13

From Pride and Prejudice to Persuasion

Satisficing in Mate Search

Peter M. Todd
Geoffrey F. Miller

Wedding is destiny,
And hanging likewise.

John Heywood, Proverbes

I married the first man I ever kissed. When I tell this to my children they just about throw up.

Barbara Bush, First Lady

In 1611, the first wife of astronomer Johannes Kepler (1571–1630) died of cholera in Prague. Liberated from an arranged and unhappy marriage, Kepler immediately began a methodical search for a replacement. Though short, unhealthy, and the son of a poor mercenary, Kepler had an MA in theology from Tübingen, succeeded Tycho Brahe as imperial mathematician of the Holy Roman empire, and had recently become famous for explaining how eyeglasses can correct myopia (*Ad Vitellionem Paralipomena*, 1604), documenting a supernova (*De Stella Nova*, 1606), and demonstrating that the orbit of Mars is an ellipse (*Astronomia Nova*, 1609). He was a good catch. Relentlessly courting, Kepler investigated 11 possible replacements in the two years after his wife's death. In a letter to Baron Strahlendorf written shortly after marrying candidate number five in 1613, Kepler described this methodical mate search. Friends urged Kepler to choose candidate number four, a woman of high status and tempting dowry, but she rejected him for having toyed with her too long, so Kepler was free to settle with his most-preferred number five. Kepler chose well: His new wife, though not of the highest rank or dowry, was well-educated, bore him seven children, and provided the domestic infrastructure for Kepler to publish four more major works laying the empirical foundations for Newton's law of gravity, and, incidentally, to save his mother from being burned at the stake as a witch in 1620 (see Ferguson, 1989; Koestler, 1960).

Kepler's experience illustrates some of the major themes in the literature on search strategies that has emerged over the past several decades in statistics, economics, and biology. So far in this book we have focused on decision problems where all of the alternatives are simultaneously presented, and one only needs to search through information to guide one's choice. In many real-world choice problems, though, an agent encounters options in a temporal sequence, appearing in random order, drawn from a population with parameters that are only partially known ahead of time. In this case, the search for possible options, rather than just for information about present alternatives, becomes central. For mate choice in particular, the structure of the search task requires that one choose a prospect that fits one's real criteria for success rather than irrelevant ideals suggested by well-meaning acquaintances, given limited time for investigating each possibility, and some risk that the prospect, being an agent in his or her own right, will reject one's offer of union. The three disciplines that have investigated search tasks most thoroughly have emphasized different subsets of these issues.

Statisticians have focused on the "secretary problem" (Ferguson, 1989; Gilbert & Mosteller, 1966), in which one must pick the very best secretary from a sequence of applicants that appear in random order drawn from an unknown distribution of quality. Once rejected, applicants cannot be recalled. The secretary problem seeks perfection, with a payoff of one for picking the very best applicant and zero for picking anyone else. It also ignores search costs such as time, ignores the problem of mutual choice (the possibility that the applicant you like will not like you), and assumes you know the exact number of applicants who will arrive. But it directly addresses what to do about the uncertainty that the next prospect one encounters might be far superior to the best seen so far. It can be shown that the solution to the secretary problem demands sampling a certain proportion of the applicants, remembering the best of them, and then picking the next applicant who is even better. The optimal number to sample is $1/e$ (37%). Following this "37% rule" finds the very best applicant about 37% of the time (see Ferguson, 1989).

Economists have developed models of job search for best salaries and consumer search for lowest prices, emphasizing the importance of search costs and the acceptability of less-than-perfect options (Lippman & McCall, 1976). These search models, like those of the statisticians, usually ignore mutual choice, but they do not assume the total number of prospects is known, nor do they assume that only the best will do. On the other hand, these models generally assume that you can backtrack to pick previously seen options. With assumptions differing markedly from those of the secretary problem, the solution is also quite different. The general solution to this type of search task is to set a "reservation price" at which the marginal cost of further search equals the marginal expected improvement over the best prospect seen so far. That is, one should keep looking for a better salary or a lower price until the effort of looking further is

likely to be more costly than the amount of improvement you could achieve, and then return to the best seen. This in turn can depend critically on the standard deviation of the distribution of salaries or prices, which may need to be estimated from previously observed options (Martin & Moon, 1992). This type of solution, requiring involved computations to determine when to stop search, falls into the class of constrained optimization methods discussed in chapter 1.

Biologists spent considerable effort in the 1980s amassing support for Darwin's (1871) claim that animals engage in mate search with enough discrimination and persistence to impose sexual selection pressures on one another (see Andersson, 1994). Several researchers developed detailed models for search behavior (e.g., Johnstone, 1997; Wiegmann et al., 1996), often with less theorem-proving zeal than the statisticians or economists, but more attention to the empirical testability of their models. These models usually incorporate search costs, and sometimes lack of knowledge about the distribution of potential mates (Mazalov et al., 1996). Much recent effort has gone into distinguishing whether different species use a best-of-*N* rule or a threshold criterion rule in mate search (e.g., Fiske & Kalas, 1995; Forsgren, 1997; Valone et al., 1996). The best-of-*N* rule means sampling a certain number *N* of prospects and then choosing the best of those seen, whereas a threshold criterion rule, like the 37% rule and the reservation price rule, means setting an aspiration level and picking the first prospect that exceeds it. Simon (1990) has termed the latter aspiration-setting mechanism "satisficing," defined as "using experience to construct an expectation of how good a solution we might reasonably achieve, and halting search as soon as a solution is reached that meets the expectation" (p. 9). As indicated in chapter 1, Simon sees satisficing as one of the main forms of bounded rationality available in situations where the complete set of possible alternatives to choose from is not, or cannot be, known.

