., learning, mutual choice, and diffusion – see [6,11]).

In diffusion models, mating happens by "contagion" from other people who are already mated. The model developed by Hernes [20] for example is based on the idea that (first) marriages are influenced by two opposite 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. The combined effect of both these forces on the diffusion process produces a unimodal pattern for marriage hazard rates. 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, Bruederl and Diekmann [5] proposed a log-logistic model with immunity, in which the diffusion process affects only a certain share of a cohort. Billari [4] 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 subgroups 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 Fig. 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 [3]) We compare the fitted Coale and McNeil, Hernes, and log-logistic with immunity models against a weakly 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 [27]) 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 data as well.

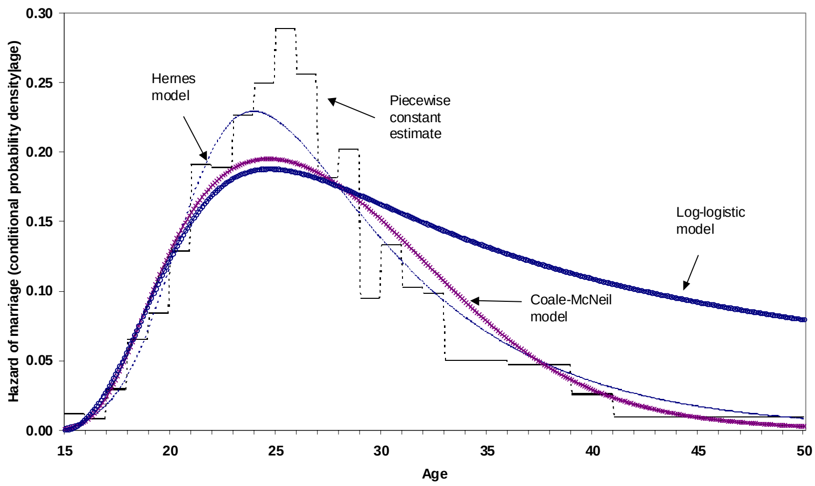

Fig. 2. Hazard function for marriage (continuous time = according to formal demographic models fitted to cohort data (Italian women born 1946-55, FFS individual-level data). *Source: [3]*

Finally, in the third type of model, typically based on economic search theory, agents act according to some search mechanism to seek mates in a reasonable (or perhaps somehow optimal) manner [7]. Individuals instantiated in these models can select possible mates, for instance by making and accepting marriage offers, in potentially heterogeneous ways, 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. Keeley [24], 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 extensive knowledge and full rationality on the part of the individual, and thus this model is subject to the common criticisms of such unrealistic assumptions [8]. In the next section, we start from the same interest in modeling individual search processes, but assuming instead that individuals act according to *bounded* rationality.

3 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?

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 [1]. 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 are not seen 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.

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 assessments of the costs of foregone opportunities and complex calculations such as Bayesian updating of probability estimates, and could thus require considerable time and computational resources.

But to make choices in a useful amount of time, real agents must employ simple mechanisms for limiting the 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 [33]. And indeed, there is considerable evidence that people faced with sequential search tasks use simple rules to make their choices [21,26,29,13] . 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 imposed by the task environment and human cognitive capacities [31]. 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* [16], 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 [13]). 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 – [16]. One of the main ways to gather empirical evidence for the use of simple heuristics is through laboratory experiments in which individuals make decisions in a controlled setting. 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, they need a stopping rule. What kinds of simple stopping rules are reasonable for our marriage model? For more or less 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 [31]. 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.[3]

3.1 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 initially 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*[4]. In the dowry problem it is the searcher's aim to

[3] 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).

[4] It is also known in the statistical literature as the *secretary problem* [15,29]. 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

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 [15] 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 rate 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 [22,29].

Instead, Todd [32,34] 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 [34] 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 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 has to be extended (along with determining whether or not mate search conforms to such situations in the first place). Seale and Rapoport [29] 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

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.

Todd [13] 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 suggests that different heuristics may be employed across the different variations of the one-sided search task, and hence we must gather more evidence to determine just what psychological mechanisms are 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 [2] recognized that the simple search rules proposed by Todd [32] 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 [2] 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 (hence selecting or "marrying") 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. [5] 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 Fig. 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.

