Available online at www.sciencedirect.com
SCIENCE DIRECT®
ELSEVIER
European Journal of Operational Research xxx (2005) xxx–xxx
EUROPEAN JOURNAL OF OPERATIONAL RESEARCH
www.elsevier.com/locate/ejor

How much information do we need?

Peter M. Todd *

Center for Adaptive Behavior and Cognition, Max Planck Institute for Human Development, Berlin 14195, Germany

Abstract

Modern technology is succeeding in delivering more information to people at ever faster rates. Under traditional views of rational decision making where individuals should evaluate and combine all available evidence, more information will yield better decisions. But our minds are designed to work in environments where information is often costly and difficult to obtain, leading us to use simple fast and frugal heuristics when making many decisions. These heuristics typically ignore most of the available information and rely on only a few important cues. Yet they make choices that are accurate in their appropriate application domains, achieving *ecological rationality* through their fit to particular information structures. This paper presents four classes of simple heuristics that use limited information—recognition-based heuristics, one-reason decision mechanisms, multiple-cue elimination strategies, and quick sequential search mechanisms—applied to environments from stock market investment to judging intentions of other organisms to choosing a mate. The findings that ecological rationality can be achieved with limited information are also used to indicate how our mind's design, relying on decision mechanisms tuned to specific environments, should be taken into account in our technology's design, creating environments that can enable better decisions.

Keywords: Heuristics; Fast and frugal decision making; Environment structure; Ecological rationality; Limited information

1. Introduction

Humans are a rather impatient lot, willing to make snap judgments and jump to conclusions on the basis of very little information. Even when more information *is* readily available, many decisions are made on the basis of quick impressions without bothering to gather further data. The same holds true when the opportunity arises to expand the set of possible choice options: People often avoid seeking additional alternatives and instead settle for one of the things that is already available. These failures to search for further information or alternatives occur at all levels of decision making, from the relatively inconsequential to the rather major. For instance, people choose products and buy stocks on the basis

* Tel.: +49 30 82 406 347; fax: +49 30 82 406 394.
E-mail address: ptodd@mpib-berlin.mpg.de

0377-2217/$ - see front matter © 2005 Published by Elsevier B.V.
doi:10.1016/j.ejor.2005.04.005

of name recognition alone (Goldstein and Gigerenzer, 1999, 2002), and people looking for information on the web tend to give up on a site after less than two clicks rather than searching it more deeply to see if it will be useful (Huberman et al., 1998). Of slightly more consequence, people on the market for apartments are happy to check only 6 of 25 cues made available to them before making a choice (Saad and Russo, 1996). Even more dramatically, it is commonly claimed that people fall in love at first sight without finding out more about their partner. While this may or may not be true, surveys in the United States have shown that more than a third of individuals in their 30s a decade ago married the first person they had sex with—a fact that economists have pointed to as evidence of human inadequacy in information search (Frey and Eichenberger, 1996).

Findings such as these usually lead to the claim that people are acting irrationally in such situations (Piattelli-Palmarini, 1996; Gigerenzer and Todd, 1999). The correct thing to do according to traditional norms of rationality is to collect all the available information and combine it appropriately, or to consider all the possible alternatives until the costs of doing so outweigh the potential benefits. Doing anything less than this would risk error and poor judgment. But is it actually so wrong to make decisions in the rapid manner humans frequently employ? How many options *should* we consider? How much information *do* we need?

The surprising answer being found in more and more recent decision-making research is this: for many situations, not that much. Instead of needing to process all the facts and consider all the options, people can often make surprisingly good decisions using simple "fast and frugal" heuristics, shortcut choice strategies that ignore a lot of information. The trick is to ignore the appropriate pieces of information, that is, the unnecessary bits. Or put the other way, the trick is to search for the few pieces of information or alternatives that will be most useful and process them appropriately. Simple heuristics are being uncovered that accomplish this trick in a variety of decision domains. This paper introduces some of these decision mechanisms in domains ranging from food choice to mate choice and shows how the study of heuristics can aid our understanding and practice of making good decisions—and making tools to help reach good decisions—with limited information.

2. Three views of human rationality

Modern technology is being used to deliver a broader range of information to people in a broader range of circumstances than ever before, often with the aim of helping people to make decisions or otherwise influencing their behavior. But how do people actually end up processing the information they are flooded with into decisions and actions? Without knowing this, we cannot say very conclusively how best to help decision makers in any particular context, nor what and how much information would best accomplish this goal. While it is obvious that we must take human psychology into account in figuring out what and how to communicate to people, there is often disagreement as to what is the nature of that psychology. It comes down to competing views of human rationality.

The traditional view of *unbounded* rationality says that decisions should be made by gathering and processing all available information, without concern for the human mind's computational speed or power. This view is found surprisingly commonly in perspectives ranging from Homo economicus in economics to the "GOFAI" (good old-fashioned AI) school of artificial intelligence (Goodie et al., 1999). According to this view, information technologies should either shower people with all the information that might possibly be relevant for making a particular decision and let them work out the optimal inference themselves, or they should gather as much information as possible and then make the decision for the user by weighing and adding it all into a final conclusion. This view of unbounded rationality at work can be seen in various World Wide Web decision aid sites, such as selectsmart.com, which gathers extensive data by asking users dozens of questions about their preferences for everything from what kind of pet they should get to what kind of industrial drill sharpener is best for them, and then processes all that information into a final ordered list of possibilities for the user to buy. This unbridled approach to information processing certainly fails to capture how most people make most decisions most of the time; as a consequence, it is not only a poor basis for building psychological models, but also a poor basis for building decision tools meant to be used or understood by real people (Katsikopoulos and Fasolo, submitted for publication).

In contrast, decision making via simple heuristics fits into the realm of *bounded* rationality—studying how people, and other animals, can make reasonable decisions given the constraints that they face, such as limited time, limited information, and limited computational abilities. Herbert Simon championed this view of cognition, arguing that because of the mind's limitations, humans "must use approximate methods to handle most tasks" (Simon, 1990, p. 6). These methods include recognition processes that largely obviate the need for further information, heuristics that guide the search for information or options when it is necessary and determine when it should end, and simple decision rules that make use of the information found—we shall see examples of each of these methods below.

Simon's notion of bounded rationality, originally developed in the 1950s, was enormously influential on psychologists and economists who followed, but it was interpreted in two distinct ways: First, a number of researchers accepted his assertion that the mind *does* work that way—but assumed at the same time that it is often flawed in doing so. We would, and should, all be unboundedly rational, if only we could. Under this view, the simple heuristics that we so often use can often lead us astray, making us reach biased decisions, commit fallacies of reason, and suffer from cognitive illusions (Piattelli-Palmarini, 1996). The very successful "heuristics-and-biases" research program of Tversky and Kahneman (1974; Kahneman et al., 1982) has embodied this interpretation of bounded rationality and led to much work on how to "debias" people so they could overcome their erroneous heuristic decision making.