All of the above approaches tend to consider a single searcher assessing passive goods waiting to be chosen. But one of the major problems in mate search is coping with mutual choice. It is fairly easy to develop satisficing rules that work well for nonmutual search, for instance, shopping around for tomatoes or televisions that will not object to being bought. There has been much less research on finding satisficing rules for mutual search under uncertainty. One more literature is relevant in this regard: the tradition of economic game theory research on "two-sided matching" (Roth & Sotomayor, 1990), which is largely the study of mutual choice, but with certainty and complete knowledge.

As with most game-theoretic analysis, this tradition has focused on finding equilibria, or sets of strategies that are mutually optimal against one another. It can be shown that if a finite set of men and women have consistent, transitive preferences for one another, then there exists at least one "stable matching" in which no one who has a mate would prefer somebody else who would also prefer him or her in return. The two-sided matching literature also shows, however, that there are often multiple equilibria, or different possible stable matchings given a particular set of men and women with particular preferences. Although each is stable in the sense that there is no rational incentive for divorce and remarriage, different equilibria fulfill people's preferences to different degrees: Some are "male-optimal" (making men as happy as they could be given the actual preferences of women), some are "female-optimal" (making women as happy as they could be given men's preferences), and some are neither. There is a simple search method called the "deferred acceptance procedure" that is guaranteed to produce a stable matching efficiently given mutual choice and perfect and complete information about everyone in the population (Gale & Shapley, 1962). But whether such equilibria exist (or ever occur) for real populations, and whether any algorithms exist for finding them in realistic situations of imperfect, incomplete information, remains to be shown.

How do all these statistical, economic, and biological models illuminate Kepler's courtship plan, or more generally, human choice behavior when presented with a sequence of options? Mate search can be considered a rather difficult but extremely important type of decision making under uncertainty. The models mentioned above have identified some of the difficulties: uncertainty about the distribution of mate values one will encounter, ignorance of the order in which prospects will be met, difficulty of backtracking to previously rejected prospects, search costs, time limits and temporal discounting, and, above all, the mutual choice problem that mating must be mutually acceptable to both parties.

Different fields address or ignore these difficulties in different ways. Statisticians and economists tend to treat mate search as an interesting pretext for developing optimality theorems relevant to job search and consumer search, rather than treating mate search as a central adaptive problem in human life. Biologists view things differently, because mate search and mate choice drive sexual selection, an evolutionary process perhaps equal to natural selection in its power and creativity. With the resurgence of interest in sexual selection theory since the late 1970s (see Andersson, 1994; Bateson, 1983a; Cronin, 1991), and evolutionary psychology since the late 1980s (see Barkow et al., 1992; Buss, 1994; Miller & Todd, 1998), research has begun to focus on the role that sexual selection via mate choice has played in shaping many aspects of the human mind (Miller, 1998; Ridley, 1993; Wright, 1994). In studying mate search then, we are studying an interesting, difficult problem of decision making under uncertainty that, perhaps uniquely among such problems, is likely to have had a strong causal influence on human evolution.

As with so many problems of human decision making, the rationality and efficiency of human mate choice, including the process of search, has been questioned. Frey and Eichenberger (1996) argued that people do not search long enough when seeking a mate, taking the incidence of divorce and marital misery as indicators of insufficient search. Rapoport and Tversky (1970) questioned whether people adhere to the reservation price rule for searching given a known distribution of values and a known search cost. However, the sequential search literature is not dominated by these sorts of worries about the ways that people deviate from known optimal strategies, in part because the optimal strategies are not known for many realistic search situations, and in part because psychologists have paid much less attention to search tasks than to other decision-making tasks. Psychologists and economists who have studied search have often focused on the simple heuristics that people actually appear to use. Hey (1982, 1987) has identified a number of these rules, such as the "one-bounce rule," by which people seeking high values keep checking values as long as they increase, but stop as soon as they decrease and select the previous value. Martin and Moon (1992) used computer simulation to assess the relative performance of different simple strategies in a consumer search setting and found that some rules can come within 3% of the normative standard.

In this chapter, we follow in the footsteps of these researchers and look for simple satisficing search heuristics that perform adaptively in the specific domain of biologically realistic mate search problems. We also evaluate different heuristics in simulation across a variety of search conditions using a variety of performance measures. Following a similar historical trend in sexual selection theory in biology (see Cronin, 1991), we begin with the rather male-centered case of one-sided search, and then proceed to the more realistic case of mutual search, emphasizing female choice as well as male. (To keep our analyses simple, in this chapter we do not go into the effects of possible sex differences in mate search strategies, though these could certainly have interesting and important consequences.)

Through our analyses we find that, even for simple cost-free, nonmutual search as in the secretary problem, the 37% rule is outperformed on many criteria by heuristics that sample significantly fewer prospects. These heuristics do not even need to know the expected number of prospects one will encounter: A simple satisficing heuristic called "Try a Dozen" works well across a large range of numbers of prospects. We also find that when mutual choice enters the picture, these types of search strategies tend to perform very poorly. Only individuals who are very highly valued themselves can get away with applying the 37% rule or the Try a Dozen heuristic in mutual choice situations. (Kepler was lucky in this respect: His high mate value helped ensure that his "Try Eleven" strategy would yield good results.)

Instead, search heuristics that take into account one's own mate value perform much better in mutual choice, producing faster, more frequent, higher-quality matings for individuals. Even if one's own mate value is not known initially, good search efficiency can be attained using a simple adaptive heuristic that adjusts one's aspiration level based on the number of offers and rejections received from others during an initial sampling period. If one also pays attention to the mate values of those who do or do not show interest, it becomes easier to learn one's mate value, which can be used as a basis for effective search strategies that deliver close to the best mate that could be hoped for given mutual choice. In brief, search strategies such as the 37% rule and the Try a Dozen heuristic that work well without mutual choice perform extremely poorly given mutual choice, falling far behind mutual choice strategies that allow one to learn one's own mate value from others' reactions. In keeping with the idea of ecological rationality, we find that the satisficing heuristics for mate search that do best in a given environmental situation—whether one-sided or mutual search—are those that exploit the structure of the information in their environments, relying solely on mate values in the former case, and on expressions of interest or disinterest in the latter.

