

Searching for the next best mate

Peter M. Todd
Center for Adaptive Behavior and Cognition
Max Planck Institute for Psychological Research
Leopoldstrasse 24, 80802 Munich GERMANY
ptodd@mpipf-muenchen.mpg.de

Abstract. How do we humans go about choosing a mate? Do we shop for them, checking prices and values and selecting the best? Do we apply for them, wooing several and taking the best that accepts us in return? Or do we screen them, testing one after another in succession until the right one comes along? Economists and other behavioral scientists have analyzed these mate-choice approaches to find their optimal algorithmic solutions; but what people really do is often quite different from these optima. In this paper, we analyze the third approach of mate choice as applicant screening and show through simulation analyses that a traditional optimal solution to this problem—the *37% rule*—can be beaten along several dimensions by a class of simple “satisficing” algorithms we call the *Take the Next Best* mate choice rules. Thus, human mate search behavior should not necessarily be compared to the lofty optimal ideal, but instead may be more usefully studied through the development and analysis of possible “fast and frugal” mental mechanisms.

1 Introduction

Some people view the process of finding a mate as similar to that of shopping for tomatoes: you know what you want, but you do not know what you will find, so you drive around to different stores, squeezing the merchandise and checking the price at each one. You drive to as many stores as necessary to find the price you like, taking into account the cost of your time and energy as you go. Finally, you return to the store with the best price (usually the last one you visited), where the tomatoes will be ready and waiting to go home with you, and you make your purchase.

Others view the mating game as something more like applying to college: you know all the options ahead of time, and you know how much you like each one, but you do not yet know how much they like you in return. So you court them all, and wait for their replies. The cost of courting is trivial compared to the ultimate cost of the four (or more) year long relationship with them, so this is not taken into account. The coy schools eventually make

their choices as well, and you end up with the first one on your preference list that also favored you.

The real situation in human mate choice is, of course, somewhere in-between these two idealized cases. As in tomato shopping, the full list of candidates is never known ahead of time; but as in college admissions, the candidates are rating and deciding on you in return. Unlike tomato shopping, the costs associated with search are not simply time and energy, and the candidates are not waiting for you to make up your mind; unlike college shopping, there are real costs associated with considering extra candidates, and the rank ordering of candidates cannot be determined all at once.

Instead, we can characterize mate choice as a sequential search process, more akin to hiring a new employee for a job opening: you can interview a succession of candidates, but in this hypothetical setting you must decide whether or not to hire each one as you see them—you cannot go back to earlier candidates that you passed over, because they have gone off and been hired somewhere else. You cannot know the range of qualifications of candidates ahead of time; instead, you must figure this out as you conduct interviews, and use that information to decide when to stop and hire the current candidate. You may not even know how many candidates you could possibly look at, further complicating your search process. In this situation, how are you to decide when to stop searching and settle down with the current candidate? (Note that this assumes the current candidate is indeed available to you—but a final challenge may lurk: when you at last decide whom you want to hire, that person may not be impressed enough with your firm to accept your offer. We return to this challenge in the last section.)

In this paper, we explore some possible answers to the question of how to go about a sequential mate search. Economists and others have analyzed the mate search and matching problem from a variety of viewpoints, including those represented in these three stories, and have proposed optimal algorithms for each. But what people do is often quite different from what optimality would predict. We often behave in a bounded rational fashion, using simple algorithms and heuristics to achieve a “satisficing” solution to a problem in a faster, less information-intensive way. It is important to understand these algorithms and their information processing abilities, because this can also shed light on the nature of the adaptive problem these algorithms evolved to solve, and the kinds of information environments they evolved in. For example, as we will discuss later, it is commonly believed that humans do not search long enough when looking for a mate (Frey & Eichenberger, 1996); but the class of simple sequential search algorithms we explore here challenges this view, showing how robust search performance can be achieved without looking too long, and indicating that our evolved search behaviors could in fact be well-adapted to the mate environment in which they operate.