In stark contrast, a growing number of researchers are finding that people can and often do make *good* decisions with simple rules or heuristics that use little information and process it in quick ways (Payne et al., 1993; Gigerenzer and Todd, 1999; Gigerenzer and Selten, 2001). This second view of bounded rationality argues that our cognitive limits do not stand in the way of adaptive decision making (though other environmental factors can, as we will turn to below); in fact, not only are these bounds not always hindrances, they can even be beneficial in various ways (Hertwig and Todd, 2003). The applied implication of this perspective is that if people can successfully use fast and frugal heuristics to process only a few pieces of information when making decisions, then striving to deliver them greater and greater amounts of information may not achieve the desired end of aiding good decisions—or at least not as cheaply and effectively as could otherwise be possible. Thus, this view of bounded human rationality prescribes figuring out what information people will actually use and focusing delivery on those items—a less-is-more, simplicity-based approach that is beginning to catch on in applications in medical communication (Hoffrage et al., 2000), law (Gigerenzer, 2002), technology industries (Kluth, 2004), business and marketing (Fasolo et al., in press), and elsewhere.

To figure out what information *is* actually important for people's decisions in different situations, we must consider the source of our bounded rationality (Todd, 2001). The usual assumption is that the constraints that bound our rationality are internal ones, such as limited memory and computational power. But this view leaves out most of the picture—namely, the external world and the constraints that it imposes on decision makers. There are two particularly important classes of constraints that stem from the nature of the world: First, because the external world is uncertain—we never face exactly the same situation twice—our mental mechanisms must be robust, that is, they must generalize well from old instances to new ones. One of the best ways to be robust is to be simple, for instance, by employing a mechanism containing few parameters. As a consequence, external uncertainty can impose a bound of simplicity on our mental mechanisms.

Second, because the world is competitive and time is money, or at least energy, our decision mechanisms must generally be fast. The more time we spend on a given decision, the less time we have available for other activities, and the less likely we are to outcompete our rivals in the endless arms race of life. To be fast, we must minimize the information or alternatives we search for in making our decisions. That is, the external world also constrains us to be frugal in what we search for.

But the external world does not just impose the bounds of simplicity, speed, and frugality on us—it also provides the means for staying within these bounds. A decision mechanism can stay simple and robust by relying on some of its work being done by the external world—that is, by counting on the presence of certain useful patterns of information in the environment. Some observable cues are useful indicators of particular aspects of the world, such as red color usually indicating ripe fruit. Our minds are built to exploit such patterns and thereby reduce the need for gathering and processing extra information. But, as the research in the heuristics-and-biases program has demonstrated, such reliance on particular expected information patterns can lead us astray if we are presented with environments that violate our expectations. We evolved in environments where bits of information were endowed with reliability by the more-or-less immutable physical objects they emanated from, things like a family member's face, or a predatory cat's snarl. In the modern digital world where images, sounds, and other sensations can be built up from scratch and delivered to us as purportedly useful information, the bits have been dissociated from the physical atoms, and the expected patterns of reliable relationships need no longer hold. This points again to the importance of giving decision makers the right information in the right presentation to facilitate their inferences and choices.

Emphasizing the role of the environment for bounding, constraining, and empowering human cognition leads to a new conception of *ecological* rationality (Todd et al., 2000). The goal in studying ecological rationality is to explore how simple mental mechanisms can yield good decisions by exploiting the structure inherent in the particular decision environment where they are used. This is the research program of Gigerenzer et al. (1999; Todd and Gigerenzer, 2000), who have so far focused primarily on laying out the contents of the "adaptive toolbox" of decision mechanisms that people use in a variety of task environments (Todd et al., 2000). From this foundation, one can then design environments that convey useful information in the way most appropriate for people to make good decisions. Two main types of simple heuristics in the adaptive toolbox have been explored to date: those that make decisions among currently available options or alternatives by limiting the amount of information they seek about the alternatives; and those that search for options themselves in a fast and frugal way. Both types rely on even simpler building blocks that guide the search for information or options, stop that search in a frugal manner, and then decide on the basis of the search's results. Next we will consider three examples of the first sort of information-searching decision heuristics, before finishing with heuristics for sequential search over alternatives.

3. The recognition heuristic—ignorance-based decision making

Within the realm of fast and frugal heuristics that seek and use only a limited amount of information, perhaps the simplest possible heuristic actually relies on a *lack* of knowledge. Consider the following inferential task: Which city is larger, San Diego or San Antonio? When Goldstein and Gigerenzer (1999, 2002) asked students at the University of Chicago this question, 62% of them got it right. Then they asked students at the University of Munich the same question and were surprised to find that 100% of the German students chose the correct answer (San Diego). Goldstein and Gigerenzer knew that this could not just be another example of Americans knowing less about the geography of their own country than do foreigners—the well-educated University of Chicago students knew more about each city than did the German students. Moreover, most of the German students had not even heard of San

Antonio. But this is apparently exactly what enabled them to do so well on this decision.

Goldstein and Gigerenzer found similar patterns on a wide variety of questions, and explained this choice behavior as the outcome of a very simple decision rule: the recognition heuristic. This heuristic says that for choosing between two objects on some criterion, when one object is recognized and the other is not, then pick the recognized one. Clearly, this heuristic will only work in some circumstances—specifically, when the environment is structured so that recognized objects are more often higher on the choice criterion. This is indeed the case for the city size question—larger cities are more likely to be talked about, or mentioned in the media, and hence recognized. In fact, the recognition rates of cities in Germany and the United States were more closely associated with the number of times they were mentioned in newspaper stories than with their actual populations. This is a reminder that the environment structure on which our decisions are based is often socially received rather than directly perceived.

Recognition rates are also correlated with size, or status, or importance, in a wide range of other domains beyond city size, from tall buildings to winning soccer teams. But the recognition heuristic can only be applied in any of these domains when some of the objects to be chosen between are *not* recognized. This is what tripped up the University of Chicago students—because they recognized both San Diego and San Antonio, they could not use this heuristic, and instead they had to rely on other knowledge they had about each city—knowledge that proved to be fallible more often than did the Germans' recognition knowledge alone. The German students in contrast were able to capitalize on their systematic ignorance. Goldstein and Gigerenzer (1999, 2002) showed more precisely that for a given decision task there is an intermediate amount of ignorance versus knowledge that yields the highest performance for making decisions with the recognition heuristic. If one knows more than this peak amount (e.g., recognizes more cities), decision performance can actually decrease. This leads to the unexpected less-is-more effect: Less knowledge can yield more accurate decisions when using the recognition heuristic. For instance, American students tested on size comparisons between pairs of the biggest American cities got 71% correct (median score), but when tested on pairs of the biggest German cities, where they knew much less, they scored 73% correct.

