

Aspiration formation and satisficing in isolated and competitive search*

Werner Güth^a, Ev Martin^a, and Torsten Weiland^{a,†}

^a Max Planck Institute of Economics, Strategic Interaction Group,
Kahlaische Str. 10, 07745 Jena, Germany

November 7, 2006

Abstract

We experimentally explore individual and interactive decision making in a sequential search task and test whether generally accepted principles of bounded rationality (aspiration formation, satisficing, and aspiration adjustment) adequately explain the observed search behavior. Subjects can, at a cost, employ screening and selection methods facilitating their search and revealing their aspirations. The majority of subjects seems to follow the single threshold heuristic after extensive experimentation. Contrary to popular theories of sequential search, aspiration levels are set below the maximum value of all previously inspected alternatives. In a competitive search subjects tend to experiment less before engaging in satisficing and generally state lower aspirations. Finally, systematic satisficing seems to significantly enhance payoffs.

Keywords: sequential search, secretary problem, optimal stopping, bounded rationality

[*JEL*] D83, C44, C61

*The authors wish to thank Evelin Wacker for programming the experiment.

†Correspondence to: Torsten Weiland, Max-Planck-Institute of Economics, Kahlaische Strasse 10, D-07745 Jena, Germany. E-Mail: weiland@econ.mpg.de

1 Introduction

The neoclassical economic paradigm assumes optimizing agents with self-serving preferences and unlimited cognitive capabilities. Yet numerous empirical and experimental studies (cf., Cohen & Bacdayan, 1994; Egidi, 1996; Pereira, 1996) have shown limited cognitive capacity of human decision makers. In view of accumulating evidence that clearly conflicts with any obvious rational choice interpretation, this approach is being increasingly questioned.

Most prominently, Simon (1955) suggested an alternative bounded rationality approach which does not deny our desire to behave rationally but, at the same time, acknowledges our cognitive constraints. Consequently, when a decision problem is too complex, the subject will try to decrease the problem's complexity by neglecting minor effects, simplifying mental representations, and using easily measurable goals. Rather than trying to identify an optimal strategy, the agent is likely to form aspirations for some measurable aspects of performance. Within the concept of bounded rationality, such thresholds are said to be the agent's aspiration levels (e.g., Schank & Abelson, 1977; Holland, Holyoak, Nisbett, & Thagard, 1986).

Much of the literature on bounded rationality is primarily interested in advancing theory (e.g., Sauermann & Selten, 1962; Selten, 1998). There are only a few experimental studies in which aspiration levels of subjects in particular decision problems are empirically quantified and in which regularities in satisficing behavior are systematically analyzed.¹ A number of studies, for instance, focuses on bargaining situations (e.g., Tietz, Weber, Vidmajer, & Wentzel, 1978; Tietz, 1992). Although such studies certainly help to advance our understanding of bounded rationality, one may argue that the authors commenced with decisions problems that were too complex from the start, involving strategic interaction, without a sound basis of individual decision making.² Complexity can, of course, arise from strategic interaction as well as from multidimensional goal variables, which is the focus of Sauermann and Selten (1962).

More in line with our study, the experimental investigation of Fellner, Güth, and Maciejovsky (2005) analyzes individual (risky) investment choices and thus rules out strategic interaction. Subjects have to invest their endowment in various more or

¹ In economics, a strand of literature deals with reference points in investment decisions, auctions, and bargaining, and their adjustment due to information feedback and learning, whereas in psychology, the pertaining literature concentrates on aspects of prospect theory, motivation, and social comparison, or self-assessment.

² In a similar vein, Camerer (2003) claims to describe behavioral theories of game playing without ever discussing individual decision making.

less risky assets. Although the aspiration choices are not incentivized, the authors find that aspirations provide a reliable and more natural way of classifying investor types and predicting investment choices. In their subsequent study, Fellner, Güth, and Martin (2006) explore task transcending satisficing, i.e., whether and, if so, how individual aspiration formation and satisficing characteristics can be transferred from one decision problem to another. Their results suggest that observing aspirations in one task helps to predict the same individual’s aspiration formation in other more or less similar tasks. Unfortunately, the support for individual satisficing is often low (e.g., Fellner et al., 2006).³

This study differs from the preceding ones in that we investigate a decision problem that is theoretically demanding, but easily and intuitively understood. In particular, we consider

- individual as well as strategic decision making and
- one essential goal variable (i.e., to be specified by one aspiration level) which is naturally related to the monetary payoff.

The decision problem has already been used for illustration by Simon (1955) in his pioneering book on bounded rationality and is now well known as the “secretary problem” (see the review of Freeman, 1983). The decision maker is asked to choose an alternative, facing a random sequence of options that appear consecutively with no possibility to recall former alternatives, i.e., either accepting the presented alternative or losing it forever. Compared to preceding experiments based on the secretary problem,⁴ we are less interested in how search is related to optimal stopping behavior, and would rather concentrate on the process aspects of sequential search. More specifically, we have chosen an experimental design allowing us to directly observe the formation of aspirations and adjustments in the satisficing process.⁵

Exploring cognitive processes⁶ on the basis of boundedly rational decision making is still in an early stage. So far, little work has been done in terms of experimental

³ Fellner et al. (2005) find that 67% of their subjects are “potential satisficers,” i.e., their aspirations could be satisfied, and that 42% of all subjects invest such that their portfolio is in line with their aspiration level (“actual satisficers”). However, in the follow-up study of Fellner et al. (2006), even though the experimental setting is largely similar, the share of potential and actual satisficing is significantly smaller (48% and 30%, respectively).