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

[5] Note that here choosing a mate is equated with marriage, and the number of potential partners seen is equated with age. The second mapping, partners linearly increasing with 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.

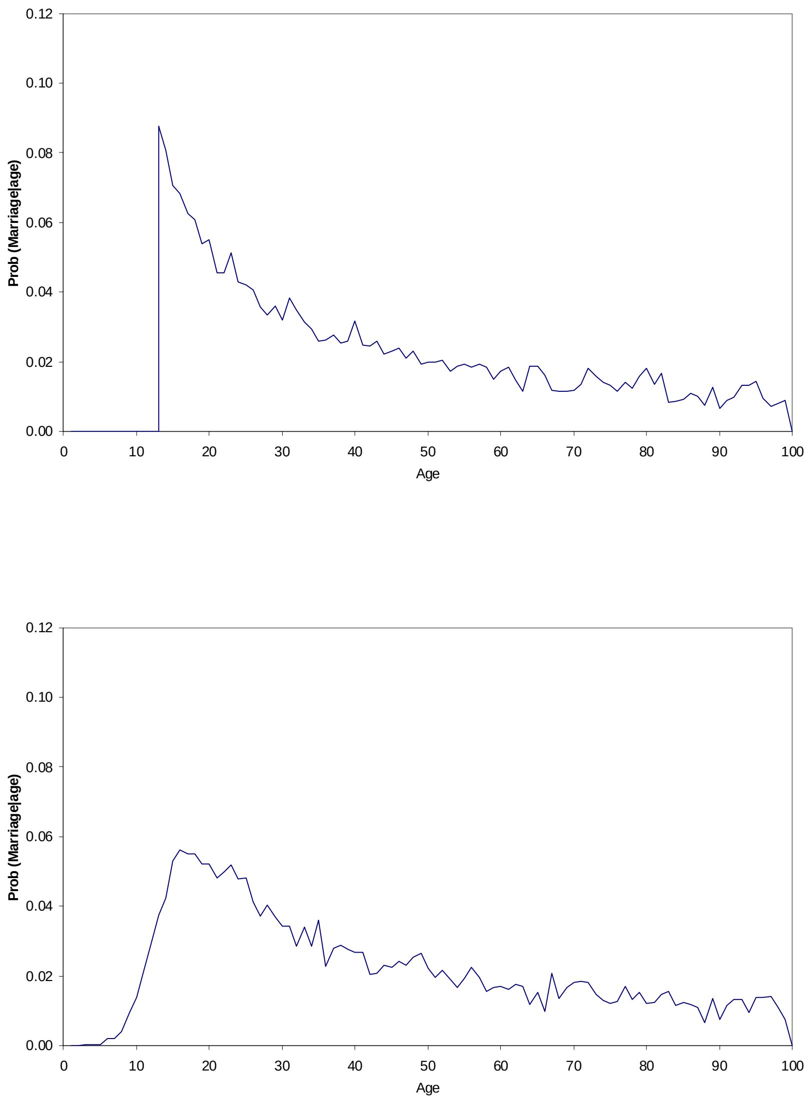

Fig. 3. Hazard function for marriage in a population of simulated agents. Each agent searches in a noncompetitive one-sided setting, using an aspiration level to guide search after an initial "adolescence" learning period. Upper figure *(a)*: The aspiration level is set to the highest value seen in the first 12 possible mates. Lower figure *(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)

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 (Fig. 3b; the same effect occurs for uniform distributions, but is not shown). [6]

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 range 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.

3.2 Two-Sided (Mutual) Search Processes

One-sided mate search, in which the members of one sex do the searching and make the decisions, is not a realistic model of the adaptive problem of mate choice actually faced by many species, including humans. (Indeed, recent evidence indicates that, even in some species where it was long assumed that females do all the choosing and males mate indiscriminately, males also actively decide whether a female is suitable to court – see [28]). The problem for these species 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

[6] 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.

the same set of potential mating partners. 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, then everyone will end up with a rather high aspiration regarding 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) would be chosen as mates, and most individuals would end up alone.