Algorithms for One-Sided Mate Search: The Dowry Problem

The idealized versions of search described in the previous section differ considerably 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 to settle down and have children with a particular person, though other definitions are possible.) There are certainly costs associated with checking out each person 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, as countless romantic tragedies attest). To further complicate matters, one does not know ahead of time what the range of potential mates may be: How can we know the first time we fall in love whether someone else might be able to incite still deeper feelings 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 we are searching, finding a mate looks like a very daunting problem indeed.

We can consider this situation in more precise detail, and in a form more closely linked to mate choice, via an alter ego of the secretary problem mentioned in the previous section: the "dowry problem." This is a well-known puzzle from statistics and probability theory (Corbin, 1980; see also Gilbert & Mosteller, 1966; Mosteller, 1987), as the number of names it goes by attests (it is also known as the "beauty contest problem" and even "Googol"). In its dowry form, the story goes like this: A sultan wishes to test the wisdom of his chief advisor, to decide if the advisor should retain his 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, the chief executioner chops off his head, and worse, he remains single. 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?

As mentioned earlier, it turns out that the algorithm the advisor should use to guarantee the highest chance of choosing correctly is the 37% rule, which in this case would work as follows: He should look at the first 37 women (or, more generally, 37% of any population of candidates he faces), 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 Ferguson, 1989; Gilbert & Mosteller, 1966; Mosteller, 1987.) This 37% rule is the best the advisor can do—it 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. With this rule, the advisor has slightly better than a one in three shot at picking the right woman and keeping his head. The other two-thirds of the time, the sultan has to look for another advisor.

The dowry problem is certainly an unrealistic reflection of human mate choice in many respects—it only involves one-sided (rather than mutual) search, it reduces search to a single dimension instead of appreciating the many facets by which we judge one another (Miller & Todd, 1998), it denies any possibility of comparing candidates simultaneously or returning to those previously seen, and so forth. But it gives us a reasonable starting point for testing some specific mate search mechanisms in a setting with at least some domain-specific structure. And we can modify some of its assumptions in useful ways to help us get a better understanding of more appropriate search mechanisms, as we will now show.

One of the major differences between the dowry problem and the real world is that in the latter, of course, our mating decisions are seldom so dramatic—we usually get to (or have to) live with whatever choice we make, even if it is 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 a lot. For instance, if applied to a set of 100 dowries covering all integer values from 1 to 100, the 37% rule returns an average value of about 81 (i.e., 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 9% fall in the bottom 25%. And it takes the 37% rule an average of 74 tests of potential mates (i.e., 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.

If the dowry problem itself is unrealistic, the 37% rule solution also has many characteristics that could make it an implausible model of how people actually choose mates. Here we will focus on two difficulties. First, it requires knowing how many potential mates, N, will be available, in order to calculate how many are in the first 37% to check and set one's aspiration level, D. 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 lifelong commitment to one of them. The evidence they cite argues against the use of the 37% rule in human mate search—but it also argues that, by not searching long enough, people are making worse mate choices than they might. If people are not using an algorithm as long-winded as the 37% rule, what might they be doing instead? Is it possible that there are any faster search rules whose performance can assuage Frey and Eichenberger's fears of poor mate choice behavior? 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.

The Consequences of Searching Less

To investigate whether any simple search heuristics exist that can outperform the 37% rule on various criteria in the standard secretary/dowry problem domain, we began by studying a class of satisficing rules derived from the original 37% rule. It turned out that even this small set of similar heuristics contained some that are better than the 37% rule on many dimensions, and so we restrict our discussion here to this class (though other types of simple search algorithms or rules will probably prove to have even better performance on some criteria). We have dubbed the class of search heuristics we consider here "Take the Next Best" (or TNB, named after the fast and frugal Take The Best decision heuristic described in chapter 4).

Take the Next Best 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—this is the searcher's aspiration level. After the first C% of potential mates have gone by, the next potential mate with a dowry greater than D is chosen. (If no greater dowry turns up, then we assume that the searcher accepts the very last individual in the sequence, which is why our performance curves in the upcoming figures fall to a final nonzero value.) This simple heuristic (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. We were interested in how the performance of these simple algorithms would change as we altered the percentage of potential mates they checked, C. Because we also wanted to be able to change the underlying assumptions of this problem, such as the distribution of dowry values, the cost of checking each potential mate, and whether or not N is even known, the mathematics quickly grew complicated, and we decided instead on a flexible simulation approach for answering 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 the interesting 37% value, where we again increased the resolution), and from 60% to 90% in increments of 10% (because we believed most of the action—that is, good performance—would occur in the lower C ranges). We ran each rule with different numbers N of potential mates, each with 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: Simply put, can the 37% rule be beaten?

Search Performance With 100 Potential Mates

The answer, even from the class of simple TNB rules, is a resounding "yes." Of course, the 37% rule picks the highest mate value most often. In figure 13-1, the "best" line shows how often the highest mate value was

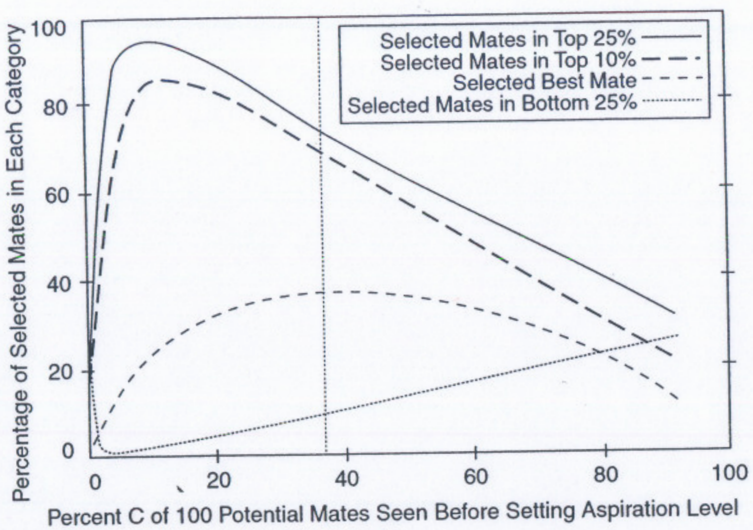

Figure 13-1: Chance of finding a mate in a particular value category, given different percentages of mates checked (out of 100 total possible mates) for setting the aspiration level before taking the next best candidate seen. The performance of the 37% rule on the various criteria is indicated by the broken vertical line.