All of this shows that people can use recognition to answer questions about things like city size in the lab. And recent research has shown that people put considerable stock in the value of recognition information for making decisions, even being swayed more in a group decision setting by colleagues who only recognize one available option (and choose that option on the basis of their recognition) than by those who have more information and recognize all available options (Reimer and Katsikopoulos, in press). But do situations in which the recognition heuristic can be applied ever arise in daily life? Clearly many large corporations think they do. For instance, the clothing company Benetton bet almost exclusively on building name brand recognition to attract customers, rather than confusing people by offering any information whatsoever about their products in their ads—and this is a strategy that paid off (Goldstein and Gigerenzer, 2002). Based on this observation, Borges et al. (1999) investigated the use of the recognition heuristic for picking companies to invest in on the stock market. They found that portfolios based on companies recognized by laypeople interviewed on the street outperformed portfolios of expert-recognized stocks as well as expert-managed mutual funds. This result may itself only apply in certain environments, such as when the market is rising in general, and it points to the problem of knowing when to use recognition and when to seek more information.

It is not yet clear how people assess the validity of recognition knowledge in particular domains. While recognition is a powerful principle that can guide accurate decisions in a wide range of domains, it is not infallible, and its validity is growing easier to manipulate, for better or worse, with modern technology. In the ancestral environments in which mammalian and human minds evolved, recognition was difficult to fake: The basic way for an individual to get to recognize another individual or object or location was for the first individual to get the thing to be recognized into close proximity, close enough to allow direct perception and storage of the relevant bits in memory. Either you went to the object, or the object came to you, and both ways could be fairly costly in terms of time and energy. Similarly, until recently advertising campaigns relied on buying a lot of paper or billboard space and somehow putting this where people would see it in order to achieve expensive name recognition. But today only the bits themselves need to be transported—the atoms can stay put. For instance, you can now buy 150 million e-mail addresses for a small sum and quickly achieve widespread name recognition, though perhaps not in the positive direction. Television shows and movies trick us into thinking the faces we recognize belong to people we actually know (Kanazawa, 2002). And studies of the overnight fame effect (Jacoby et al., 1989) have shown that people can be manipulated into thinking some unknown person is famous just by seeing the person's name in a list of real celebrities and then recognizing that name the next day. How can we protect against such cheap manipulation of our recognition heuristic in the overly flexible modern information environment?

One way is to rely on socially amassed recognition, rather than merely individual experience. This is essentially what Borges et al. did in constructing the recognition-based stock portfolios, picking those companies that many people recognized to increase the reliability of the choices. The *Google* search engine does a similar thing, if you think of links between pages as indicators of recognition—the more links, the more recognition in this sense, and the more useful the site is judged to be. Anthropologists have studied related simple mechanisms of social learning—copying the most prevalent behavior, which can be accomplished through recognition mechanisms (Todd and Heuvelink, submitted for publication)—as means for the evolution and spread of cultural innovations (Henrich and McElreath, 2003). While not living in the digital world, rats also rely on social recognition cues. They develop a preference for foods that they recognize from having smelled them on the breath of a fellow nest mate, presumably under the assumption that if the other rat ate something and is still alive to exhale its aroma, then this substance must be okay to eat (Galef, 1987; Noble et al., 2001). Online, this social recognition principle has been incorporated into restaurant recommendation sites (such as the now-defunct Boston Eats)—if a student is looking for, say, a cheap local Chinese restaurant and finds a recommendation of one from another student who lived long enough to send in a review, then it is probably okay to eat there, too. Thus, relying on the information collected by others, even if it is only whether or not something is recognized, can help overcome the ease with which recognition-based choices can be dishonestly manipulated at the individual level. This is an aspect of ecological rationality that Gigerenzer (1996) and co-workers call *social rationality*, to highlight the importance of the fact that many of the decisions we make are based on information from social environments constructed of sets of other people we have contact with.

4. One-reason decision heuristics—taking the best cue

Of course, we often have more information than just recognition available for making our decisions. What kinds of fast and frugal heuristics are appropriate in situations like the following? Imagine trying to decide between two restaurants for taking a guest to dinner. The traditional and normatively prescribed method would be to collect all the information or cues that you know or could find out about each restaurant, such as the average meal cost, distance from home, and amount of garlic in the dishes; then weight each of these cues by their importance for this decision; and finally combine all the weighted values for each alternative to come up with a final total criterion value for each. Whichever restaurant has the higher final criterion value is the one to go to, according to this weighted-additive approach to computing the expected utility of the two choices (Edwards and Fasolo, 2001).

A simpler and faster method is the following: Consider a single cue for the two alternatives, such as meal cost. Does this cue distinguish between the restaurants? If it does, then stop and choose the restaurant pointed to by the cue (e.g., the cheaper one, or the more expensive one, depending on if you want to conserve your resources or impress the guest). If the first cue does not distinguish between the alternatives, then consider a second cue, such as distance. If that cue distinguishes, then stop at this point and go with the indicated choice (e.g., the nearer restaurant). If not, consider a third cue, and so on, stopping this search for cues at the first distinguishing one found and using that cue alone to make the final decision. Mechanisms that operate in this way are called "one-reason decision heuristics," because their final decision is made on the basis of a single cue or reason alone (Gigerenzer and Goldstein, 1999). All of the heuristics in this family have the same stopping and decision rule building blocks (stop after the first discriminating cue, and use that cue alone to make the decision), but they differ in terms of the cue-search building block. For instance, the Minimalist heuristic looks at the cues in a random order, while another heuristic called Take The Best looks at cues in order of their validity, that is, how often they point to the right choice.

This one-reason decision-making approach is certainly fast and simple—but can ignoring most of the available information actually work? A growing number of studies have shown that simple heuristics of this type can indeed perform remarkably well in a variety of inferential settings, where it is possible to determine whether decisions are correct. In one case, a competition was run (Czerlinski et al., 1999) to compare the performance of two heuristics—the Minimalist heuristic using cues in random order and the Take The Best heuristic using cues in validity order—with two strategies that weight and combine all of the available cues—Dawes's Rule, which weights cues equally, and multiple linear regression, which weights cues differentially in an optimal fashion. These four algorithms were tested in 20 decision environments, ranging from judging homelessness rates to comparing professors' salaries on the basis of several specific cues. Across all 20 environments, the simple heuristics were indeed frugal, only looking up a third of the available information (and only ever using one cue to make a decision), whereas the other two strategies processed it all by design. And yet this frugality did not cost the heuristics much in terms of accuracy: When fitting the existing data, Minimalist and Take The Best scored 69% and 75%, respectively, while Dawes's Rule and multiple regression scored 73% and 77%, respectively.