⁴ See Bolle (1979) and Zwick, Rapoport, Lo, and Muthukrishnan (2003), to mention just a few.

⁵ Previous studies simply assume that after exploring a number of alternatives participants use them to form an aspiration level. Several heuristics, based on aspiration building, were tested for validity (see, e.g., Dudgey & Todd, 2001; Todd, Rieskamp, & Gigerenzer, in press).

⁶ Brain scanning might reveal what matters to us but not at all how we cognitively perceive and process cognitively and emotionally relevant features.

research on the formation of aspirations and subsequent satisficing. Thus, we felt there was a need for conducting a thorough exploration of those two aspects in a controlled laboratory experiment without having to develop a new paradigm. To render the cognitive processes that go along with systematic search more transparent and to provide the best conditions for testing aspiration formation and satisficing, we introduced a number of specific search tools that, if employed, will allow us to observe when a subject is forming an aspiration or when she is trying to satisfy a given aspiration. By our design we can even test whether subjects actually employ the offered screening and selection routines in a way that is in line with the bounded rationality approach to sequential search.

Section 2 briefly characterizes and motivates our modified version of the secretary problem and introduces the experimental design and the offered search routines. In section 3, we provide and discuss selected results from a simulation of single and multiple threshold search heuristics (henceforth *STR* and *MTR*) and derive our behavioral hypotheses for the subsequent experimental study. Section 4 lists the experimental protocol and presents selected results. Particular emphasis is placed on behavioral regularities in the search process and on the identification of key variables that affect the formation and use of aspirations. In section 5, we conclude by discussing and interpreting our main findings and point out directions for future research.

2 The model

2.1 The secretary problem

Our scenario is the well-known stopping task, colloquially known as the secretary problem (for a comprehensive review, see Freeman, 1983; Ferguson, 1989; Samuels, 1991) which we modify by incorporating a number of new features. The basic scenario can be characterized as follows:

1. The sequence of alternatives a_m with $m \in \{1, \dots, n\}$ is of known or unknown length n . This information constitutes a treatment variable and is varied within-subjects in the first two phases of the experiment.
2. Values of alternatives are randomly determined according to uniform densities where we rule out identical alternatives. Thus, the n alternatives can be ordinally ranked.

3. The sequential order of alternatives is randomly assigned such that the position of an alternative in the sequence and its value are uncorrelated.
4. Similar to preceding experiments,⁷ inspection of alternatives carries a cost. Three distinct search routines are available at a cost of $c(> 0)$. These search routines differ in their duration and their information provision.
5. For each alternative a_m that the decision maker inspects, he can ascertain its position m in the sequence and its real value v_m . Thus, he is perfectly able to compare the current alternative a_m with all previously inspected alternatives a_1, \dots, a_{m-1} .
6. If an alternative a_m is rejected, it cannot be recalled at a later point of time. If search continues until the n -th alternative, then this alternative a_n is automatically accepted.
7. Upon completion of the search task, the decision maker having selected the alternative a_i earns

$$\pi(a_i) = \frac{v_i - \min(v_m)}{\max(v_m) - \min(v_m)} * M \quad \text{with} \quad m = 1, \dots, n$$

whereby M is a positive amount that is known to participants.

This payoff function represents a special case of the rank-dependent payoff function with $M = \pi(1) \geq \pi(2) \geq \dots \geq \pi(n) = 0$ and the expression in parentheses representing the alternative's rank. Instead of merely conditioning on the alternative's ordinal rank, the payoff from an alternative directly (and linearly) depends on the alternative's absolute value in relation to the bounds of the sequence's value range. Generally, every choice generates a non-negative payoff which increases in the value v_i of the selected alternative a_i . A negative payoff may only be attained if total search cost exceed the payoff from the selected alternative.⁸

We opted for this payoff rule because it realistically represents the problem of choosing from a set of choices (e.g., labor productivity of hired staff, quality of a bought product). For the sake of generalizability we refrain from using a payoff function that concentrates on only one or few "superstars" such as the optimal choice

⁷ See Zwick et al. (2003), Kogut (1990), and Rapoport and Tversky (1970).

⁸ For further reference, this and alternative classes of payoff functions which are frequently encountered in generalized secretary problems, are introduced and discussed in more detail in Bearden, Murphy, and Rapoport (2004).

or the rank-dependent payoff rule with a small set of paid ranks. Moreover, to enable subjects to incrementally learn about the value distribution of alternatives in a sequence and to allow for the adjustment of aspirations, a cardinal payoff function is essentially necessary.

2.2 The experimental design

The experiment is divided in three phases, each including a different type of search task:

- Task (*IK*)
Subjects search **independently**, whereby they **know** the number n of alternatives in the sequence.
- Task (*IU*)
Subjects **independently** search a sequence of alternatives whose exact length is **unknown**. Instead, they are only given a range for n (as a rough indication) to discourage overhasty decisions at the beginning of a sequence. It is assured and known to subjects that the true length n of the sequence lies within some given range $\underline{n} \leq n \leq \bar{n}$.
- Task (*CK*)
Groups of four participants simultaneously search one single sequence of alternatives whose length n is publicly **known**. The **collective** search process is synchronized such that the pace of the slowest subject determines the overall speed of search in a group.⁹

The two conditions *IK* and *IU* model the standard secretary problem in which search stops once the subject selects an alternative or once all alternatives have been inspected. This scenario is certainly theoretically and empirically interesting, but the restrictive assumption made that only one subject is acting in isolation from others may not be applicable to many real-world situations. Task *CK* therefore allows for several subjects to simultaneously search a common sequence of alternatives. In the competitive environment, a subject would have an incentive to reduce the time spent on search and select a satisfactory alternative more quickly since an unduly long screening phase may decrease her chance of attaining a satisfactory alternative.