picked by a TNB algorithm, for different percentages *C* of possibilities (potential mates) checked. (See Gilbert & Mosteller, 1966, figure 1, p. 42, for the mathematically derived equivalent of this function.) The greatest chance of choosing the highest mate value or dowry comes with a *C* of (about) 37%, as expected (the maximum in the figure is not at exactly 37%, because of the stochastic nature of the simulations we ran). But this curve also exhibits a 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 30% and 50%. And the chance of finding the highest-value mate for any of these strategies is never higher than 37%, as mentioned in the first section—not very good odds.

To an animal searching for a mate, this one in three chance of getting the "best" member of the opposite sex is probably not a bet worth taking—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"). In terms of having an adaptive advantage over other competing mate seekers, it may suffice to find a potential mate with a value in the top 10% of the population relatively quickly. In figure 13-1, we see that a low value of *C*, 14%, yields the highest chance, 83%, of selecting a mate in the top 10% of the value

distribution. If one's standards are a bit more lax, just desiring a mate in the highest quartile (top 25%), then only $C = 7\%$ of the initial stream of potential mates need be checked to maximize this chance, yielding mates in that top quartile over 92% of the time. Finally, rather than being risk-seeking by searching for a mate in the top ranks of the population, an animal may be risk-averse, preferring only to minimize its chances of picking a mate 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 3%, leading to a less than 1% chance of choosing a mate in the bottom (quarter) of the barrel. The 37% rule would pick these poor mates over 9% of the time, which is much worse performance by risk-averse standards.

Alternatively, an animal might gain the greatest adaptive advantage over its competitors by simply maximizing the expected value of its selected mate. Figure 13-2 indicates how to accomplish this goal, showing mean obtained mate value plotted against the percentage *C* of the potential mates that are checked to set the aspiration level *D*. Searchers using $C = 9\%$ in this environment of 100 potential mates with mate values from 1 to 100 will select mates with the highest average mate value, nearly 92. In contrast, if searchers were to use the 37% rule, their average would drop to 81.

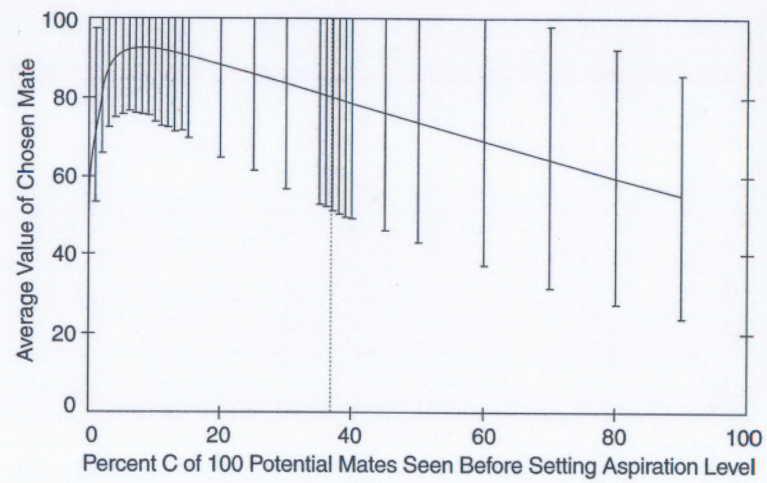

Figure 13-2: Average value of selected mate (bars indicate one standard deviation), given different percentages of mates checked (out of 100 total possible mates) for setting the aspiration level before taking the next best candidate seen. The performance of the 37% rule on this criterion is indicated by the broken vertical line.

The values of the mates selected by these search algorithms may not be the only criterion that matters to an organism seeking a mate—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 13-3, 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 ($C \times 100$) to set the aspiration level before mate selection. The 37% rule must look at 74 potential mates 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 alone would be $C = 0$, that is, pick the first potential mate encountered. When combined with the other criteria, the importance assigned to this mean search length variable will determine the precise trade-off between finding a good mate and spending time and energy looking for that mate.

Search Performance With a Greater Number of Potential Mates

All of the criteria other than the chance of picking the single best mate favor Take the Next Best rules that set their aspiration levels by looking at less than 37% of the population. Checking about 10% of the population

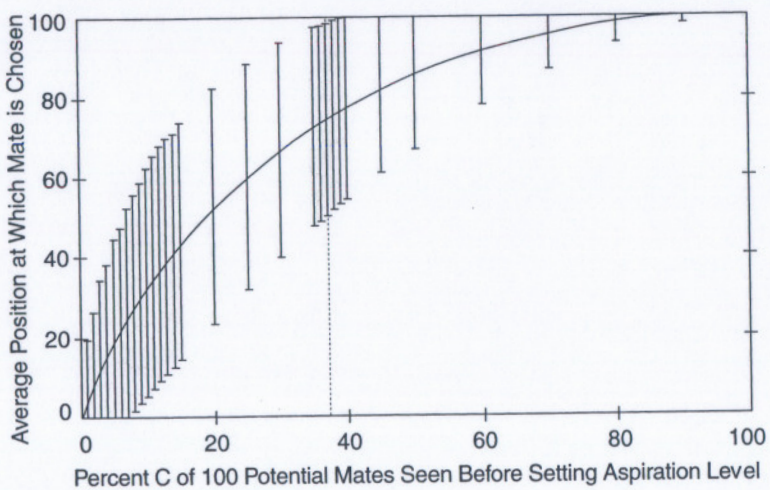

Figure 13-3: Average position at which a mate is selected (bars show one standard deviation), given different percentages of mates checked first (out of 100 total possible mates) for setting the aspiration level before taking the next best candidate seen. The performance of the 37% rule on this criterion is indicated by the broken vertical line.