But a more important measure is how well decision strategies do when applied to new data that they have not seen before, because such generalization to new situations is what decision makers must usually confront. On this dimension, Take The Best scored 71% across the 20 environments, whereas multiple regression, usually the gold standard for multi-attribute decision making, overfit the noise in the training data and hence fell further in performance than did Take The Best, to 68% accuracy (and Minimalist and Dawes's Rule scored 65% and 69%, respectively). The frugal information use and fast processing of Take The Best thus proved more robust than the precise weighting and adding of multiple regression, demonstrating that less information can be more successful in decision making between alternatives.

Not only are simple one-reason decision mechanisms accurate and robust, they also correspond to how people (and other animals) make decisions in a variety of circumstances. People use these fast and frugal algorithms in environments that have the appropriate structure, even when they must first learn how the environment is structured (Rieskamp and Otto, submitted for publication). Heuristics such as Take The Best are also particularly used where information is costly or time consuming to acquire (Rieskamp and Hoffrage, 1999; Bröder, 2000; Newell and Shanks, 2003), whether the costs come from searching for cues in the environment or from searching in memory (Bröder and Schiffer, 2003).

There is a problem lurking here, though, in applying one-reason decision strategies: How can we tell what cues a heuristic should use and in what order? As can be seen from the performance figures just given, Take The Best's validity-ordered cue search does considerably better than Minimalist's random search—but how do we come to know a more-or-less validity-ordered set of cues?

In evolutionarily important decision contexts like choosing a mate or selecting something to eat we might have some built-in knowledge of valid cues to use, such as facial symmetry or sweet taste. But we are unlikely to have innately specified cues to use, for instance, in deciding between restaurants. For decisions like this in modern environments, people must learn what cues are most useful or valid. This can be done through individual experience using simple learning rules, for instance, keeping an ordered list of possible cues and moving a cue up in the list every time it leads to a correct decision and down in the list every time it fails (Dieckmann and Todd, in press). Alternatively, people can learn a good cue order socially from other decision makers. This suggests a particular path for aiding individual choice: informing people about the cues that other successful decision makers have used, rather than about the specific choices they have made.

5. Multiple-cue decision heuristics—using few cues to choose

When there are more than two options to choose among, then more than a single binary cue must typically be used to determine a single choice. But here, too, in these situations of multi-attribute decision making (Montgomery and Svenson, 1976) it is possible to reach quick decisions using a minimal amount of information, rather than gathering and combining a large number of cues or attributes. A "fast and frugal" approach to these decision situations is to use the process of elimination, as incorporated by Tversky (1972) in his Elimination by Aspects (or EBA) choice mechanism. For instance, if there are several restaurants to be decided among, first pick a cue (or aspect) dimension somehow and a way of using that cue to discard some of the available options. In the case of EBA, the cues are picked probabilistically, and a threshold is set for determining which options are eliminated from further consideration, such as discarding all restaurants that are more than 10 km away. If there are still multiple options left to be considered, then select another cue and use it to eliminate some more possibilities—such as all restaurants not serving fish tonight. Proceed in this way, using successive cues to whittle down the set of remaining options, until only a single one remains, which is the final choice. Tversky found that this process describes well what people do in these types of preferential choice tasks.

A similar elimination process can be used to categorize objects or stimuli, where the task can be conceived of as deciding which of several possible categories the object best fits into (Berretty et al., 1999). When information may be difficult to come by, and decisions should be made quickly, a fast and frugal categorization process can be adaptive. Consider the situation of trying to decide about another individual's intentions as they approach: Does this person want to greet me, dance with me, or take my wallet? How can one judge this, especially if the person is a stranger and is not announcing her aims verbally or facially? One way is to come to a quick first guess on the basis of how she is moving, that is, using motion cues alone to make a rapid categorization (Blythe et al., 1999).

People readily ascribe intentions to other organisms just on the basis of their motions—a bird flying straight at you seems intent on attack, or a dog circling around intends to play. Heider and Simmel (1944) showed that people will even effortlessly attribute intentions to inanimate geometric shapes moving around in a simple two-dimensional cartoon. People watching such cartoons would spontaneously describe the actions as, for example, the angry triangle chasing after the adulterous circle. It is surprising that such intricate stories of internal mental states and desires would be generated on the basis of so little information, just the two-dimensional cues of whole-body motion. But this fits in with the perspective of bounded rationality, that decision makers will take shortcuts like this if there are reliable simple cues of intention from motion.

Barrett et al. (in press) wanted to show that the motion cues *are* reliable—that is, that people can *accurately* judge the intentions of moving agents. To do this, examples of motion trajectories had to be collected from organisms whose intentions could be determined. Natural examples, for instance, footage of cheetahs chasing gazelles, would be realistic but difficult to obtain and display in a controlled fashion, with all cues removed except for two-dimensional motion. Instead, the researchers developed a computerized game setting, in which two people sat at computers in separate rooms, each controlling the motion of one of two colored dots on the computer screen but able to see both. These dots had certain properties associated with them so that they would, for instance, accelerate and decelerate in a semi-natural way. The two motion generators were then instructed to move their dot with a certain intention: For instance, generator 1 would pursue and generator 2 would evade. A third person acting as a judge would watch the motion trajectories being generated by the other two people and select what intentional category they thought was being generated, out of a list of six that were used. Whenever the judge chose the correct intention—that is, whenever the generators moved their dots in a way that convinced the judge that they intended to, for instance, pursue and evade rather than play or do something else—all three participants were rewarded with a monetary payoff.

Barrett et al. used the trajectories generated in this way to test how accurately a new set of people could decide what the original intentions had been, based solely on the motion patterns they saw onscreen. Participants chose the correct intention out of the set of six possibilities nearly 80% of the time. Moreover, the researchers wanted to test whether this was just the result of cultural learning, for instance, from watching the same cartoons as children and so picking up a shared vocabulary of motion types, or if it reflected deeper evolved schemas for understanding social interactions. To do this, they showed the same trajectories to adults in a very different non-Western culture: the Shuar hunter-horticulturalists from the Amazon region of Ecuador. Adults there made the same intentional judgments as the original participants from Germany, suggesting that our cognitive adaptations for inferring intention from motion may be universal components of human psychology. Thus, the limited information available in two-dimensional whole-body motions is enough to make accurate decisions about intent, in line with the use of a fast and frugal categorization mechanism. Further studies are needed to determine the exact nature of this mechanism, such as whether it is based on elimination or some other rapid process, and the cues that are used to make this judgment, such as relative velocity, heading, and rate of turning.