Clearly, individuals must be using other rules in mutual mate search situations. To explore how different search rules can work in a two-sided setting, we must amend our earlier model (see [34] for more details on the models, search rules, and results discussed here). This time, 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. In our first runs of this model, we give each of the 200 individuals the same search strategy, all of which we based on the aspiration-level-setting approach of satisficing.

Each individual begins his or her simulated life by assessing and making (or not) practice 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 a mating offer (which however cannot result in actual pairing during this initial period). Over this time, individuals can also adjust their aspiration level up or down from an initial value. [7] After this adolescence period, males and females are paired up at random in each time period, 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 mated, and the two individuals are removed from the population. Otherwise, both individuals remain in the mating pool to try again with someone else in the next time period. This pairing-offering-mating cycle is repeated over successive periods until every individual is mated, or until every individual has had the opportunity to assess and propose to every member of the opposite sex.

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 paired up (i.e., mate), how well matched the pairs end up being, and when the pairings occur. We can compare search rules along these dimensions not only against each other, but also against

[7] 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.

empirical 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 [35]. (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, much social psychology research in the 1960's and 1970's found that there is a significant correlation between the attractiveness of the two people in married couples; estimates ranged between .4 and .6 in various studies [23]. We take attractiveness here as a rough proxy for mate value; similar within-pair correlations have been found for other traits that could be related to mate value (e.g., intelligence, height) as well. And third, we have the empirical 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 [34]. 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 here that one's own mate value is initially known. With this approach, most of the population can succeed in finding and pairing up with mates of a similar value to their own [34,25]. When we look at the hazard function for marriage, however, we again see the problematic exponentially-decreasing function (Fig. 4a) that appeared in the one-sided search case – clearly, changing the search setting to two-sided choice does not by itself lead to a realistic distribution of marriage times, as we had at first hoped. Again, though, the same alteration used before of introducing variation in learning times (here, letting the adolescence period vary normally) creates the familiar unimodal curve, as shown in Fig. 4b.

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?