⁹ This rule needed to be introduced to grant sufficient time to subjects requiring more time to process the frequently updated screen information and to input their decisions.

Alternatives in treatment CK are allocated to subjects in a “first come, first served” manner. If there is only one claim on a particular alternative, it is directly assigned to the respective subject. If, however, more than one subject opt for the same alternative, the alternative is assigned to one of them by random draw. Afterwards the search process continues for the remaining active subjects who have not yet selected an alternative. Finally, when the remaining number of alternatives becomes scarce and still more subjects participate in the search, the risk increases that an unsatisfactory, remaining option is assigned by mandatory allocation.¹⁰

The scenarios IK and CK differ from the task IU in the information that is given to subjects. Both variants of the sequential search task with a “known” (IK and CK) or “unknown” (IU) sequence length can be justified: Zwick et al. (2003), for instance, rely on the “story” that participants have to check apartments from a given list, thereby informing them about the length of the sequence. This is replicated in our conditions IK and CK . Yet, real-world decision problems can be imagined where the number of options is not known, e.g., in mate choice, recruitment of specialists, etc., and where the number of choice alternatives cannot be predicted with certainty. We model such a setting in condition IU .

In scenario CK , we only explore the case in which the sequence length n is publicly known because we want all subjects to share the same belief about when the sequence will end. Note that we can disentangle the effect of time pressure and competition, e.g., by means of a within-subject comparison of individual search behavior in the IK and CK conditions.

Insert table 1 about here

In the first two phases of the experiment, half of all subjects experience the conditions IK and IU in this order, whereas the other half performs the two conditions in reverse order. After completing the two treatments of isolated search, subjects compete in search in the third phase (task CK) of the experiment. Table 1 illustrates the treatment design and the temporal arrangement of the experimental conditions.

¹⁰ We arbitrarily rule out the possibility that a subject does not receive any alternative. If, toward the end of the sequence, the number of remaining alternatives is about to fall below the number of remaining subjects, one option per search step is automatically and randomly assigned to an active subject.

2.3 Search routines

We modify the original secretary problem by offering participants several search routines. The routines are designed such that each has its own merits. They may be applied repeatedly and interchangeably. Whenever a routine is used anew, a cost of c is incurred. We will explore the temporal sequence, duration, and frequency of the applied routines to gain more insights into the individual decision dynamics in a sequential search problem. Let us now describe and motivate the three distinct search routines:

- ***S*-routine:** This routine represents the original way of searching in the secretary problem. It allows subjects to investigate one alternative at a time. When the *S*-routine is applied (assume at position m), at cost c per application, the next alternative in the sequence a_{m+1} is revealed which may then be selected or rejected. Our intuition is that *S* will be mainly used after unsuccessful satisficing, i.e., when a subject, having searched extensively, may fear to run out of alternatives.
- ***E*-routine:** This routine allows to screen several alternatives ($e \geq 2$) in one pass and to pay the search cost c only once instead of paying it for each alternative $a_m, a_{m+1}, \dots, a_{m+e}$.¹¹ Its drawback is that only the last alternative a_{m+e} may eventually be selected, since all previous alternatives are foregone and recall is impossible. If the presented alternative a_m coincides with the last alternative ($m = n$), the routine is aborted and a_n is automatically selected. The routine is designed to allow and reveal aspirations for experimentation, e.g., in order to learn more about the underlying distribution function $U(\underline{v}, \bar{v})$ and to form an aspiration concerning the value of a satisfactory candidate. One should therefore expect that a search starts out by using *E*.
- ***A*-routine:** This routine allows to inspect an unspecified number of alternatives in a sequence in one pass while also paying search cost c just once. When using this routine, a subject specifies an aspiration level \tilde{v} which is subsequently used to evaluate the “eligibility” of screened alternatives. When this routine is applied (assume at position m), the value v_{m+1} of alternative a_{m+1} is compared with \tilde{v} . If $v_{m+1} \geq \tilde{v}$, then this alternative is selected. Otherwise, the screening process continues by comparing $v_{m+2}, v_{m+3}, \dots, v_n$ with \tilde{v} . As soon as $v_{m+k} \geq \tilde{v}$ with $k = 1, \dots, n - m$ is fulfilled, the routine stops and a_{m+k}

¹¹ If one allowed for $e = 1$, then the *S*- and *E*-routines would be identical.

is selected. If it has not occurred earlier, the routine is aborted once $m = n$. When using this routine, the subject is not informed about the actual value v_{m+k} , unless $v_{m+k} \geq \tilde{v}$ is true and search ends. In all other cases, a merely generic statement is received that v_{m+k} was less than \tilde{v} . Subjects employing the A -routine are presumed to engage in satisfying their previously established aspiration. One can think of the A -routine as delegating a search task in the sense of “Scan all upcoming alternatives and select the first one that is better or equal to \tilde{v} !” We expected the use of A after the use of E and before any use of S .

As in earlier studies (e.g., Seale & Rapoport, 1997, 2000; Zwick et al., 2003), we also rule out the possibility of recall. As a consequence, a subject is throughout allowed to select only the last alternative of the set of alternatives that she has already inspected. Moreover, although the three routines feature the same search cost c per application, they obviously differ in their cost-performance ratio due to their specific functionalities. Clearly, inspecting alternatives stepwise (S -routine) costs more overall than acquiring a general idea about the value distribution (E -routine) and/or choosing the first satisfactory alternative irrespective of all others (A -routine).