of potential mates before selecting the highest individual thereafter will result in about the highest average mate value possible, along with a high chance of choosing mates in the top quartile and top 10%, and will require a search through 34 or so potential mates before the final selection is made. This seems like quite reasonable performance, given that it only requires checking 10 individuals initially out of a population of 100. But ancestral humans may have had effective mating group sizes much larger than this, and certainly in modern environments one can expect to meet more than 100 people who could potentially become mates. So what happens with our simple search heuristics if the population size is increased to 1,000, where checking 10% means testing 100 individuals, which may start to seem less like fun and more like hard work? Because the number of individuals that must be tested by a TNB rule with a *C*% parameter goes up linearly with the total population size *N*, these rules may not end up being so fast and frugal, at least for larger populations, after all.

But figure 13-4, which shows how TNB rules fare in a population of 1,000 potential mates with mate values from 1 to 1,000, proves that our fears of linear time increase are unwarranted. As before, the greatest chance of picking the single highest-value mate comes from first checking 37% of the population. But to maximize the chances of picking a mate in

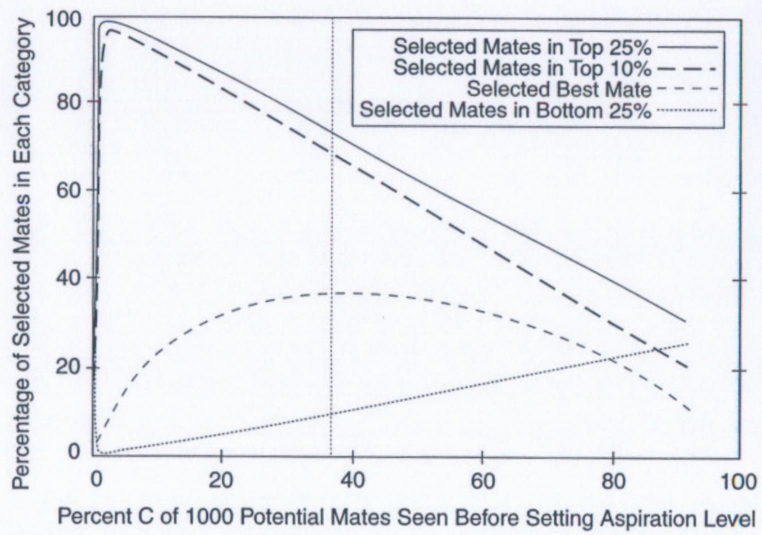

Figure 13-4: Chance of finding a mate in a particular value category, given different percentages of mates checked (out of 1,000 total possible mates) for setting the aspiration level before taking the next best candidate seen. Note that a smaller percentage of potential mates need now be checked to maximize the chances of getting a top mate. The performance of the 37% rule on the various criteria is indicated by the broken vertical line.

the top 10% (with a 97% probability), only 3% of the potential mates need to be checked to set the aspiration level D; and for a mate in the top 25% (with a 98% probability), only 1% to 2% of the potential mates need be checked. Similarly, to minimize the chances (to 0.3%) of choosing a mate in the bottom 25%, only 1% of the population needs to be checked.

Thus, to maximize potential mate value and minimize risk in this population of 1,000 potential mates, somewhere between 1% and 3% of the population, or 10 and 30 individuals, must be checked first to come up with the aspiration level D. In the previous population of 100 individuals, checking about 10 of them also resulted in top search performance judged by these criteria. So despite the tenfold increase in population size, the number of individuals to check increases only slightly. This suggests that our TNB rules can be simplified. Instead of checking a certain percentage of the potential mates to come up with an aspiration level D, we only need to check a certain absolute *number* of potential mates. This number will work for population sizes varying over a wide range—for instance, Try a Dozen ($C = 12$) is appropriate for population sizes from 100 to several thousand. This simplified search heuristic escapes the criticisms raised earlier against the 37% rule: It performs better than the 37% rule on multiple criteria, it does not need knowledge of the total population size, and it does not require checking an inordinate number of individuals before a choice can be made. These results indicate that Frey and Eichenberger's (1996) pessimism about short-searching humans ever finding an appropriate mate may be unfounded—even a little bit of search may go a long way.

On to Mutual Sequential Mate Search

That is, a little search can go a long way, if you are a despot who can force a collection of hapless potential mates to parade past you until you choose one. While we may start out with adolescent fantasies about getting the person we most desire, most of us soon discover that the mating game operates a bit differently. Imagine that you enter the game with your brand-new egocentric Try a Dozen rule, all set to find that high-value mate. You dutifully consider the first 12 people you randomly encounter, eventually turning each one down but remembering how much you liked the best. Starting with the 13th, you look at a succession of further possibilities until finally, on person 20, you find what you have been looking for: someone better than all the others you have already seen. Your rule is satisfied, and so are you. You propose to your newfound mate—and are summarily rejected. What went wrong?

The problem is, at the same time that you are evaluating prospective mates, they are evaluating you in return. If you do not meet a particular other person's standards, then no amount of proposing on your part is going to win them over (in this restricted scenario, at least). And if you and everyone else in the population have been using the Try a Dozen rule to form an aspiration level, then you and everyone else will have rather high aspirations for whom you will agree to mate with. The trouble is then that if you do not yourself *have* a high mate value, then you will not be selected by anyone else as a potential mate, and will end up alone.