The problem of sequential search has been explored extensively over the last three decades in exquisite mathematical detail (see for example Fergu-

son, 1989, and Corbin, 1980). So why is there anything left to say about this problem, and why should we say it with simulations rather than explicit derivations? The answer to both questions is the same: because we want to explore a range of psychologically plausible alternative algorithms for performing adaptively in sequential search. The mathematical analyses of this problem have focused almost exclusively on finding the optimal solution under a variety of conditions. These solutions have often been quite complex, with little regard for their feasibility when implemented by real non-mathematician organisms. (Corbin, 1980, however, does take cognitive limitations into account.) Instead of this top-down analytical approach, we start here from the bottom and work our way up, asking how far very simple search rules can get us in this domain—something that the mathematical analyses have largely not addressed. To undertake this exploration, a convenient approach is to use individual-based simulation models, in which we instantiate mate search algorithms in artificial individuals who live and mate in representative environments. This way, we can test the algorithms’ performance directly. For simple situations, this simulation approach is merely an augmentation and confirmation of (or substitute for) mathematical analysis; but when the conditions become more complex, as for instance when we add in mutual mate search between the two sexes (in later papers), mathematical tractability begins to plummet, and the usefulness of simulation models for exploration becomes a necessity.

In the next section, we briefly discuss the past optimal and behavioral, satisficing approaches to mate search when conducted as tomato shopping and college applying, and then concentrate on the mate-choice-as-applicant-hiring (or, as we reframe it, dowry-maximizing) view of human mate search and its optimal solution, the so-called *37% rule*. Then in section three we discuss a class of simple satisficing approaches to this problem, which we call the *Take the Next Best* mate choice rules. We show how well these rules do in comparison to the optimal method through simple simulations with small populations, and how their behavior changes as the population of potential mates grows larger. We also indicate the implications of these results in terms of the “fast and frugal” (Gigerenzer & Goldstein, 1996) satisficing view of human behavior. Finally, we conclude with a discussion of the further directions that we can take to explore the mate search and choice algorithms that we humans really use.

2 Different types of mate search: optimal and actual behavior

Darwin (1871) began the serious evolutionary-minded study of how animals including humans choose their mates, and of the effects over time of these search and choice behaviors. In the past few years, research in this area has

exploded, both in terms of how humans (e.g. Buss, 1994) and other animals (e.g. Bateson, 1983) choose mates, and the evolutionary effects of sexual selection through mate choice (e.g. Ridley, 1993, and Miller, in press, for the effects on human evolution, and Andersson, 1994, and Cronin, 1991, for the effects on animal evolution; Miller & Todd, 1995, and Todd & Miller, 1993, in press, present simulation models of the evolutionary effects of mate choice). Because human mate search and choice are interesting examples of human choice behavior in general, and because they often have profound economic impact, economists have become interested in this form of behavior as well (see e.g. Becker, 1991; Roth & Sotomayor, 1990; Frey & Eichenberger, 1996; and Harrison & McCabe, 1996).

2.1 Mate search as shopping and applying

Economists have largely been interested in formulating the optimal decision rules for behavior in mate search and choice situations, and then seeing how well human behavior matches those optima. Thus, for instance, when the mate search task can be considered one of shopping for mates with a cost for search effort, the optimal rule is to keep checking potential mates until the first one that meets some precalculated reservation value based on the distribution of mate values and the search cost (Rapoport & Tversky, 1970). However, people do not generally seem to adhere to this optimal behavior. Hey (1982, 1987) has identified several non-optimal satisficing “rules of thumb” that people often use instead, such as the “one bounce” rule (keep checking values as long as they go up each time; as soon as they go down, stop and take the previous value). These non-optimal rules, though, are simpler to calculate (no need to compute the reservation value), and in practice achieve performance close to the optimum. (Martin & Moon, 1992, have since gone further, exploring a wider range of simple search rules both experimentally and through computer simulations in an approach similar to that we take here.)

In contrast, when the mate search task is described as one of mutual choice like applying for colleges, the procedure to follow to ensure that everyone gets a mate they are happy with is the Gale-Shapley algorithm (Gale & Shapley, 1962). This procedure produces a set of mated pairs from a population of males and females, and guarantees that no male and female who are not mated in a pair would both be happier if they were together. In this sense, the matings produced by the Gale-Shapley rule are stable (and optimal, if stability is the optimality criterion). This procedure requires that everyone has a fixed preference ranking of all their potential mates, and that they report this ranking accurately for use in the algorithm. However, as evidenced by the high divorce rates in many countries (Frey & Eichenberger, 1996), stable pairings are rather seldom obtained in reality, arguing against the use of a Gale-Shapley type algorithm by mate-seekers on an individual basis. And even when a centralized Gale-Shapley procedure is used to make pair

matches, people often misrepresent their true preference rankings to gain an advantage, and yet end up making the matching procedure work further against their desires (Harrison & McCabe, 1996).