By using the various routines, participants will hopefully provide conclusive data on their reasoning process, e.g., by the protocol of the sequentially effected search routines. Using the E -routine means that a subject engages in systematic experimentation before even considering to stop search and therefore, thus indicating an aspiration for experimentation. If the A -routine is applied after the E -routine has been used, this would suggest satisficing after forming aspirations. According to empirical findings in the literature¹² and intuition, at least experienced participants should

- (i) start their search by relying on the E -routine, then
- (ii) switch to the A -routine while basing their specification of \tilde{v} on some value of v_t (with $k = 1, \dots, m$) observed so far,
- (iii) possibly readjust the level of \tilde{v} after searching unsuccessfully for some time, and

¹² Gilbert and Mosteller (1966), Seale and Rapoport (1997), Todd (1997), and Seale and Rapoport (2000), to mention just a few.

- (iv) incrementally inspect alternatives (S -routine) when “panicking”, i.e., when fearing to run out of attractive alternatives.

(Not) knowing the number n of remaining alternatives (condition IU) might considerably affect the average experimentation period ($r * n$ with $0 < r < 1$). More specifically, we conjecture that not knowing the number n of alternatives will induce subjects to shorten experimentation and start satisficing earlier. This may, however, also crowd out “panicking,” i.e., the use of the S -routine after a long unsuccessful search. With respect to the competition treatment (CK), we expect participants to decrease the length of experimentation and to start satisficing earlier.

3 On single and multiple search thresholds

Before turning to our experimental results, we briefly consider search heuristics that are commonly debated in the literature on sequential search and will serve as benchmarks when analyzing our data. We distinguish the single (STR) and the multiple (MTR) threshold heuristics.¹³

3.1 The STR and MTR search heuristics

Both approaches commence with screening a subset of alternatives to develop some understanding of the value range for the alternatives $v_i \in [\underline{v}, \bar{v}]$ and the type of distribution producing them.¹⁴ After sufficient sampling (option E), a threshold value \tilde{v} is defined with which subsequent alternative values v_k with $k = t + 1, \dots, n$ are compared. If the value of the inspected alternative a_i is greater than, or equal to, this threshold value ($v_k \geq \tilde{v}$), alternative a_i is selected and the search ends. The solution of the classical secretary search task to maximize the probability of hiring the best candidate (Gilbert & Mosteller, 1966) is such a rule. In our interpretation, the STR heuristic describes the process in which a subject first experiments, then defines an aspiration level \tilde{v} which is maintained until it is either matched by a subsequent alternative or until the sequence ends.

¹³ Clearly, there also exist numerous other search heuristics such as the “horse race decision rule,” the “successive undesirable applicant decision rule,” and others (see Bearden et al., 2004). We decided to neglect these, however, since we wanted to concentrate on the setting and readjusting of (aspiration) thresholds in sequential search processes.

¹⁴ Most theoretical and experimental studies rely on the uniform distribution from which the values of the various alternatives are drawn. Changing the underlying distribution may affect the efficiency of certain search heuristics.

The multiple threshold approach allows the decision maker to (repeatedly) modify the initial threshold value \underline{v} . The search and adjustment procedure works as follows: The values of alternatives are first screened before determining more or less cutoff levels r_1, r_2, \dots, r_k with decreasing threshold levels. Usually, the maximum of the observed alternatives at the relative (sequence) position r_1 then defines the first aspiration level \tilde{v} . Search continues for a given time interval until a candidate a_k is found that satisfies $v_k \geq \tilde{v}$. If not, \tilde{v} is re-evaluated by $\tilde{v} = \max(v_k)$ with $k = 1, \dots, r_2 * n$. Again, search continues until a qualifying alternative is found. If not, \tilde{v} is again downgraded. The idea is to incrementally adapt to lower aspiration levels \underline{v} if the more ambitious ones cannot be satisfied within a certain time interval.¹⁵

Experimental findings (Seale & Rapoport, 1997, 2000) suggest that subjects who readjust their threshold level (henceforth \tilde{v}'), usually reduce it ($\tilde{v}' < \tilde{v}$).¹⁶ Clearly, this heuristic can diminish the risk of ending with the last alternative a_n whose expected value $E(a_n) = \frac{\bar{v}-v}{2}$ under the uniform distribution is rather small. Chow et al. (1964) point out in their theoretical study that it is indeed sensible to decrease one’s aspiration toward the (expected) end of the sequence of alternatives.

3.2 Simulating heuristics

To judge the efficiency¹⁷ (measured by the mean or median of the chosen v_i in relation to the interval bounds) of the *STR* and the *MTR*, we ran a Monte Carlo simulation of both heuristics.¹⁸ The simulation exercise also revealed that the particular rank-dependent payoff rule of our scenario is distinctly different from both the “best choice” (Gilbert & Mosteller, 1966) and the classical “rank-dependent” (Chow et al., 1964) payoff function.¹⁹

We first estimate the optimal cutoff with value r^* under the *STR* heuristic which

¹⁵ For a comprehensive description of the multiple threshold rule, see Lindley (1961) and Chow, Moriguti, Robbins, and Samuels (1964) who extensively discuss the class of *MTR* search heuristics.

¹⁶ Under *MTR*, a decision maker sets an initial threshold and subsequently updates \tilde{v} once or repeatedly. At present, there is still controversy in the literature if indeed subjects employ the more “sophisticated” *MTR* heuristic at all (e.g., Zwick et al., 2003) since the usually *S*-routine data are rather inconclusive.