Homing in on the cues and heuristics underlying judgments of intention from motion can lead to a number of applications. For instance, it will help in figuring out how to generate super-stimulus motion trajectories that give everyone a strong impression of intention. These trajectories can be used in making convincing animation, but also in interesting physical applications, such as designing robots that engage humans and trigger desired responses based on how they move with respect to their observers; for instance, robot guides in museums can engage in leading behavior in the hope that visitors will readily follow. This work can also be extended to analysis of intentions of other moving agents; drug companies in particular would be interested in an automated way of telling what rats are doing after they have been given an injection of some new substance—does it make them fight, or court, or play? Driving patterns could also be assessed using such an approach, where the moving body is someone else's car—are they cutting you off to flirt or fight?

6. Sequential search heuristics—looking for a good option

The above examples indicate how good decisions can be made among alternatives by searching for and using little information about each. But what about situations in which the alternatives themselves must be sought—wouldn't more search in such cases, finding more alternatives to choose from, be better than less search? Economists prescribe looking for alternatives until the cost of further search outweighs any potential benefits (Stigler, 1961) and then taking the best alternative seen so far. But often the world does not allow such an approach, limiting our knowledge of the costs of search, the benefits that future options might bring, and even the ability to return to a previously seen option. Are there fast and frugal heuristics that can be applied for such sequential decision tasks?

This is an important type of decision to study, because sequential search is ubiquitous, occurring whenever resources being sought are distributed in time or space and so cannot be considered (or at least not encountered) simultaneously. Searching for mates or friends, houses or habitats, jobs, parking spaces, shopping bargains, or restaurants to eat at all involve sequential decisions of this sort. The problem is that whatever option you currently have available—for instance, the restaurant that you are standing in front of—another possibly better option could become available in the future, so how can you decide when to stop searching and stick with the current (or some previous) option?

Search problems can be characterized by a number of features of the search environment and the knowledge and capabilities of the searcher (Schotter and Braunstein, 1981), including the following: First, is there a fixed search horizon, that is, a fixed number of alternatives that can be looked through? Second, what is the distribution of values of alternatives in the environment and what does the searcher know about this distribution? Third, are the alternatives or options ephemeral, that is, do they disappear once they have been seen and passed by, or do they stay around so that they can be returned to, or recalled, later? Fourth, are there search costs for evaluating each alternative? Fifth, are their switching costs for leaving behind a previously chosen alternative and selecting a new one? And sixth, what is the payoff function, that is, what reward does the searcher receive, based on the value of the chosen alternative? Given some particular set of these characteristics defining a specific search environment, the question then is, how can search be stopped appropriately? We are particularly interested in answers to this question that are in the form of psychological mechanisms—specifically simple heuristics—but as is done in studying the behavioral ecology of animals, it can be useful to first explore optimal approaches to help guide the search for the psychological shortcuts.

One class of search domains is of particular interest because it captures many important real-world decisions. In situations where there is competition for specific alternatives, as when buying unique items such as antiques or houses, looking for a job or job candidate, or seeking a mate, once the searcher has passed by an alternative and decided not to pick it, there may be no chance of changing one's mind and returning to that alternative later, because someone else will have bought the house the searcher rejected or married the person previously spurned. That is, these search settings have little or no possibility of recall—feature 3 in the list above. Also in such situations, the searcher probably will not know the range of possible alternatives ahead of time—feature 2 above—and will have to learn about this distribution as the search progresses. What approach can one take to search in such an environment? We can look at a specific simple model of this kind of search and see what the optimal approach is, what simpler decision methods work well, and what people actually do (see Todd and Miller, 1999, for more details). A problem of this form has been well studied in probability theory (Ferguson, 1989), where it is known as the secretary problem in the job search domain, or the dowry problem in the mate search domain. As the dowry problem, it goes like this:

A sultan wishes to test the wisdom of his chief advisor, to decide 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. If he chooses correctly, he gets to marry that woman and keep his post; if not, the chief executioner chops off his head. The advisor can see one woman at a time and ask her dowry; then he must decide immediately if he thinks 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 her pass, she is 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 greatest chance of picking the woman with the highest dowry?

In a search situation like this, where the distribution of available alternatives is unknown, there is no recall and no switching, then searching with an aspiration level can be appropriate—what Simon (1956, 1990) called *satisficing*. In particular, search can be divided into two phases: In the first phase, alternatives are just looked at without selecting any of them, so that the searcher can gather information about the available options. This information is used to set an aspiration level—the minimum value that the searcher will try to get in further search. The second phase then consists of looking at additional alternatives, until one is found that exceeds the aspiration level set in phase 1. Search is stopped at that point and that alternative is chosen. Once the aspiration level is set, the length of the second search phase is out of the searcher's control. But how long should the *first* phase be, and how should the aspiration level be set when it is done?

In the case of the dowry problem, the searcher is trying to maximize the chance of picking the single best alternative, here in terms of the highest dowry. The *optimal* way to set the aspiration level is to search long enough in phase 1 that enough information is obtained about the available values to make a good decision, but not *so* long that the searcher passes by the best alternative in phase 1 without selecting it. The length of phase 1 that optimizes this balance is to look at N/e of the available alternatives (Ferguson, 1989), where N is the search horizon length or number of alternatives and $e \approx 2.718$ is the base of the natural logarithm system. This comes out to 37% of N, so in other words, the optimal approach is to follow the 37% rule: In phase 1, look at 37% of the upcoming alternatives; then set the aspiration level to equal the highest value seen among all those alternatives; and finally continue search in phase 2 until an alternative is found that exceeds the aspiration level. This method gives a better than 1 in 3 chance (a 37% chance, in fact) of picking the highest dowry. This is a simple heuristic, and it is relatively successful, but it has a drawback—it is not particularly fast or frugal. In fact, the mean search time required when using the 37% rule—that is, the sum of phase 1 and phase 2—is 74% of the search horizon. Thus, for instance, if people faced a mate search situation akin to the dowry problem, then they would certainly have to do a lot of search to behave optimally, on average going through about three-quarters of the potential mates they might ever meet before making a final choice. Rising divorce rates notwithstanding, this is probably not a widespread strategy.