The optimal matching research begun by Gale and Shapley has spawned a large literature, as reviewed by Roth and Sotomayor (1990). But the work in this area assumes that all parties involved have complete knowledge about their own preferences and the available options to choose among. In contrast, in the job search literature, preferences and options are usually assumed to be at least partially unknown, and the search process itself entails ongoing costs. Roth and Sotomayor (p. 247) indicate that these two kinds of models can be profitably brought together, to explore more realistic mutual search situations, but point out that this is still an open research area. This mutual partial-information costly-search model is also exactly what we need in the case of mate search. Here, we next consider the simplified version of this problem in which one party searches a stable set of sequential alternatives, but in section 4 we indicate the direction we are heading in for exploring the mutual search case.

2.2 A third view: sequential mate search with loss of candidates

We have just considered the idealized versions of mate choice as tomato shopping and applying to colleges, and the satisficing approaches people take to these sorts of problems. But just because we have found evidence for how people really behave in these kinds of settings does not mean that we know more about how people really choose mates: if these idealized versions of mate choice differ from real mate choice in important ways, then people's behavior in these idealized problems could also differ from their behavior in real mate choice in significant ways.

And these idealized versions do indeed differ from the situation that presents itself to men and women searching for a mate, at least in many modern Western cultures. This type of mate choice usually consists of a sequential search through successive potential mates, in which each one is evaluated and decided on in turn, in a process that can take minutes, hours, days, or years. (Here the decision can be thought of as whether or not to settle down and have children with a particular person, though other definitions are possible.) There are certainly costs associated with each person that one checks out during this search; but perhaps the most significant cost is that it is difficult, and often impossible, to return to a potential mate that has been previously discarded (because they remain in the "mating pool" and are likely to pair up with someone else in the meantime while one's own search continues—as countless romantic tragedies of missed opportunities attest). To further complicate matters, one does not know ahead of time what the range of potential mates may be: how can we know, during our first love, whether someone else might be able to incite still deeper passions, if we just keep

searching long enough to find them? We cannot even tell how many more potential mates we may encounter! Given these restrictions on the search process and lack of knowledge about the space over which we are searching, finding a mate looks, understandably, like a very daunting problem indeed.

Earlier we compared this sequential mate search problem to a firm trying to hire someone from a stream of job applicants. We can present this situation in more precise detail, and in a form more closely linked to mate choice, via the “dowry problem,” a well-known puzzle from probability theory (Mosteller, 1987; Gilbert & Mosteller, 1966; Corbin, 1980; this is also known as the “secretary problem” or “beauty contest problem” and even “Googol”). A sultan wishes to test the wisdom of his chief advisor, to know if he should retain this cabinet position. The chief advisor is seeking a wife, so the sultan takes this opportunity to judge his wisdom: the sultan arranges to have 100 women from the kingdom brought before the advisor in succession, and all the advisor has to do to retain his post is to choose the woman with the highest dowry (marriage gift from her family). If he chooses correctly, he gets to marry that woman, and keep his post; if not, he remains single, and loses his head. The advisor can see one woman at a time and ask her dowry; then he must decide immediately if she is the one with the highest dowry out of all 100 women, or else let her pass by and go on to the next woman. He cannot return to any woman he has seen before—once he lets them pass, they are gone forever. Moreover, the advisor has no idea of the range of dowries before he starts seeing the women. What strategy can he possibly use to have the highest chance of picking the woman with the highest dowry?

It turns out that the algorithm the advisor should use, to guarantee the highest chance of choosing correctly, is as follows: He should look at the first 37 women, letting each one pass, but remembering the highest dowry from that set—call this value D . Then, starting with the 38th woman, he should select the first woman with a dowry greater than D . (For derivations of this procedure, see Mosteller, 1987; Gilbert & Mosteller, 1966; Ferguson, 1989). If the advisor follows this procedure, he will have a 37% chance of succeeding in choosing the woman with the highest dowry. This is the best he can do—no other algorithm would improve his odds of picking correctly.

More generally, if the advisor knows that he will be presented with a succession of women from a total set of N women, he should check the first 37% of those women, remembering the highest dowry, and then select the first one after those 37% who exceeds the previous highest value. This 37% rule, to repeat, finds the highest value more often than any other algorithm (again, 37% of the time), and thus is, in this sense, the optimal solution to this problem. The advisor has slightly better than a 1 in 3 shot at picking the right woman and keeping his head. The other two-thirds of the time, well, the sultan has to look for another wise man.