¹⁷ When discussing the “efficiency” of a search heuristic, we refer to that heuristic’s mean (expected) payoff that is computed according to the payoff function stated in section 2.1.

¹⁸ The simulation was conducted using the open source statistics package *R*. A copy of the source code is available from the authors upon request.

¹⁹ Evidently, the payoff function has a direct and pronounced implication on the optimal cutoff value r^* (optimal threshold vector) under *STR* (*MTR*).

allows to derive the optimal aspiration level \underline{v}^* after having seen $r^* * n$ alternatives.²⁰ For that purpose, we set up a simulation that generates 1000 random sequences of alternatives of length $n = 30$ with distribution $U(0, 1)$.²¹ In each iteration, the first m alternatives with $m = r * n$ of the random sequence $a = (a_1, \dots, a_n)$ are inspected and the maximum value $\tilde{v} = \max(v_k)$ with $k = 1, \dots, m$ is determined. The search then continues until a candidate a_k is found that fulfills $v_k \geq \tilde{v}$ or the end of the sequence is reached. In the former (latter) case, v_k (v_n) is selected.

Insert figure 1 about here

The efficiency of the *STR* for all levels of r with $r \in [0, 1]$ is shown in figure 1. The solid (dashed) line indicates the mean (median) search efficiency at each level of r whereas the dotted lines denote the first and third quartile of the search efficiency distribution. The curve of the mean selected value has its maximum at $r^* = 0.17$ at which the mean (median) efficiency equals 86.2% (94.7%). Our derivation of the best cutoff value r^* differs considerably from the “Golden Rule” of sequential search (for an overview, see Sardelis & Valahas, 1999) that suggests to screen the first 37% ($1/e$) of a random sequence of length n before trying to satisfy the threshold value $\underline{v}^* = \max(v_k)$ with $k = 1, \dots, r^* * n$. This difference seems to be mainly caused by the diverging payoff functions (ordinally rank-dependent vs. cardinally valued alternatives).²²

Is there (at least) one specification of a *MTR* heuristic that stochastically dominates the *STR* heuristic in a direct comparison when applied to the identical set of randomly generated sequences of alternatives? An exhaustive answer would have to systematically compare the *STR* (with optimal cutoff value r^*) with all possible combinations of the *MTR* (defined by the threshold vector $r^* = r_1^*, \dots, r_k^*$ of variable length k). Here we only compare the *STR* heuristic with $r^* = 0.17$ against a *MTR* heuristic with the arbitrary threshold vector $r^* = (0.17, 0.65, 0.80, 0.90)$ with three aspiration adaptation steps.²³ Even the arbitrarily parametrized *MTR* features two advantages over the *STR* heuristic with the best cutoff value r^* :

²⁰ Note that we employ r^* as the relative duration of experimentation which implies that $0 < r^* < 1$ must be satisfied.

²¹ Additional simulations showed that decreasing the sequence length ($n = 10$) does not affect the benchmark prediction whereas increasing n shifts the optimal (relative) cutoff value r^* toward zero ($r^* = 0.12$ for $n = 50$). This finding suggests that r^* and n are related in a non-linear way.

²² Dudey and Todd (2001) also confirm that the optimal cutoff value crucially depends on the employed payoff rule.

²³ Another simulation was run in which the *STR* and *MTR* heuristics were tested on the identical set of 1000 random sequences of alternatives.

Insert figure 2 about here

First, the *MTR* on average finishes search earlier (see figure 2, left and middle panels). The two density plots exhibit the probability that the search is ended at a particular time t with $0.17 \leq t \leq 1$.²⁴ The second, and probably more important benefit is that the variance of outcomes in *STR* is significantly greater than in *MTR* ($\sigma_{STR} = 0.213$ vs. $\sigma_{MTR} = 0.175$, $p < 0.001$, F-test). Figure 2 (right panel) reveals that the first quartile value of the distribution of selected values in *MTR* is greater than the one in *STR* (0.862 vs. 0.828).²⁵ Since more search processes are concluded in the interval $0.65 \leq t < 1$ under the *MTR* heuristic, only a smaller number of searches actually lasts until $t = 1$, which explains the lower variance of payoffs in *MTR*.

Altogether, these properties render the *MTR* heuristic more attractive for risk-averse decision makers. We conclude from our simulation exercise that, given some experience, subjects are likely to prefer the *MTR* heuristic with multiple threshold revisions to the *STR* heuristic in which a simple threshold is established and subsequently maintained.

4 Experimental results

4.1 Experimental protocol

The experiment was computerized using z-tree (Fischbacher, 1999) and conducted with students from various faculties of the University of Jena who were recruited through ORSEE (Greiner, 2004). In total, 64 subjects participated in the four sessions of the experiment.

On entering the laboratory, participants received a set of instructions (see Appendix) explaining their tasks and how their performance was related to their payoff. Subjects were not immediately informed about the whole sequence of search tasks. Additional information on the further procedure of the experiment was provided after the end of the first and the second phase. An on-screen questionnaire was used to verify that the game rules were fully understood. After the questionnaire entries were checked and remaining questions were resolved, the actual experiment started.

²⁴ t denotes the relative search duration and ranges from 0.17 (stop search after having screened the first 17% of alternatives) to 1 (search the entire sequence).

²⁵ Apart from the difference in variance, (absolute) selected values are not significantly different under the *STR* and *MTR* heuristics (0.947 vs. 0.951, $p = 0.215$, MWU-test). The linked crosses within the two box plots indicate the sample averages.