Do people actually use the 37% rule in these types of search settings? Seale and Rapoport (1997) experimentally investigated behavior in the secretary problem setting, looking for simple heuristics that could explain what participants did. They proposed three such heuristics, namely, a *cutoff rule*, a *candidate count rule* and a *successive non-candidate rule*. The cutoff rule is a generalization of the optimal 37% rule solution, where searchers simply pass by a certain number of options and then select the next encountered top-ranked option (so the 37% rule is a cutoff rule with the cutoff set at 37% of the possible alternatives). Defining each option that is top-ranked at the moment it is assessed as a *candidate*, the candidate count rule simply implies choosing the jth candidate seen. The successive non-candidate rule, on the other hand, chooses the first candidate that is interviewed after observing at least k consecutive non-candidates—that is, it stops searching after the gap between successive candidates has grown sufficiently large. All of these heuristics require only minimal cognitive resources (mainly counting and comparing options against the best seen so far), and the cutoff rule and successive non-candidate rule in particular can perform very well on dowry/secretary-type search problems given appropriate parameters.

Seale and Rapoport compared the predictions of the three search heuristics with the actual behavior of their participants when searching through sequences of 80 values (presented as relative ranks). The cutoff rule came out the best, being most consistent with observed search behavior for 21 out of 25 participants. However, a majority of the participants stopped earlier than prescribed by the optimal solution, using a cutoff of less than 37% of the 80 options, which led to success rates for finding the best option of

30–32% (compared to the optimally expected rate of 37%). Such early stopping is a common finding in search experiments, and it is usually thought to mean that, although people are being faster and more frugal in their search than the optimal approach prescribes, they are consequently not doing as well as they could.

But people *are* doing a good job of searching if in these tasks they are operating with a slightly different goal of always picking a high value, rather than the very highest. This is a reasonable assumption, because few search problems that people face in reality have exactly the form of the dowry or secretary problem. In particular, there are almost no situations where choosing the single best option yields maximal payoff while all other options yield zero payoff—picking the second-best job may give a slightly lower salary, for instance, and picking the second-best house may give a slightly smaller yard. When searchers are rewarded for finding any option in the top 10% of the available distribution, or when they receive a payoff proportional to the quality of the option chosen (and thus just try to select a high-valued option, not the highest), much less phase 1 and phase 2 search is required to perform well with a simple cutoff rule (Dudey and Todd, 2002; Todd and Miller, 1999). People put in such different payoff settings also adjust their behavior accordingly: In one small study where people searched through a set of 100 numbers and were rewarded whenever they stopped and chose a value in the top 10% of the distribution, people searched through 28 values on average (covering both phase 1 and phase 2) and stopped appropriately about 95% of the time. For rewards proportional to the value chosen, people searched even less and still selected values that were on average at the 94th percentile of the possible range (Dudey and Todd, 2002).

Thus, little search (and hence little information) is needed for a variety of sequential choice situations—at least in situations such as those presented so far where the searcher is in complete control of the choice process. But many sequential choice problems are actually not like this: Few of us are sultans, able to line up a selection of potential mates or job candidates and one-sidedly declare which one we will have. Most of the time, these kinds of searches are two-sided, which means the searchers are being searched by others at the same time, and choice must therefore be mutual. Job applicants must select their employer and be selected in return; two individuals seeking to marry must both decide to take the plunge together. This added challenge can be solved by the searchers learning their own value or rank position within their pool of fellow searchers and using this self-knowledge to determine how high they should aim their search aspirations (Kalick and Hamilton, 1986), rather than merely setting an aspiration level based on the values of a small sample of available options as in the one-sided approaches covered above. Todd and Miller (1999) presented a range of simple heuristics that do just this, for instance, heuristics for learning one's mate value through the acceptances and rejections encountered during an adolescent dating period or more generally a phase 1 search period (see Simão and Todd, 2002, for another approach to this problem). These heuristics, like the one-sided mechanisms already discussed, can perform well with little search, quickly learning appropriate aspiration levels based on the searcher's own quality.

With this knowledge of how people search for alternatives in different domains, we can design tools to help people search more effectively. For instance, we can order the alternatives presented to people in a way that allows them to end a search quickly if their goal is a rapid decision, or in a way that encourages them to search longer if their goal is to find the very best alternative or get a greater understanding of the search space. In designing these tools it is important to take into account individual differences in search style as well. There appear to be significant sex differences in search behavior in some domains; for example, Dudey and Todd (2002) found in their particular task that men searched quite a bit less than women. In this case, men were taking the riskier strategy, rather than searching longer and having a higher chance of finding the best option. This tendency could be reversed if it is made *more* risky to search longer, such as by saying the search is on a bad Web connection that could cut out at any moment and cause the searcher to lose all results found so far—in that case, if men are risk-seekers then they may be likely to search longer than women (see Slovic, 1966, for a difference of this sort between boys and girls).

These artificial agents acting on our behalf may be preferable to searching in person in some domains, because they could be made less susceptible to misjudging the distribution of available alternatives. For instance, when you see a photo of a beautiful person, your brain may be rewired in a way that changes your perception of the distribution of available mates or potential mating competitors, depending on your inclinations. This is because we evolved in an environment in which the only way we could perceive the information pattern corresponding to an adult human face was for someone to grow one for about 20 years and put it in front of us—and thus by definition all these faces we saw belonged to people who were potential mates or competitors. But today face-information can be thrown at people in any number of cheap ways, so that we are now flooded with images of beautiful people, all of which we take in as data, and our conception of the distribution of available mates can become skewed in the unattainable direction (Buss, 2000). The same thing holds for houses, or luxury goods, or jobs, watching the lifestyles of the rich and famous on television or in movies. This can ultimately lead us to search far too long for alternatives we will never be able to afford. Software search agents need not be so easily misled.

7. Further directions

In this paper we have seen several ways that people can use simple fast and frugal heuristics to make good decisions with little information in a variety of domains, and some of the implications this has for the design of systems that help people gather and make use of that limited information. In each case, the success of the heuristic has relied on the presence of certain structures or patterns in the environment—such as whether recognition alone is reliable, or what multiple cues are valid, or whether there is competition for items being searched for. The ongoing study of simple heuristics and their implications must be based on understanding the structure of information in different decision environments and the way this matches the structure of the decision heuristics themselves—that is, understanding ecological rationality.

An important class of environments awaiting further study from this perspective is that of social domains. What heuristics do people use when the information they can gather comes from other individuals, and how do the heuristics employed match the distribution of social information? Furthermore, what simple mechanisms do groups use to process the information spread among their individual members to come up with a final choice, and how do those mechanisms fit the group information structure? Researchers are beginning to make progress on these questions, for instance, finding that a fast and frugal search through information about one's social circle can be used to come up with good estimates of the prevalence of population-wide occurrences such as instances of different diseases (Pachur et al., in press). Looking at inferences made in small groups, Reimer and Katsikopoulos (in press) have found that recognition knowledge is accorded a special status not only for individual decision making but also in group deliberation. Social information gathering and processing when individuals are structured into differentially connected webs and networks (Watts, 2003) should reveal other patterns of heuristic use and ecological rationality.