We can observe these effects by constructing a new simulation to explore how different mate search rules will work in a mutual search situation. We create a population of 100 males and 100 females, each with a distinct mate value between 0.0 and 100.0, and each with accurate knowledge of the mate values of members of the opposite sex, but not necessarily knowing his or her own mate value. We give each of the 200 individuals the same search strategy, and first let them assess some specific number of members of the opposite sex during "adolescence." During this time, individuals can adjust their aspiration level, if their search rule uses one. After this adolescence period, males and females are paired up at random, at which point they can either make a proposal (an offer to mate) to their 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. This pairing-offering-mating cycle is repeated until every individual is mated, or until every individual has had the opportunity to assess and propose to every member of the opposite sex. We are interested in who gets paired up in this setting using different search rules; other criteria, such as how long this pairing process takes, are also of interest, but we will not discuss them here.

Figure 13-5 shows the number of mated pairs that will form in a population of 100 males and 100 females all using a particular mate-search strategy. If everyone uses Take the Next Best with $C = 1\%$, checking one individual to set their aspiration level, about half of the population will pair up. But as we increase the adolescence period (number of potential mates first checked), the number of mated pairs falls drastically. Thus if, instead, everyone uses the Try a Dozen variant and checks 12 potential mates for their aspiration level, only about eight mated pairs will be formed. The reason for this can be found in figure 13-6, where we can see the mean mate value of all mated individuals. For individuals using Take the Next Best rules, the longer the adolescence (number of mates to check, C), the higher the average mate value of all those who succeed in getting mated. That is, TNB rules give everyone in the population aspirations that are too high, so only the individuals who actually have the highest mate values will find mutually agreeing mates. Everyone else ends up spending Saturday night watching television.

But why not use TNB and check only a single individual? Then your aspiration level will not be too high, and nearly half of the population gets mated, which might be more reasonable. The problem lies in a third measure of population-level mating success: the average difference in mate value between partners in a mated pair. This is graphed in figure

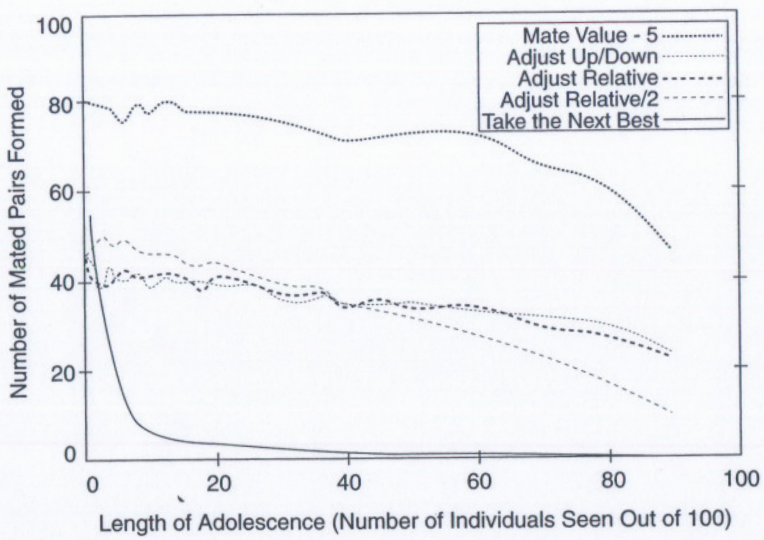

Figure 13-5: Number of mated pairs formed in a population with all individuals using a particular mutual sequential mate search strategy, graphed against the length of the adolescence period (during which an aspiration level can be learned). Higher values indicate more successful mate search strategies.

13-7. Here we can see that, even though TNB rules with very short adolescence periods do yield a good number of mated pairs, those pairs are rather mismatched—there is an average difference of nearly 25 between partners' mate values. Such a large difference would make the pairings formed very unstable in the game theory sense discussed in the first section: Many individuals would be inclined to switch partners. So how can we find a mutual sequential mate search rule that not only yields a high proportion of the population finding good mates (high values in figure 13-5), and finds mates for individuals from a wide and unbiased range of mate values themselves (values around 50 in figure 13-6), but also succeeds in pairing up individuals who are well matched to each other in terms of mate value (low values in figure 13-7)?

Now imagine that, considerably chastened by your earlier failure on the mating market, you reconcile yourself to be more realistic this time, and only aspire to a mate with a value similar to your own, rather than some lofty Hollywood-inspired ideal. In fact, out of humility you set a threshold five points *below* your own mate value, proposing to any individual with a mate value above this level. Now how will you fare, and how will everyone else do if they use similarly humble thresholds? In figure 13-5, we see that this strategy results in a high proportion of the

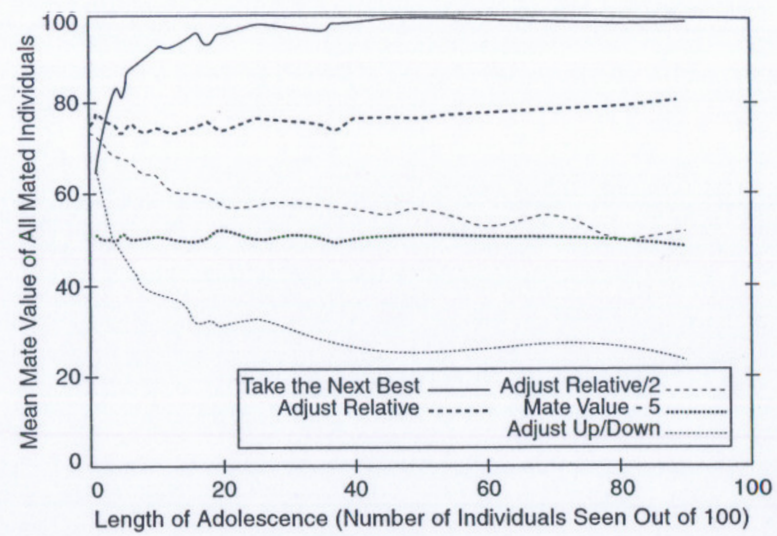

Figure 13-6: Mean mate value of all mated individuals in a population with all individuals using a particular mutual sequential mate search strategy, graphed against the length of the adolescence period. Middle values (around 50) indicate more successful egalitarian mate search strategies (for instance, those that enable more than just the elite to find mates).

population finding mates. In this case, adolescence does not involve learning or adjusting your aspiration level, because that value is fixed, but only represents an extended nonfertile period during which you meet people but cannot propose to them (and you still cannot go back to them later either). The length of adolescence has little effect on the performance of this humble mate search strategy. Only when adolescence gets very long does it start to reduce the number of mated pairs, simply because there is no longer enough of the population left to search through to ensure finding a good-enough partner.