Each of the three phases comprised four periods during which subjects were provided with a history of all previously inspected alternatives, which was displayed on their computer screen.²⁶ The final payoff of participants was determined as the sum of payoffs from three randomly drawn periods whereby it was publicly known that exactly one period from each of the three phases would be selected. In addition to the experiment-related payoff, participants received an initial endowment of €3.00.

The following parametrization was applied: The number of alternatives n in a sequence ranges from 15 to 35. The ratio of the maximum (\bar{v}) and minimum (\underline{v}) value of alternatives in a sequence (\bar{v}/\underline{v} with $\bar{v} > \underline{v} > 0$) ranges from 1.01 to 622.75. The cost c of initializing a search method equals $c = 20$ ECU, and a maximum payoff of $M = 1000$ ECU can be attained in a period (if $v_i = \bar{v}$), which corresponds to €5.00.

4.2 Commonly observed search heuristics

In the first part of our analysis, we focus on individual search heuristics to detect frequently recurring search profiles. For this purpose, we categorize the large number of distinct individual search strategies into nine generic search profiles whereby each of them accounts for at least 1% of all observations.²⁷ The remaining unclassified search attempts are subsumed in a tenth profile (“other”). All search profiles are listed in table 2 according to their relative frequency throughout the experiment. The table summarizes the mean duration of every applied search routine within a profile as well as the overall duration of the search profile.

Insert table 2 about here

According to the concept of boundedly rational decision making, subjects should at first generate a (simplified) mental model of the decision problem they face and then determine an achievement level that they intend to accomplish in the subsequent satisficing phase. We do indeed find support for this hypothesis in our data.

²⁶ Whenever the A -routine is selected, the value of the inspected alternative is not stated in the history (see section 2.3).

²⁷ For the purpose of classification, the labels of search profiles are condensed such that repeated applications of the same search method are summed up in one single letter, i.e., although they are structurally not identical, the observed search heuristics “ EEA ” and “ $EEEA$ ” are thus pooled as “ EA .” This pooling is supported by the fact that a search cost is only charged when initializing a new search routine (the S -routine is (re)initialized whenever used).

In 42% of all observed search processes, subjects initially inspect a series of alternatives from an ex-ante unknown distribution via the E -routine (henceforth simply E), before they determine an aspiration level and delegate the search task to the satisficing routine (A).²⁸ Surprisingly, in 23% of all cases, subjects begin their search with screening a series of alternatives and then directly choose the last alternative of the reported sequence (E).²⁹ In view of the uncertainty concerning the selected alternative, such participants seem mainly interested in keeping search cost low. In 8% of all cases, subjects first screen a set of alternatives and then change to individual search (ES), which can be interpreted as satisficing with frequent aspiration adaptations. Further search heuristics involve incrementally forming an aspiration with a subsequent delegation of the search task (SA), exclusive incremental search (S), and the “intuitive” sequential use of all three search methods (EAS).³⁰ In 4% of all cases, the search is immediately delegated which means that the subject specifies an aspiration level without having any idea of the sequence’s underlying value distribution.

By pooling related search profiles, we note that about 80% of all subjects start their search via the E -routine. The A -routine is the most widely used search method to reach a final decision, accounting for 53% of all observations and indicating the subject’s initial (or adjusted) aspiration level. Incidents of aspiration adaptations in which subjects revise their previously set aspiration are less prominent but far from rare. A minority of subjects subsequently readjusts (and lowers) their initial aspiration level once or even repeatedly (see section 4.4).

Observation 1: *Most search processes start out with systematic experimentation to form aspirations, continue with satisficing, and end with aspiration adaptation.*

²⁸ As stated in section 2.3, subjects are given three search routines to facilitate their search. They may inspect one alternative at a time (S), screen a series of alternatives at once (E), or define a minimal threshold that must be met or exceeded by the value of an alternative in order to be automatically selected (A).

²⁹ Exclusively relying on the E -routine is risky since the expected value of the last observed (or any other randomly drawn) alternative a_k merely amounts to $E(v_k) = \frac{\bar{v}-v}{2}$. Based on our data, we argue that the E -profile is rather a fragment of the longer search profile EA . If v_k is substantially lower than $\max(v_i)$ with $i = 1, \dots, k$, subjects will usually proceed in their search, whereas they immediately choose alternative a_k if $v_k \approx \max(v_i)$ is satisfied.

³⁰ We label the ($EA\dots$) search pattern as “intuitive” since it reflects the single threshold rule (STR) in case of EA or the more sophisticated multiple threshold rule (MTR) in case of EAS , respectively.

4.3 Forming an initial aspiration

In the following, we discuss essential variables influencing the value of the individual threshold level \tilde{v} and estimate \tilde{v} by a regression model. We say that subjects are forming an aspiration if they screen alternatives groupwise (E) or search by inspecting alternatives incrementally (S), and that the subject has formed an aspiration once (A) is selected for the first time. This definition allows us to unambiguously determine the aspiration level in four of the ten search profiles (EA , SA , EAS , and ESA), accounting for 54% of all search processes. All subsequent statements relating to aspirations and their fulfillment will exclusively refer to this subset of observations.

The share of subjects who set an (initial) aspiration monotonically increases from 32.8% in the first period and peaks at 70.3% in period 8. With competition the share shrinks to 45.3% in the final period. We further observe that a subject's propensity to form an aspiration at all depends neither on (not) knowing the sequence length n (conditions IK and IU) nor on the presence of competing agents (condition CK). However, with experience (i.e., with increasing periods, coefficient *period*) subjects tend to use an aspiration level more frequently (see table 3).³¹

Insert table 3 about here

Moreover, subjects rather consistently (do not) use a search threshold throughout the experiment. Subjects who (do not) define an aspiration in one period will do so again in the following period with a probability of 78.0% (71.4%).