2.3 Looking for better search rules

Of course, in the real world, our mating decisions are seldom this dramatic—we usually get to (or have to) live with whatever choice we make, even if it was not the “best” one. To the sultan’s advisor, the performance of the 37% rule on those occasions when it did not pick the highest dowry did not matter—he was killed in any case. But to a population of individuals all using such an algorithm to choose their mates, what this rule does the other 63% of the time would matter rather a lot. For instance, if applied to a set of 100 dowries ranging from 1 to 100, the 37% rule returns an average value of about 82 (that is, the mean of all dowries chosen by this rule). Only 67% of the individuals selected by this rule lie in the top 10% of the population, while 8% fall in the bottom 25%. And it takes the 37% rule an average of 74 tests of potential mates (that is, double the 37 that must be checked before selection can begin) before a mate is chosen. (These figures are all discussed in the next section.) If any of these performance figures could be improved upon by some other sequential choice algorithm, that algorithm could well prove more adaptive for a population of mate-choosers, allowing them to pick better mates more often, or more quickly, or with a smaller chance of picking a total loser, and we might therefore reasonably expect it to evolve in preference to the 37% rule.

To the extent that humans use any such sequential search algorithm when choosing mates, the more adaptive ones (if they exist) would make a better yardstick against which to compare human behavior. The 37% rule has at least two disadvantages as a model of human behavior: First, it requires knowing how many potential mates, N , there are to check through, in order to calculate how many are in the first 37% to look at before making a final choice. In our ancestral environments, N may not have varied much—our hunter-gatherer tribe sizes may have been fairly constant, so that an evolved value of N could be used in this algorithm. But in modern metropolitan environments, the number of potential mates has increased enormously, and to work optimally, the 37% rule would have to take those new large values of N into consideration. Second, this rule requires checking through a large number of individuals before a decision can be made—74 out of 100 in the previous example. Even assuming a rather quick assessment of someone’s mate potential, perhaps a few dates over a month’s time, the search time involved becomes extensive.

Thus, using the 37% rule for human mate search may require information that is difficult to obtain (an accurate value for N), and a large number of individuals to be checked and consequently a long search time. On the other hand, Frey and Eichenberger (1996) argue that one of the paradoxes of marriage is that people search too little for their marriage partners, checking too few individuals before making a life-long commitment to one of them. The evidence they cite argues against the presence of the 37% rule in human mate search—but it also argues that, by not searching long enough, people

are making suboptimal mate choices. If people are not using an algorithm as long-winded as the 37% rule, what might they be doing instead? And is it possible that there are any faster search rules that still perform well, so that Frey and Eichenberger's fears of suboptimal mate choice can be countered? If so, will these rules prove more complicated? In the next section, we explore the answers to these questions, and discover that we can in fact do more, in mate choice, with less.

3 The performance of simpler sequential mate search algorithms

Optimal rules have been derived for the problem of sequential search in a variety of settings that are more or less relevant to mate search. For instance, Gianini-Pettitt (1979) analyzed the situation in which the number of potential mates to be searched through is unknown. Corbin (1980) came up with methods for specifying search processes when the searcher can be satisfied with outcomes that are not necessarily ranked as number one (that is, where getting the second best, or third best, etc., all have some non-zero utility as well). She also considered the case where past candidates might still be available for the searcher to return to, with some probability—in this situation, it is sometimes best to search through the whole set of potential candidates and try to go back to the best one once the end of the sequence has been reached. Gilbert and Mosteller (1966) looked at situations where the distribution of candidate values is known or unknown, and at what happens when the searcher can select several candidates (i.e., be a polygamist, and try to get the best mate among the set of candidates chosen). And Quine and Law (1996) found optimal algorithms for maximizing the rank of a selected candidate (as opposed to maximizing the probability of selecting the absolute top candidate).

However, these optimal solutions generally require longer search, or more knowledge, or greater computation, than even the basic 37% rule. For instance, Quine and Law's (1996) algorithm for maximizing the selected rank has the searcher first check some number of individuals, and then take the next best one after up to some next point in the sequence, and then select any candidate in the top *two* relative-rank positions up to another later point in the sequence, and then select any candidate in the top *three* positions, and so on, successively lowering one's standards as the search progresses. While this makes intuitive sense (and works quite well), it involves a lengthy search (on average) and the non-trivial computation of multiple criterion-shifting points, all of which makes this algorithm ultimately less psychologically plausible. Instead, we wanted to start with very simple search rules which most people would agree *could* be psychologically plausible (or at least could be easily followed by an organism), and see how well—or how far below optimum—these

simple rules could also do in the search task. If they do well, we will have a simpler baseline against which we can compare human behavior (and a candidate set of rules to look for evidence of in human behavior), and if not, then we can begin to enhance them, adding slowly to their complexity until they do perform adaptively.