This mate-value-based humble search strategy also does well on our other measures. Because most of the population gets paired up, the average mate value of those mated is around 50 (figure 13-6). It also succeeds in pairing individuals with very similar mate values (figure 13-7), making for a stable arrangement. That is, this strategy successfully sorts the population by mate value as it pairs the individuals. So this seems like a good mutual sequential mate search strategy to use. But there is a problem: Knowing one's own mate value is not necessarily an easy thing. We cannot be born with it, because it is both context sensitive (it depends on the others around us) and changes over time as we develop. We cannot simply observe ourselves to determine it, because we do not see ourselves in

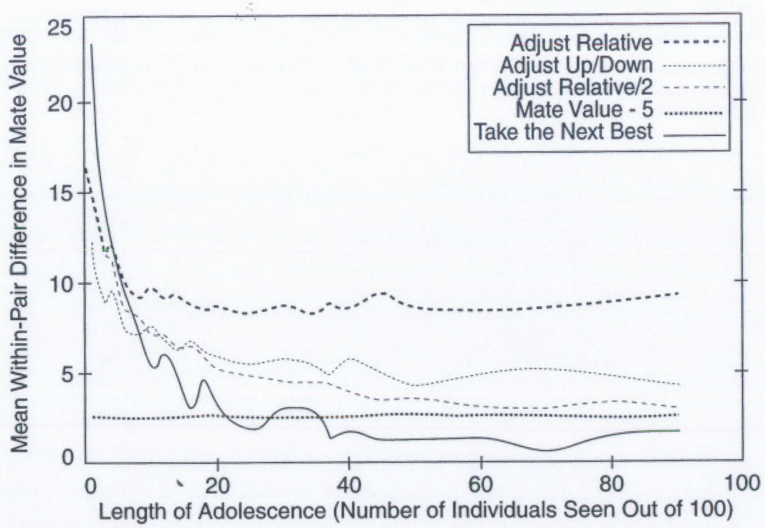

Figure 13-7: Mean difference between the mate values of partners in mated pairs formed in a population with all individuals using a particular mutual sequential mate search strategy, graphed against the length of the adolescence period. Lower values indicate mate search strategies that are more successful at forming well-matched pairs.

the same way that the others who judge us as potential mates see us. We do not even know the proper criteria on which to judge ourselves from the perspective of the opposite sex. Without this initial knowledge, then, we must somehow estimate our own mate value, if we are to use it to form our aspiration level.

Thus we must take another step toward making our mate search strategy less and less self-centered. We started by just considering what we thought of everyone else (Take the Next Best), then we used what we thought of ourselves (self-based aspiration level), and now we will look at what others think of us (adjusting our self-perception based on feedback). The first feedback-based method to try is to raise our aspiration level (the same as our self-perceived estimate of our own mate value) every time we get a proposal from someone else, and lower our aspiration level every time someone else does not propose to us. We will do this for a certain adolescence period again (i.e., use this feedback from a certain number of individuals we first encounter). The amount of adjustment we make to our aspiration level on each instance of feedback is inversely determined by the total length of our adolescence: If we have a short adolescence, we should make more adjustment (learn quickly) at each step, while if we have a long adolescence, we can learn more slowly. Thus, starting with an aspiration level of 50, we use *adjustment* = 50/(1 + *C*)

where *C* is the number of people we check out, and get checked out by, in adolescence. This rule pairs up about 40% of the population ("Adjust up/down" in figure 13-5), but preferentially in the lower half of the population (the mean mate value of mated individuals is about 25 in figure 13-6). What is happening here?

The problem with this aspiration-adjustment heuristic is that it is vain. Whenever a proposal comes from anyone, no matter what that person's mate value, the individual being proposed to gets excited and raises his or her aspiration level. Thus, individuals with mate values above 50 will get a lot of offers and then raise their aspirations to be too high, while those with mate values below 50 will more often get rejections, lower their aspirations, and as a consequence continue to boost the egos of the other half of the population. But individuals in the lower half, with the crushed aspirations, also succeed in finding mates, whereas those in the too-proud top half of the population often do not.

Instead of just taking someone else's word for it on whether you have a high mate value, you should also consider the source: What is the mate value of the other individual who is assessing you? If that person's mate value is higher than you think your own is, and he or she still proposes to you, then you should raise your own self-assessment, under the assumption that the other is well-calibrated and so is giving you accurate feedback about your own mate value. (You would always expect offers from individuals with mate values *lower* than your own self-image, so you should not use their offers to boost your self-image.) Similarly, if you get a refusal (lack of offer) from an individual with a mate value lower than your own self-perception, then this should make you think twice about your self-image, and lead you to lower your aspirations as well. (Again, lack of offers from those who are higher value than you think you are should not affect your self-image.)

If we make adjustments of the same size as the previous strategy, but now relative to the mate value of the other individual, we get about 40% of the population paired up again ("Adjust relative" in figure 13-5), but now it is the *top* half of the population that finds each other (mean paired mate value is about 75 in figure 13-6). However, they do not do a very good job of matching up—this strategy gives the worst mismatch between mated partners (about 10 points difference, in figure 13-7). The problem this time is that we are still using an adjustment that is independent of the mate values involved. The adjustment here is a fixed value depending only on the length of the adolescence period. But it does not make good sense to make the same upward adjustment both when someone with a mate value of 100 proposes to you, and when someone with a mate value of 60 proposes to you (assuming your self-image starts at 50, say). You should be much more excited about the former offer than the latter, and you should raise your self-estimate correspondingly higher. In the next strategy, we do just that.