Observation 2: *About half of all subjects set an aspiration level to facilitate their search. The propensity to use an aspiration slightly increases with experience but appears to be primarily determined by idiosyncratic characteristics.*

Which factors, then, determine the value of the aspiration level? To answer this question, we first compare the distributions of (relative) aspiration levels across periods and define the relative aspiration level as the ratio $v_{rel} = \frac{\tilde{v}-v}{\bar{v}-v}$ which considers the smallest (v) and that largest (\bar{v}) value in a sequence which has been observed for this subject so far. Figure 3 exhibits a series of box plots of observed relative aspirations which are grouped by period.³²

³¹ The initial model included a dummy marking that the sequence length is known and another marking the competition treatment. Insignificant coefficients were then iteratively removed from the model until only significant coefficients remained.

³² Observations classified as extreme outliers ($Q_1 - 3 * IQR$ or $Q_3 + 3 * IQR$) were removed from the sample since statements concerning the mean of key variables (e.g., relative aspiration level) would otherwise have been unduly biased. In all, 16 (out of 411) data points were removed.

Insert figure 3 about here

Relative aspiration levels start out in phase 1 with a median that equals 87.5% of the maximal, previously observed alternative and monotonically increase to 98.2%, while the spread of the stated aspiration levels gradually decreases. During periods 5 to 8, v_{rel} stabilizes at a high level of about 92.2%, after which it decreases again to 83.8% during the final periods (condition *CK*) while its variance continuously increases. For the last two periods the median of v_{rel} rises again to roughly the initial value but with considerable variance.

Figure 4 shows the empirical cumulative distribution function of relative aspiration levels where no distinction is made with respect to the period of observation. What can be learned from this figure is that 83% of all (initial) aspirations are set below or at the maximum value of all previously inspected alternatives. A proportion of 23.0% of all set aspirations v_{rel} lie between 0.8 and 0.9 of \bar{v} , 31.4% between 0.9 and 1, and only 16.2% exceed 1. In contradiction to both the familiar specifications of the *STR* and *MTR* heuristics, subjects do not generally set \tilde{v} equal to \bar{v} but allow for a certain discount. This can be interpreted as an attempt to attenuate the risk of an overly high aspiration.

Insert figure 4 about here

In a further step, we estimate a linear mixed-effects regression enabling us to more thoroughly investigate the factors determining relative aspiration levels. The fitted coefficients are presented in table 4. The distribution of the dependent variable v_{rel} features a mean of $\mu = 87.80$ (with $\sigma = 21.71$) where a value of 100 signifies $v_{rel} = \bar{v}$.

Insert table 4 about here

We find that the competitive search condition (dummy $\delta_{phase=3}$) significantly and negatively affects v_{rel} : subjects are on average satisfied with a lower v_{rel} when knowing that they compete in search. In the model selection, we also tested the distinction between the conditions in which the sequence length is exactly (*IK* and *CK*) or only imprecisely known (*IU*). Moreover, we inspected whether the vaguely or precisely announced sequence duration (\underline{n} , \bar{n} in *IU* or n in *IK* and *CK*) and the current position in the sequence (m) might affect subject behavior. Yet all mentioned coefficients and dummies turned out to be insignificant.

Observation 3: *Relative aspiration levels are generally stable and only affected by the presence of competitors.*

Another mixed-effects regression relies on a very similar set of coefficients to predict the average timespan in which subjects form their (initial) aspiration. The estimated coefficients of this regression are shown in table 5.

Insert table 5 about here

To put the marginal effects of the coefficients into perspective, note that subjects screen on average 10.7 alternatives ($\sigma = 4.8$) before deciding on an aspiration level. Again, the lower (\underline{n}) and upper (\bar{n}) bound for the sequence length n exhibit a significant and positive effect on the duration of aspiration formation, as if subjects determine their experimentation phase as a constant fraction of the presumed sequence length n . Only a minority of subjects determines their aspiration level by screening a (roughly) constant number of alternatives in each period.³³ Moreover, subjects spend slightly more time on experimentation in phase 2 (irrespective of whether condition *IK* or *IU* applies) and less time in phase 3 (condition *CK*) compared to the first phase.

Insert figure 5 about here

All in all, subjects on average spend more time ($0.42 = r > r^* = 0.17$) on establishing their (initial) aspiration level than suggested by the simulation results of the *STR* heuristic (see figure 5).³⁴

Observation 4: *The amount of time that subjects spend on forming an aspiration depends mainly on sequence length and the presence of competitors. Generally, subjects tend to spend more time on screening alternatives than seems optimal (in light of the simulation in section 3).*

Provided that subjects use the highest value during experimentation as reference for their initial aspiration level, this should lead to observed aspiration levels being higher than predicted threshold levels based on the *STR* heuristic. However, as

³³ 27.5% of the subjects who repeatedly form an aspiration always rely on an experimentation phase of roughly the same absolute duration (which only varies by $\sigma \leq 2$).

³⁴ When comparing observed with predicted screening durations, we find that subjects systematically exceed the predicted length of experimentation (10.66 vs. 4.86, $p < 0.001$, paired Wilcoxon signed rank test).

a consequence of frequent aspiration “discounting,” we observe that subjects commonly set their minimal threshold level below both the maximal value of observed alternatives and the benchmark prediction.³⁵

4.4 Adapting an existing aspiration

How often do subjects adjust their aspiration level \tilde{v} in the course of the search process? To answer this question, we first have to define when an aspiration level is adapted. We say that a subject adapts her aspiration \tilde{v} whenever the A -routine is employed, i.e., as long as (A) is not interrupted, assume that subjects prefer to maintain \tilde{v} .