To investigate whether or not any simple sequential search rules exist that can outperform the standard 37% rule in the standard “secretary problem” search domain in various ways, we began by studying a class of rules derived from the 37% rule. It turned out that even this small set of rules contained some which are better than the 37% rule on many dimensions, and so we restrict our discussion here to this type of rule (with the understanding that other types of search algorithms or rules, which we are currently exploring, may prove to have even better performance). We have dubbed the class of rules we consider here the *Take the Next Best* (TNB) mate-choice rule family (after the *Take The Best* or TTB decision algorithm of Gigerenzer and Goldstein, 1996, which is a prime example of a simple satisficing mental mechanism).

TNB rules work in direct analogy to the 37% rule as follows: for some specified C , the first $C\%$ of the N total potential mates are checked (without being selected), and the highest dowry D is remembered. After the first $C\%$ of potential mates have gone by, the next potential mate with a dowry greater than D is the one chosen. (If no greater dowry turns up, then we assume that the searcher accepts the very last individual in the sequence—this is why our performance curves in the figures to come do not end up going to zero.) This simple sort of algorithm (of which the 37% rule is one specific example) has minimal cognitive requirements: it only uses memory for one value at a time (the current highest dowry), only needs to know N and C and calculate $N \times C / 100$, and only needs to be able to compare two dowry values at a time. What we want to know is how the performance of these simple algorithms changes as we change the percent of potential mates checked, C . Because we also wanted to be able to change the underlying assumptions of this problem, such as the distribution of dowry values, or the cost of checking each potential mate, or whether or not N is even known, the mathematics quickly grew complicated, and we decided on a much more flexible simulation approach to these questions.

We tested the behavior of TNB search algorithms with values of C from 0% (corresponding to always choosing the first potential mate) to 15% in increments of 1%, from 20% to 50% in increments of 5% (except around 37% where we again increased the resolution), and from 60% to 90% in increments of 10% (because we believed more of the action—good performance—would occur in the lower C ranges). We ran each rule in a variety of “mate environments,” or different numbers N and distributions of dowry values. For each such mate environment, we ran each rule against 10,000 different randomly-created dowry (or mate value) lists. We collected statistics on the distribution

of mate values selected by each algorithm (including the mean, standard deviation, quartile distributions, and number of times the single best dowry value was chosen) and positions at which mates were selected (the mean and standard deviation). With these values in hand, we can answer the questions posed at the end of the previous section: can another sequential mate search algorithm beat the 37% rule?

3.1 Performance of TNB rules with 100 potential mates

The answer, even from the class of simple-minded TNB rules, is a resounding “yes.” It is true, the 37% rule yields the highest mate value most often when the distribution of mate values is created without replacement (i.e. no duplicate values). But when we start relaxing the mate choice criterion slightly, as we do in Figure 1, this rule begins not to look so good. In the figure, the “Top 1” line shows the number of times out of 10,000 that the highest mate value was picked by a TNB algorithm, for different numbers of “dates” (potential mates) checked (i.e. different values of C —here C is out of 100, so it is equivalent to percentage as well). (See Gilbert & Mosteller, 1966, Figure 1, p. 42, for the mathematically-derived equivalent of this function.) This curve exhibits a classic flat maximum, so that it does not much matter what exact value of C is used—the results are largely the same for C between 20% and 40%. The greatest chance of choosing the highest mate value or dowry comes with a C of 37%, as expected, but this highest-value mate is only found in 37% of the cases. And things get worse from here.