As we have become less self-centered in our strategies, we have also added more information about the other potential mates we are interacting with. First, we looked at whether they proposed to us; next, we considered their proposals and the direction of their mate values relative to our own (i.e., we only needed to know if the values were bigger or smaller, not the exact values themselves); and now we consider their proposals, and the actual difference between their mate values and our self-estimate of our own mate value. If someone proposes to us whose mate value is higher than our self-image, then we raise our self-image (and hence our aspiration level) by half of the difference between the two. If a potential mate encountered during adolescence does not propose, and that person's mate value is lower than our self-image, then we lower our self-image by half of the difference. In this way, we put more weight on the feedback we get from individuals who are further away from our self-image.

When we do so, we end up with the best aspiration-learning strategy out of those we have considered so far. For short to medium adolescence lengths, this strategy pairs up about half of the population ("Adjust relative/2" in figure 13-5). The more learning it can do (i.e., the longer adolescence is), the closer it comes to pairing up an even distribution of individuals (figure 13-6). And with more learning, the mismatch between mated partners falls to near that of the humble mate-value-based strategy (figure 13-7).

Thus, we are getting close to a reasonable mutual sequential mate search strategy. It involves estimating one's own mate value by using the feedback of offers and refusals from members of the opposite sex, assuming that we know their mate values. But note that this kind of simple strategy does *not* assume that we know, or calculate, anything about the population as a whole. We do not have to keep track of means or standard deviations of the mate values encountered, for instance (as some of Martin & Moon's, 1992, strategies required). We also do not have to calculate optimal search times (as many of the approaches to the secretary/dowry problem required)—instead, most of the criteria seem to reach asymptote after checking about 20 individuals. And we do not need prior knowledge of the entire population, distinguishing this approach from that considered in two-sided matching problems (Roth & Sotomayor, 1990). Just seeing one individual after another, and learning about ourselves in the process, is enough.

Further Directions

We have presented here a collection of simple satisficing heuristics for one-sided and mutual search that can learn appropriate aspiration levels after checking only a few possible choices. As such, these heuristics fit into our overall framework of bounded rationality: They use as little of the available information as possible and still yield satisfactory performance. These rules are also ecologically rational, relying on the structure of information in the environment—here, the pattern of proposals and rejections made by members of the opposite sex—to bootstrap their adaptive choice behavior.

Of course, we have still left much out of this discussion of mate search. Populations are never fixed, and the mating game does not proceed in discrete periods during which everyone in a predetermined set must pair up or give up—rather, new individuals are always being introduced, which has an effect on the overall mating success of different strategies (Johnstone, 1997). The distribution of mate values we have used here is uniform, but in the real world it is probably closer to a normal distribution. How will different distributions of mate values affect the performance of different strategies? We have given everyone in the population precisely the same impression of all the members of the opposite sex (all females rank all males the same way, and vice versa), but this is not realistic either: There will typically be some degree of agreement about who is a good catch and who is not, but there will also be a large amount of idiosyncracy in individual mate preferences. Some of these individual preference differences will be based on purely aesthetic criteria, but some will also have important fitness consequences (such as preferences for mating with distant, but not close, relatives—see Bateson, 1983b).

This leads to another issue we must address: What *are* the most important dimensions over which search algorithms such as 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 chosen mate value, or distribution of chosen mate values, can be improved upon. Furthermore, finding a mate at all, in the mutual search case, could require selecting an individual with a mate value close to one's own. But we need to support these claims. 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.

The ultimate goal is to look for evidence of particular strategies in the actual evolved search behavior that humans and other animals use, as others have done experimentally in settings including mate choice (e.g., Alatalo et al., 1988; Harrison & McCabe, 1996; Hey, 1982, 1987; Martin & Moon, 1992; Rapoport & Tversky, 1970; Sethuraman et al., 1994). There is always the concern that experimental situations may not tap into the mental mechanisms used in real-world behavior, though, so it is also important to look for evidence of different search algorithms in the real observed mate search behavior of people and other animals. Our simulations are intended to guide these investigations of real behavior, by indicating what kinds of psychologically plausible, simple but effective search rules we can reasonably expect, and so should look for.

All of this is not to say that love has no place in mate choice, that it is all down to percentages and aspiration levels and adaptive self-assessments. Love can be a way of making any particular choice stick, lessening or erasing any perceived mismatch between partners and making further search seem blissfully unnecessary, even unthinkable. Love and other emotions are important parts of behavioral mechanisms, rather than unique undefinable forces that are orthogonal or even antagonistic to adaptive behavior. But love can—and indeed may be designed to—obscure the operation of the decision mechanisms in mate choice, so that the entire process seems unfathomable when one is caught up in it. Choosing a mate should not be a scientific affair. But we hope that scientific research *can* be used to reveal some of the patterns in behavior underlying the way that people search for, and find, each other.

14

Parental Investment by Simple Decision Rules

Jennifer Nerissa Davis
Peter M. Todd

There was an old woman who lived in a shoe. She had so many children she didn't know what to do.

Traditional Nursery Rhyme

The old woman may not have had much to give her children, but somehow she still had to figure out how to divide the broth that she did have among them. How could she do this? How do parents decide how to divide their time, money, and energy among their children? They could try to perform some sort of complex analysis, estimating all the future costs and benefits from now until their children become independent for all possible current choices, if such a thing were realistically calculable. Although the task of figuring out "optimal" solutions to these sorts of problems may be terribly complex or even impossible, given the large amount of computation and prediction of uncertain future events required, this does not mean parents must perform complex calculations to invest wisely. Instead they can rely on simple rules to guide their investment in their children. In this chapter we present results of a study designed to test just how successful such simple rules can be.

Parental Investment

Economists and behavioral ecologists have both addressed the problem of how parents should divide investment among their children. The models they have created, however, typically require information that is at best difficult to calculate and at worst actually unknowable. For example,