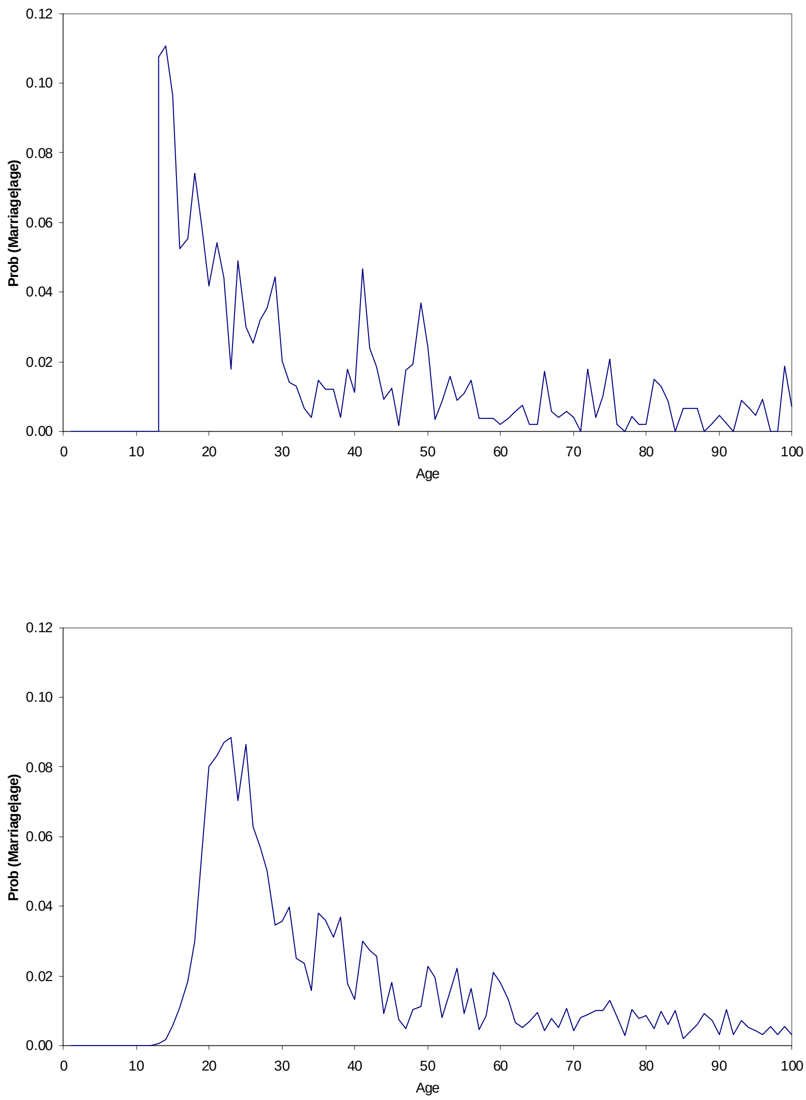

Fig. 4. Hazard function for marriage in a population of simulated agents. Each agent searches in a competitive two-sided setting, using an aspiration level slightly below its own mate value to guide search after an initial "adolescence" period (which in this case involves no learning). Upper figure *(a)*: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

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 the majority of the population to form mated pairs of individuals with similar mate value [34]. In fact, setting an aspiration level by searching through many *more* than this during adolescence (e.g., over 20 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 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 (Fig. 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 (Fig. 5b).

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. 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 (adolescence) periods resulted in the unimodal age-at-marriage

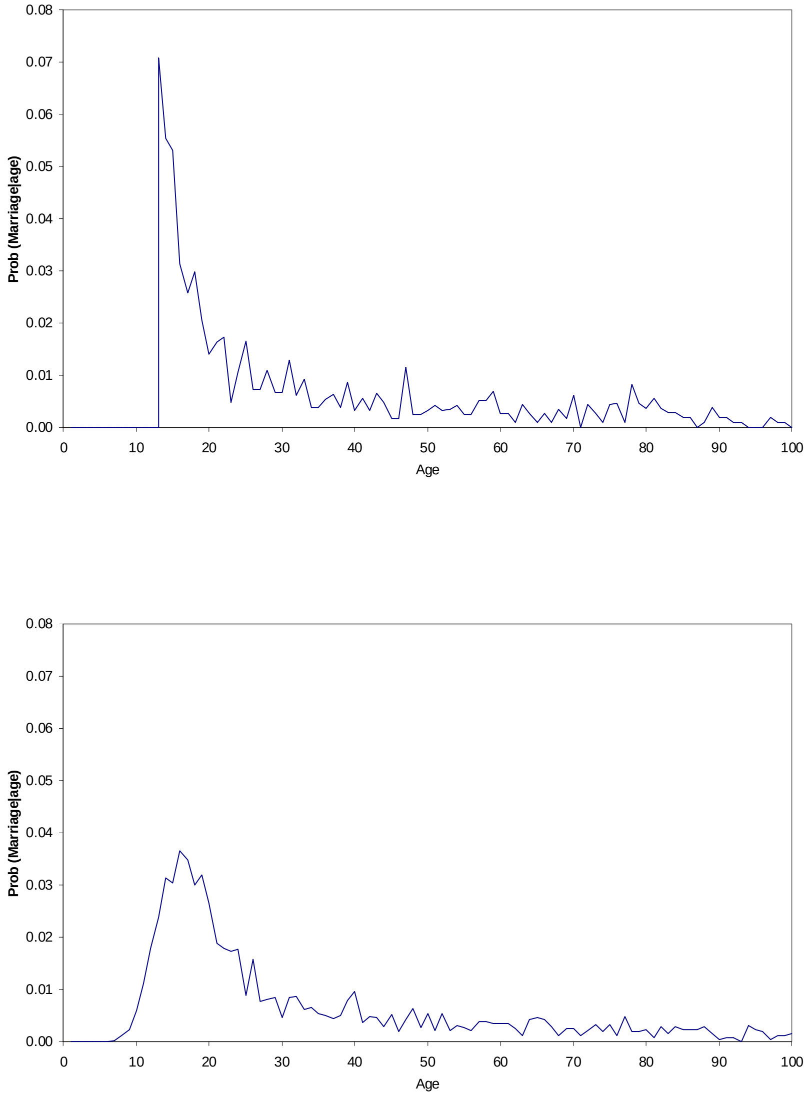

Fig. 5. Hazard function for marriage in a population of simulated agents. Each agent searches in a competitive two-sided setting, using an aspiration level to guide search after an initial "adolescence" learning period. Upper figure *(a)*: The aspiration level is adjusted according to acceptances and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

curve; varying the distribution of mate values (quality levels) or initial aspiration levels did not have an appreciable effect.