An animal searching for a mate is probably not, as we mentioned earlier, only willing to settle for the “best” member of the opposite sex—other “pretty good” potential mates will often be selected instead, to save search time or energy (or even because the animal cannot perceptually distinguish between “best” and “pretty good”). It may suffice, in terms of having an adaptive advantage over other competing mate-seekers, to find a potential mate with a mate value in the top 10% of the population relatively quickly. In Figure 1, we see that the value of C that returns the most chosen mates in the top 10% of the value distribution is only 12%, and checking the first 12% of the sequence of mates before making a final choice results finding a top-10% mate over 77% of the time. If one’s standards are a bit more lax, and it is only desired to obtain a mate in the highest quartile (top 25%) with the greatest chance, then only $C=8%$ of the initial stream of potential mates need be checked, yielding mates in that top quartile over 90% of the time. Finally, rather than being risk-seeking in searching for a mate in the top ranks of the population, an animal may be risk-averse, preferring only to minimize its chances of landing a mate with a value in the bottom quartile of the population, where the mutants lie. From the line marked “Bottom 25%” we can see that the way to achieve this goal is to use a much lower C of 2%, yielding only 1% of chosen mates being from the bottom (quarter) of the barrel. The 37% rule would pick these poor mates over 11% of the




Figure 1: Chance of finding a mate in a particular category, given different number of mates checked first before taking the next best, out of 100 total possible mates.

time—much worse performance by risk-averse standards.

If instead an animal would gain the most adaptive advantage over its competitors by simply maximizing its average selected mate value over multiple uses of its TNB sequential search algorithm, Figure 2 shows how to accomplish this goal. By using $C=8\%$ in this environment of 100 potential mates, over time the searcher will select mates with an average mate value of 89. In contrast, if the searcher were to use the 37% rule, they would only average 74 (which is, in fact, the same performance level as if they checked only 1% of the population before taking the next best candidate!).

The mate value selected by these search algorithms may not be the only criterion that matters to a mating animal—the time and energy spent searching may also strongly influence the adaptiveness of the algorithm used (see, e.g., Pomiankowski, 1987; Sullivan, 1994). In Figure 3, we see how many total potential mates must be looked at, on average, before the final mate is chosen, varying as a function of the number of potential mates checked, C , before mate selection can begin. For the 37% rule operating in this environment, 74 potential mates must be looked at on average before a final mate is selected. With lower values of C , the number of mates that must be looked at falls off rapidly, with increasing advantage as C decreases. The optimal value of C according to this criterion is $C=0$, i.e., pick the first potential mate




Figure 2: Average value of selected mate (bars indicate one standard deviation), given different number of mates checked first before taking the next best, out of 100 total possible mates.

encountered. When combined with the other criteria, the exact cost assigned for each potential mate looked at will determine the precise tradeoff between finding a good mate and spending time and energy looking for that mate.

3.2 Performance of TNB rules with a greater number of potential mates

All of these criteria, beyond maximizing the chance of picking the single best mate, point to TNB rules with C other than 37%. Checking about 10% of the population of potential mates before selecting the next highest one to follow will result in just about the highest average mate value possible, most of the chosen mates in the top quartile and three quarters of them in the top 10%, and a search through 36 or so potential mates before the final selection is made. This seems like quite reasonable performance, especially because it means only checking 10 individuals initially out of the whole list of 100. But what happens if the population size is increased to 1000? Then checking 10% means testing 100 individuals, which starts to seem less like fun and more like hard work. If the number of individuals that must be tested by a TNB rule goes up linearly with the total population size (which it will do of course if C is a percentage, rather than a cardinal number), these rules may not end




Figure 3: Average position at which a mate is selected (bars show one standard deviation), given different number of mates checked first before taking the next best, out of 100 total possible mates.

up being so “fast and frugal,” at least for larger populations, after all.

But Figure 4, testing TNB rules in a population of 1000 potential mates with mate values from 0 to 1000 (with possible duplicates), shows that our fears of linear time increase are unwarranted. As before, the greatest chance of picking the single highest-value mate (or biggest dowry) comes from first checking 25-30% of the population. But to maximize the chances of picking a mate in the top 10% (with a 95% probability), only 3% of the potential mates need to be checked first to find D ; and for a mate in the top 25% (with a 98% probability), only 1% of the potential mates need be checked. Similarly, to minimize the chances (to 0.4%) of choosing a rotten mate in the bottom 25%, only 1% of population needs to be checked. And finally, to maximize the actual mean value of the selected mate, only 3% of the population should be checked, as shown in Figure 5.