Social environments made up of other decision makers with whom one must deal point to another challenge for understanding ecological rationality. In addition to looking at how our decision heuristics are shaped by the structure of the environment, we need to consider how our decision environments are in turn shaped by the decisions everyone makes. For instance, fast and frugal mate choice and search mechanisms can impact what cues potential mates display to entice each other and can lead to large population-level patterns such as the distribution of ages at which people first get married (Todd and Billari, 2003). But this decision-mechanism/environment co-evolution can take place much more rapidly and extensively when the environment is not just biological, but cultural. The way products become known and are chosen clearly impacts the structure of the market of available products—using recognition-based choice mechanisms can, for instance, result in just a few products being much more popular than most others, in a J-shaped distribution of selections made (Heuvelink, 2004; Todd and Heuvelink, submitted for publication; see also Janssen and Jager, 2003, for related studies combined with social network structure). Even the pattern of available parking spots in a large lot is created by the decisions made by others who have come there before you, and the strategy that you use to find a good spot will perform differently depending on the strategies that everyone else has already used in creating the environment structure that you encounter (Hutchinson et al., in preparation). Thus, exploring the nature of such pattern-constructing/pattern-exploiting feedback loops between heuristics and their environments will be important for our understanding of ecological rationality more broadly.

Knowing—or changing—where these loops are going can help us provide people with better tools for making decisions. More generally, knowing how people use simple heuristics to reach good decisions on the basis of little information can also help us create better communication technologies. But this knowledge can also be used to help divert people's decisions in directions they might not realize, and might not like. So, while a little information can go a long way, we need to work to ensure that it goes the way that people want.

References

Barrett, H.C., Todd, P.M., Blythe, P.W., Miller, G.F., in press. Accurate judgments of intention from motion cues alone. Evolution and Human Behavior.

Berretty, P.M., Todd, P.M., Martignon, L., 1999. Categorization by elimination: Using few cues to choose. In: Gigerenzer, G., Todd, P.M., the ABC Research Group (Eds.), Simple Heuristics that make us Smart. Oxford University Press, New York, pp. 235–254.

Blythe, P.W., Todd, P.M., Miller, G.F., 1999. How motion reveals intention: Categorizing social interactions. In: Gigerenzer, G., Todd, P.M., the ABC Research Group (Eds.), Simple Heuristics that make us Smart. Oxford University Press, New York, pp. 257–285.

Borges, B., Goldstein, D.G., Ortmann, A., Gigerenzer, G., 1999. Can ignorance beat the stock market? In: Gigerenzer, P.M., Todd, P.M., the ABC Research Group (Eds.), Simple Heuristics that make us Smart. Oxford University Press, New York, pp. 59–72.

Bröder, A., 2000. Assessing the empirical validity of the "Take The Best" heuristic as a model of human probabilistic inference. Journal of Experimental Psychology: Learning, Memory, and Cognition 26, 1332–1346.

Bröder, A., Schiffer, S., 2003. "Take The Best" versus simultaneous feature matching: Probabilistic inferences from memory and effects of representation format. Journal of Experimental Psychology: General 132, 277–293.

Buss, D.M., 2000. The evolution of happiness. American Psychologist 55, 15–23.

Czerlinski, J., Gigerenzer, G., Goldstein, D.G., 1999. How good are simple heuristics. In: Gigerenzer, G., Todd, P.M., the ABC Research Group (Eds.), Simple Heuristics That make us smart. Oxford University Press, New York, pp. 97–118.

Dieckmann, A., Todd, P.M., in press. Simple ways to construct search orders. In: Forbus, K., Gentner, D., Regier, T. (Eds.), Proceedings of the 26th Annual Conference of the Cognitive Science Society. Erlbaum, Mahwah, NJ.

Dudey, T., Todd, P.M., 2002. Making good decisions with minimal information: Simultaneous and sequential choice. Journal of Bioeconomics 3, 195–215.

Edwards, W., Fasolo, B., 2001. Decision technology. Annual Review of Psychology 52, 581–606.

Fasolo, B., McClelland, G.H., Todd, P.M., in press. Escaping the tyranny of choice: When fewer attributes make choice easier. Marketing Theory.

Ferguson, T.S., 1989. Who solved the secretary problem? Statistical Science 4, 282–296.

Frey, B., Eichenberger, R., 1996. Marriage paradoxes. Rationality and Society 8, 187–206.

Galef Jr., B.G., 1987. Social influences on the identification of toxic foods by Norway rats. Animal Learning and Behavior 15, 327–332.

Gigerenzer, G., 1996. Rationality: Why social context matters. In: Baltes, P., Staudinger, U.M. (Eds.), Interactive Minds: Life-span Perspectives on the Social Foundation of Cognition. Cambridge University Press, Cambridge, pp. 319–346.

Gigerenzer, G., 2002. Reckoning with Risk: Learning to Live with Uncertainty. Penguin, London.

Gigerenzer, G., Goldstein, D.G., 1999. Betting on one good reason: The Take The Best heuristic. In: Gigerenzer, G., Todd, P.M., the ABC Research Group (Eds.), Simple Heuristics that make us Smart. Oxford University Press, New York, pp. 75–95.

Gigerenzer, G., Selten, R. (Eds.), 2001. Bounded Rationality: The Adaptive Toolbox. MIT Press, Cambridge, MA.

Gigerenzer, G., Todd, P.M., 1999. Fast and frugal heuristics: The adaptive toolbox. In: Gigerenzer, G., Todd, P.M., the ABC Research Group (Eds.), Simple Heuristics that make us Smart. Oxford University Press, New York, pp. 3–34.

Gigerenzer, G., Todd, P.M., the ABC Research Group, 1999. Simple Heuristics that make us Smart. Oxford University Press, New York.

EOR 6863 ARTICLE IN PRESS No. of Pages 16, DTD = 5.0.1
13 May 2005 Disk Used

Goldstein, D.G., Gigerenzer, G., 1999. The recognition heuristic: How ignorance makes us smart. In: Gigerenzer, G., Todd, P.M., the ABC Research Group (Eds.), Simple Heuristics that make us Smart. Oxford University Press, New York, pp. 37–58.

Goldstein, D.G., Gigerenzer, G., 2002. Models of ecological rationality: The recognition heuristic. Psychological Review 109, 75–90.

Goodie, A.S., Ortmann, A., Davis, J.N., Bullock, S., Werner, G.M., 1999. Demons versus heuristics in artificial intelligence, behavioral ecology, and economics. In: Gigerenzer, P.M., Todd, P.M., the ABC Research Group (Eds.), Simple Heuristics that make us Smart. Oxford University Press, New York, pp. 327–355.