4 Implications and Conclusions

By combining the top-down demographic approach to the topic of marriage (or mate search) with the bottom-up psychological modeling approach, we have been able to illuminate both. 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 the population-level outcomes of individual search mechanisms has provided a more psychologically satisfying (and satisficing) explanation of the demographic data, beyond merely pointing to latent stages or diffusion processes.

However, the psychological-mechanism-level explanation is still not entirely satisfying. There are of course many ways in which our simple model of individual search is unrealistic. While this starting point has at least allowed us to expand the range of reasonable explanations to consider for the demographic data, more psychologically plausible models would further enhance our understanding of the processes involved. For instance, our current model assumes that there is an abrupt switch from a learning period to a mate-choice period. Even if this were so, how could we measure it empirically? It seems more reasonable to propose either a gradual shift from learning to choosing, or else a life-long learning model in which aspiration levels change continuously with experience until a final choice is made (e.g., as in [30], which also allows individuals to change their minds and switch partners before marriage). The strict cutoff in the current model (accept any offer from an individual above a certain quality level, reject any below that) is also unrealistic – a graded acceptance function (or the introduction of noise into the quality-appraisal process) should result in more individuals finding mates, more quickly. Furthermore, the issue of cultural differences must be addressed; while similar age-at-first-marriage patterns have been observed in many countries at different points in time, there are clearly differences in marriage traditions, such that a sequential search model (at least at the individual level, as opposed to perhaps the parental or family level) may well not apply in many cases.

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 (not least because the definition of “partner” as used in our models remains vague). 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 between models, but again there are the problems of defining and collecting demographic data about mate quality. 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. Even just knowing more about the distribution of mate qualities of potential partners available in a given cohort would be important.

Clearly, we have opened up more questions than we have answered. But the fact that this research project has generated these new avenues to be pursued is another advantage of bringing together demography and psychology in this way. 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. Pulling the top-down and bottom-up approaches together in the middle can generate both new understanding, and new opportunities for further exploration.

References

1. Bergstrom, C.T., Real, L.A. (2000): Towards a theory of mutual mate choice: Lessons from two-sided matching. Evolutionary Ecology Research 2(4), 493-508
2. Billari, F.C. (2000a): Searching for Mates Using 'Fast and Frugal' Heuristics: a Demographic Perspective. In: Ballot, G., Weisbuch, G. (Eds.): Applications of Simulation to Social Sciences. Hermes Science Publications, Oxford, UK, 53-65
3. Billari, F.C. (2000b): L'analisi delle biografie e la transizione allo stato adulto. Aspetti metodologici e applicazioni ai dati della Seconda Indagine sulla Fecondità in Italia. Cleup editrice, Padova, Italy
4. Billari, F.C. (2001): A log-logistic regression model for a transition rate with a starting threshold. Population Studies 55, 15-24
5. Brüderl, J., Diekmann, A. (1995): The Log-Logistic Rate Model. Two Generalizations With an Application to Demographic Data. Sociological Methods & Research 24, 158-86
6. Burch, T.K. (1993): Theory, Computers and the Parameterization of Demographic Behavior. In Proceedings of the International Population Conference, Montreal, Canada. International Union for the Scientific Study of Population, Liege, Vol. 3, 377-388
7. Burdett, K., Coles, M.G. (1999): Long-Term Partnership Formation: Marriage and Employment. The Economic Journal 109, 307-334
8. Chase, V.M., Hertwig, R., Gigerenzer, G. (1998): Visions of rationality. Trends in Cognitive Science 2 (6), 206-214
9. Coale, A.J. (1971): Age Patterns of Marriage. Population Studies 25, 193-214
10. Coale, A.J., McNeil, D.R. (1972): The distribution by age of the frequency of first marriage in a female cohort. Journal of the American Statistical Association 67, 743-749
11. Coale, A.J., Trussell, J. (1996): The Development and Use of Demographic Models. Population Studies 50: 469-484
12. Diekmann, A. (1989): Diffusion and Survival Models for the Process of Entry into Marriage. Journal of Mathematical Sociology 14: 31-44
13. Dudey, T., Todd, P.M. (in press): Making good decisions with minimal information: Simultaneous and sequential choice. Journal of Bioeconomics
14. Epstein, J.M., Axtell, R. (1996): Growing Artificial Societies: Social Science from the Bottom Up. MIT Press, Cambridge, MA
15. Ferguson, T.S. (1989): Who solved the secretary problem? Statistical Science 4(3), 282-296
16. Gigerenzer, G., Todd, P.M., the ABC Research Group (1999): Simple heuristics that make us smart. Oxford University Press, New York
17. Gilbert, N., Troitzsch, K.G. (1999): Simulation for the Social Scientist. Open University Press, Buckingham, UK