Thus, to maximize potential mate value and minimize risk in this population of 1000 potential mates, somewhere between 1% and 3% of the population, or 10 and 30 individuals, should be first checked to come up with the target value for D . In the previous population of 100 individuals, also checking about 10 of them resulted in top search performance along these criteria. So despite the tenfold increase in population size, the number of individuals to check has increased only slightly. This suggests that our TNB




Figure 4: Chance of finding a mate in a particular category, given different number of mates checked first before taking the next best, out of 1000 total possible mates. Note that a smaller percentage of potential mates need now be checked to maximize the chances of getting a top mate.

rules can be simplified, so that instead of requiring that $C\%$ of the population is checked to come up with a value for D , now only C —a number independent of population size over a wide range, for instance $C=12$ (“Take a dozen”) for population sizes from 100 to 1000 or so—need be checked. Furthermore, this answers the criticisms of the 37% rule raised at the end of the last section: a simple sequential search rule that performs better than the 37% rule in several aspects does indeed exist, and it does not need knowledge of the total population size, nor require that an inordinate number of individuals be checked before a choice can be made. Finally, these results indicate that Frey and Eichenberger’s (1996) pessimism over short-searching humans ever finding their ideal mate may be unfounded: even a little bit of search may go a long way.

4 Further directions

Of course, we have left much out of this discussion of mate search. In focusing on the processes by which one individual can find his or her semi-ideal mate, we have neglected all aspects of mutual choice. Being able to have a high




Figure 5: Average value of selected mate (bars indicate one standard deviation), given different number of mates checked first before taking the next best, out of 1000 total possible mates. Note that a smaller percentage of potential mates need now be checked to maximize the average selected mate value.

probability of finding a good-quality mate quickly is certainly important; but in our species at least, that is usually only half the battle. We must also worry about being chosen by the other individual in return (as emphasized in the mate-choice-as-college-application scenario). This is the two-sided matching situation of Gale & Shapley, 1962, and Roth & Sotomayor, 1990, which must be combined with the limited-knowledge costly-search situation we have addressed here. We can (and will) examine the mutual choice behavior of TNB rules by incorporating them into a whole population of mate-searching individuals and recording the number of iterations it takes for everyone in the population to either get matched up or fail to find a mate. We will compare this behavior to that of an optimal matching rule like the Gale-Shapley algorithm in terms of population stability, number of unmated individuals, and time to settle. It may be that a simple, fast and frugal algorithm like TNB can outperform others along important dimensions such as these.

But this is another issue we must address: what *are* the most important dimensions over which search algorithms like these should be compared? Here we have argued that finding the absolute best individual in a population is not necessarily the most adaptive goal, if the search time, or mean mate

value, or distribution of mate values, can be improved upon. But we need to support this claim. One way to approach this problem is to create evolutionary simulations in which different algorithms compete with each other for mates and offspring, and see which types of algorithms win out over time. This approach, though, will only succeed in telling us something about real evolved human (or animal) behavior to the extent that we successfully incorporate the relevant ecological details (of how mate value maps onto number of offspring, for instance) into our model.

Another approach to this question is to study the actual evolved search behavior that humans and other animals use, as others have done in different settings (e.g. Hey, 1982, 1987; Harrison & McCabe, 1996; Martin & Moon, 1992; Rapoport & Tversky, 1970; Sethuraman, Cole, & Jain, 1994; Alatalo, Carlson, & Lundberg, 1988). We are interested in pursuing this line of evidence as well, but there is always the concern that experimental situations may not tap into the mental mechanisms used in real-world behavior. Alternatively, we can look for evidence of different search algorithms in the real observed mate search behavior of people and other animals; we are currently exploring the options of existing mate search databases.

But our main goal will continue to be to seek simple but effective possible mental algorithms for the mate search task, and test these algorithms through simulation in a variety of mate population environments. In this way, we hope to discover further “fast and frugal” algorithms, like the class of TNB rules we have presented here, that produce adaptive behavior with a minimum of information and computation. Finding a good-enough mate, as making a variety of choices, does not require optimal and expensive cognitive algorithms. Whether hoping for a ripe tomato, a willing college, or a vast dowry, it may not take much to get the next best thing.

Acknowledgments

Thanks to Ron Todd, Geoffrey Miller, David K. Smith, and the members of the Center for Adaptive Behavior and Cognition (ABC) in Munich for helpful comments on this work.