Heider, F., Simmel, M., 1944. An experimental study of apparent behavior. American Journal of Psychology 57, 243–259.

Henrich, J., McElreath, R., 2003. The evolution of cultural evolution. Evolutionary Anthropology 12, 123–135.

Hertwig, R., Todd, P.M., 2003. More is not always better: The benefits of cognitive limits. In: Hardman, D., Macchi, L. (Eds.), Thinking: Psychological Perspectives on Reasoning, Judgment and Decision Making. Wiley, Chichester, pp. 213–231.

Heuvelink, A., 2004. Monkey see, monkey do... Modeling the formation of social consensus. Unpublished master's thesis, Department of Cognitive Artificial Intelligence, Utrecht University, Utrecht, The Netherlands.

Hoffrage, U., Lindsey, S., Hertwig, R., Gigerenzer, G., 2000. Communicating statistical information. Science 290, 2261–2262.

Huberman, B.A., Pirolli, P.L., Pitkow, J.E., Lukose, R.M., 1998. Strong regularities in World Wide Web surfing. Science 280 (5360), 95–97.

Hutchinson, J., Fanselow, C., Todd, P.M., in preparation. Car parking as a game between simple heuristics.

Jacoby, L.L., Kelley, C., Brown, J., Jasechko, J., 1989. Becoming famous overnight: Limits on the ability to avoid unconscious influences of the past. Journal of Personality and Social Psychology 56, 326–338.

Janssen, M.A., Jager, W., 2003. Simulating market dynamics: Interactions between consumer psychology and social networks. Artificial Life 9, 343–356.

Kahneman, D., Slovic, P., Tversky, A. (Eds.), 1982. Judgment under Uncertainty: Heuristics and Biases. Cambridge University Press, New York.

Kalick, S.M., Hamilton, T.E., 1986. The matching hypothesis reexamined. Journal of Personality and Social Psychology 51, 673–682.

Kanazawa, S., 2002. Bowling with our imaginary friends. Evolution and Human Behavior 23, 167–171.

Katsikopoulos, K.V., Fasolo, B., submitted for publication. New tools for decision analysis.

Kluth, A., 2004. Make it simple: A survey of information technology. The Economist (October 30, Suppl.), 1–20.

Montgomery, H., Svenson, O., 1976. On decision rules and information processing strategies in multiattribute decision making. Scandinavian Journal of Psychology 17, 283–291.

Newell, B.R., Shanks, D.R., 2003. Take The Best or look at the rest? Factors influencing "one-reason" decision-making. Journal of Experimental Psychology: Learning, Memory, and Cognition 29, 53–65.

Noble, J., Todd, P.M., Tuci, E., 2001. Explaining social learning of food preferences without aversions: An evolutionary simulation model of Norway rats. Proceedings of the Royal Society of London B: Biological Sciences 268, 141–149.

Pachur, T., Rieskamp, J., Hertwig, R., in press. The social circle heuristic: Fast and frugal decisions based on small samples. In: Forbus, K., Gentner, D., Regier, T. (Eds.), Proceedings of the 26th Annual Conference of the Cognitive Science Society. Erlbaum, Mahwah, NJ.

Payne, J.W., Bettman, J.R., Johnson, E.J., 1993. The Adaptive Decision Maker. Cambridge University Press, New York.

Piattelli-Palmarini, M., 1996. Inevitable Illusions: How Mistakes of Reason Rule Our Minds. Wiley, New York.

Reimer, T., Katsikopoulos, K.V., in press. The use of recognition in group decision-making. Cognitive Science.

Rieskamp, J., Hoffrage, U., 1999. When do people use simple heuristics and how can we tell? In: Gigerenzer, G., Todd, P.M., the ABC Research Group (Eds.), Simple Heuristics That Make Us Smart. Oxford University Press, New York, pp. 14–167.

Rieskamp, J., Otto, P., submitted for publication. SSL: A theory of how people learn to select strategies.

Saad, G., Russo, J.E., 1996. Stopping criteria in sequential choice. Organizational Behavior and Human Decision Processes 67, 258–270.

Schotter, A., Braunstein, Y.M., 1981. Economic search: An experimental study. Economic Inquiry 19, 1–25.

Seale, D.A., Rapoport, A., 1997. Sequential decision making with relative ranks: An experimental investigation of the "secretary problem". Organizational Behavior and Human Decision Processes 69, 221–236.

Simão, J., Todd, P.M., 2002. Modeling mate choice in monogamous mating systems with courtship. Adaptive Behavior 10, 113–136.

Simon, H.A., 1956. Rational choice and the structure of environments. Psychological Review 63, 129–138.

Simon, H.A., 1990. Invariants of human behavior. Annual Review of Psychology 41, 1–19.

Slovic, P., 1966. Risk-taking in children: Age and sex differences. Child Development 37, 169–176.

Stigler, G.J., 1961. The economics of information. Journal of Political Economy 69, 213–225.

Todd, P.M., 2001. Fast and frugal heuristics for environmentally bounded minds. In: Gigerenzer, G., Selten, R. (Eds.), Bounded Rationality: The Adaptive Toolbox (Dahlem Workshop Report). MIT Press, Cambridge, MA, pp. 51–70.

Todd, P.M., Billari, F.C., 2003. Population-wide marriage patterns produced by individual mate-search heuristics. In: Billari, F.C., Prskawetz, A. (Eds.), Agent-based Computational Demography. Springer, Berlin, pp. 117–137.

Todd, P.M., Gigerenzer, G., 2000. Simple heuristics that make us smart. Behavioral and Brain Sciences 23, 727–741.

Todd, P.M., Heuvelink, A., submitted for publication. Shaping social environments with simple recognition heuristics. In: Carruthers, P. (Ed.), The innate mind: Culture and cognition.

Todd, P.M., Miller, G.F., 1999. From pride and prejudice to persuasion: Satisficing in mate search. In: Gigerenzer, G., Todd, P.M., the ABC Research Group (Eds.), Simple Heuristics that make us Smart. Oxford University Press, New York, pp. 287–308.

Todd, P.M., Fiddick, L., Krauss, S., 2000. Ecological rationality and its contents. Thinking and Reasoning 6, 375–384.

Todd, P.M., Gigerenzer, G., the ABC Research Group, 2000. How can we open up the adaptive toolbox? (Reply to commentaries) Behavioral and Brain Sciences 23, 767–780.

Tversky, A., 1972. Elimination by aspects: A theory of choice. Psychological Review 79, 281–299.

Tversky, A., Kahneman, D., 1974. Judgment under uncertainty: Heuristics and biases. Science 185, 1124–1131.

Watts, D.J., 2003. Six Degrees: The Science of a Connected Age. Norton, New York.