18. Goldstein, J.R., Kenney, C.T. (2001): Marriage Delayed or Marriage Forgone? New Cohort Forecasts of First Marriage for U.S. Women. American Sociological Review 66, 506-19
19. Henz, U., Huinink, J. (1999): Problems concerning the parametric analysis of the age at first birth. Mathematical Population Studies 7(2), 131-145
20. Hernes, G. (1972): The Process of Entry into First Marriage. American Sociological Review 37, 173-182
21. Hey, J. D. (1982): Search for Rules for Search. Journal of Economic Behavior & Organization 3, 65-81
22. Kahan, J.A., Rapoport, A., Jones, L.V. (1967): Decision making in a sequential search task. Perception & Psychophysics 2, 374-376
23. Kalick, S.M., Hamilton, T.E. (1986): The matching hypothesis reexamined. Journal of Personality and Social Psychology 51(4), 673-682
24. Keeley, M.C. (1979): An Analysis of the Age Pattern of First Marriage. International Economic Review 20 (2), 527-544
25. Miller, G.F., Todd, P.M. (1998): Mate choice turns cognitive. Trends in Cognitive Sciences 2, 190-198
26. Moon, P., Martin, A. (1990): Better heuristics for economic search—experimental and simulation evidence. Journal of Behavioral Decision Making 3, 175–193.
27. Raftery, A.E. (1995): Bayesian model selection in social research. In: Marsden P.V. (Ed.): Sociological Methodology 1995. Basil Blackwell, Oxford, 111-196
28. Sæther, S.A., Fiske, P., Kålås, J.A. (2001): Male mate choice, sexual conflict and strategic allocation of copulations in a lekking bird. Proceedings of the Royal Society of London B: Biological Sciences 268(1481), 2097-2102
29. Seale, D.A., Rapoport, A. (1997): Sequential decision making with relative ranks: An experimental investigation of the "Secretary Problem". Organizational Behavior and Human Decision Processes 69, 221-236
30. Simão, J., Todd, P.M. (in press): Modeling mate choice in monogamous mating systems with courtship. Adaptive Behavior
31. Simon, H.A. (1990): Invariants of human behavior. Annual Review of Psychology 41, 1-19
32. Todd, P.M. (1997): Searching for the next best mate. In: Conte, R., Hegselmann, R., Terna, P. (Eds.): Simulating social phenomena. Springer-Verlag, Berlin, 419-436
33. Todd, P.M. (2000): Fast and frugal heuristics for environmentally bounded minds. In: Gigerenzer, G., Selten, R. (Eds.): Bounded rationality: The adaptive toolbox (Dahlem Workshop Report). MIT Press, Cambridge, MA, 51-70
34. Todd, P.M., Miller, G.F. (1999): From pride and prejudice to persuasion: Satisficing in mate search. In: Gigerenzer, G., Todd, P.M., the ABC Research Group: Simple heuristics that make us smart. Oxford University Press, New York, 287-308
35. United Nations (1990): Patterns of First Marriage: Timing and Prevalence. United Nations, New York
36. Vaupel, J.W., Carey, J.R., Christensen, K., Johnson, T.E., Yashin, A.I., Holm, N.V., Iachine, I.A., Khazaeli, A.A., Liedo, P., Longo, V.D., Zeng Yi, Manton, K.G., Curtsinger, J.W. (1998): Biodemographic Trajectories of Longevity. Science 280(5365), 855-860