References

- Alatalo, R. V., Carlson, A., & Lundberg, A. (1988). The search cost in mate choice of the pied flycatcher. *Animal Behaviour*, *36*(1), 289-291.
- Andersson, M. (1994). *Sexual selection*. Princeton, NJ: Princeton University Press.
- Bateson, P.P.G. (Ed.) (1983). *Mate choice*. Cambridge, UK: Cambridge University Press.
- Becker, G. (1991). *A treatise on the family*. Cambridge, MA: Harvard University Press.
- Buss, D. M. (1994). *The evolution of desire: Strategies of human mating*. New York: Basic Books.
- Corbin, R. M. (1980). The secretary problem as a model of choice. *Journal of Mathematical Psychology*, *21*(1), 1-29.
- Cronin, H. (1991). *The ant and the peacock: Altruism and sexual selection from Darwin to today*. Cambridge, UK: Cambridge University Press.
- Darwin, C. (1871). *The descent of man, and selection in relation to sex* (2 vols.). London: John Murray.
- Ferguson, T.S. (1989). Who solved the secretary problem? *Statistical Science*, *4*, 282-296.
- Frey, B. S., & Eichenberger, R. (1996). Marriage paradoxes. *Rationality and Society*, *8*(2), 187-206.
- Gale, D., & Shapley, L. (1962). College admissions and the stability of marriage. *American Mathematical Monthly*, *69*, 9-15.
- Gianini-Pettitt, J. (1979). Optimal selection based on relative ranks with a random number of individuals. *Advances in Applied Probability*, *11*, 720-736.
- Gigerenzer, G., & Goldstein, D. (1996). Reasoning the fast and frugal way: Models of bounded rationality. *Psychological Review*, *103*(4), 650-669.
- Gilbert, J.P., and Mosteller, F. (1966). Recognizing the maximum of a sequence. *American Statistical Association Journal*, *61*, 35-73.
- Harrison, G.W., & McCabe, K.A. (1996). Stability and preference distortion in resource matching: An experimental study of the marriage problem. In R.M. Isaac (Ed.), *Research in experimental economics*, vol. 8. Greenwich: JAI Press.
- Hey, J. D. (1982). Search for rules for search. *Journal of Economic Behavior and Organization*, *3*, 65-81.
- Hey, J. D. (1987). Still searching. *Journal of Economic Behavior and Organization*, *8*, 137-144.
- Martin, A., & Moon, P. (1992). Purchasing decisions, partial knowledge, and economic search: Experimental and simulation evidence. *Journal of Behavioral Decision Making*, *5*(4), 253-266.
- Miller, G. F. (in press). Sexual selection in human evolution: Review and prospects. To appear in C. Crawford & D. Krebs (Eds.), *Evolution and human behavior: Ideas, issues, and applications*. Mahwah, NJ: Lawrence Erlbaum.

- Miller, G. F., & Todd, P. M. (1995). The role of mate choice in biocomputation: Sexual selection as a process of search, optimization, and diversification. In W. Banzhaf & F. H. Eeckman (Eds.), *Evolution and biocomputation: Computational models of evolution*. (Lecture notes in computer science 899, pp. 169-204.) Berlin: Springer-Verlag.
- Mosteller, F. (1987). *Fifty challenging problems in probability with solutions*. New York: Dover. (Reprint of the 1965 Addison-Wesley edition.)
- Pomiankowski, A. (1987). The costs of choice in sexual selection. *Journal of Theoretical Biology*, *128*, 195-218.
- Quine, M.P., and Law, J.S. (1996). Exact results for a secretary problem. *Journal of Applied Probability*, *33*, 630-639.
- Rapoport, A., & Tversky, A. (1970). Choice behaviour in an optimal stopping task. *Organizational Behaviour and Human Performance*, *5*, 105-120.
- Ridley, M. (1993). *The red queen: Sex and the evolution of human nature*. London: Viking.
- Roth, A. E., & Sotomayor, M. (1990). *Two-sided matching: A study in game-theoretic modeling and analysis*. Cambridge, UK: Cambridge University Press.
- Sethuraman, R., Cole, C., & Jain, D. (1994). Analyzing the effect of information format and task on cutoff search strategies. *Journal of Consumer Psychology* *3(2)*, 103-136.
- Sullivan, M. S. (1994). Mate choice as an information gathering process under time constraint: Implications for behaviour and signal design. *Animal Behaviour*, *47(1)*, 141-151.
- Todd, P. M., & Miller, G. F. (1993). Parental guidance suggested: How parental imprinting evolves through sexual selection as an adaptive learning mechanism. *Adaptive Behavior*, *2(1)*, 5-47.
- Todd, P. M., & Miller, G. F. (in press). Biodiversity through sexual selection. To appear in C. Langton (Ed.), *Artificial Life V*. Cambridge, MA: MIT Press/Bradford Books.