On average, 54% of all subjects employ the A -routine at least once during a search process. Of those, 77.6% define \tilde{v} exactly once, whereas the remaining subjects subsequently modify \tilde{v} once (18.7%) or twice (3.4%).³⁶ Of the altogether 92 cases in which an aspiration is adapted, almost all concede a lower achievement level.³⁷ From the initial mean aspiration level of 87.8% of the largest observed value, a first (second) aspiration adaptation decreases the threshold to 81.2% (66.1%). We find that initial aspiration levels are generally higher than revised aspirations (initial vs. once revised aspiration levels, $p < 0.001$; once vs. twice revised aspirations, $p < 0.001$, both paired, one-sided MWU-tests). Thus, unsuccessful search clearly seems to induce subjects to lower their aspirations. Moreover, the duration between setting and readjusting an aspiration varies considerably across subjects. Those subjects who modify their aspiration at all, maintain their initial aspiration \tilde{v} on average for seven search steps before readjusting it.

Observation 5: *Aspirations are rarely adapted and, if so, they are generally lowered.*

4.5 On search efficiency and payoffs

Already in the first experiment phase, the majority of subjects sample the unknown value distribution of alternatives before fixing an aspiration level \tilde{v} . As a result, 80% of subjects start their search with the E -routine, which is particularly designed

³⁵ Observed relative aspiration levels v_{rel} are significantly smaller than their corresponding benchmark values based on the STR heuristic (mean of 0.820 vs. 0.861, $p < 0.001$, MWU-test). To allow for a direct comparison between the two statistics, the relative aspiration level is computed by using the global rather than the observed minimum and maximum values in the sequence.

³⁶ Throughout the experiment, only one subject adjusted her initial aspiration three times.

³⁷ In two cases, the new aspiration is slightly higher than the one it replaces.

for this purpose whereas only 11% of subjects incrementally inspect the sequence of alternatives at a higher cost (S). Strikingly, all of the five most prominent search profiles – which, with only one exception, start out with E – feature a high degree of efficiency and tend to be employed even more frequently in later periods of the experiment. The other search profiles cause considerably more search cost, perform rather poorly, and lose in importance over time. Table 6 summarizes the relative frequency of the various search profiles across phases and lists aggregate outcome statistics.

Insert table 6 about here

The S -routine is less frequently used in later periods due to its high cost ($\beta_0 = 22.88$, $\beta_1 = -0.58$, $R^2 = 0.33$, $p = 0.030$).³⁸ The less frequent use of S is partially offset by a steady increase of A whose share decreases only marginally toward the end of the experiment ($\beta_0 = 31.36$, $\beta_1 = 1.17$, $R^2 = 0.34$, $p = 0.027$). The E -routine, to the contrary, has no pronounced time trend and features little variation in its share throughout the experiment ($\beta_0 = 51.67$, $\beta_1 = 0.38$, $R^2 = 0.17$, $p = 0.10$).

Observation 6: *Search heuristics are selected according to their (pay-off) efficiency. Highly efficient search profiles hold or even increase their share over time.*

To identify the main determinants of search efficiency, table 7 presents a corresponding regression model and details the fitted coefficients.³⁹

Insert table 7 about here

Generally, the mean (median) search efficiency over all periods equals 75.3% (86.1%).⁴⁰ The coefficient *period* denotes the period of observation. Surprisingly, we observe a continuous decrease in efficiency over periods, although this effect is rather weak. The negative slope is mainly due to the fact that the last phase of the

³⁸ To avoid inflated counts due to an unusually high individual usage frequency (particularly of the S -routine), only binary values (i.e., routine was (not) used by the subject) are aggregated for each period.

³⁹ The right panel of table 7 shows the estimates of a quantile regression (with $\tau = 0.5$) for the same model. We present this information as a reference, since the distribution function of efficiency is highly asymmetric, which questions the mean as an adequate predictor.

⁴⁰ As a comparison, note that by implementing the benchmark heuristics STR and MTR , using the parametrization as described in section 3, mean (median) search efficiency would have amounted to 0.839 (0.983) for STR and 0.918 (0.973) for MTR .

experiment is the relatively less efficient competition treatment (*CK*). Computation of aggregate statistics reveals that both efficiency and period payoffs are lower in *CK* ($\mu_{eff} = 0.735$, $\mu_{pay} = 684$) when compared to the first two phases featuring conditions *IK* ($\mu_{eff} = 0.780$, $\mu_{pay} = 729$) and *IU* ($\mu_{eff} = 0.746$, $\mu_{pay} = 690$).⁴¹

The first of the two dummies, $\delta_{period=1}$, marks the first period, whereas the second, δ_{aspir} , points out that the observed subject has formed an aspiration. The effect of $\delta_{period=1}$ accounts for inexperience at the start of the experiment and affects search efficiency only in period 1. The significant dummy δ_{aspir} supports the claim that forming an aspiration in fact increases payoff. Relying on one (or multiple) aspiration threshold in the search process yields clear advantages as it shortens search duration and generally decreases the variance of search outcomes. Together, these aspects suggest that satisficing behavior is a reasonable heuristic in search tasks.

Insert figure 6 about here

Mean earnings in the experiment amounted to €10.81 ($\sigma = 2.50$). We further tested if there was a structural relationship between the subject's cumulative payoff and her propensity to use aspirations in the search process.⁴² Assigning subjects to two groups by median split of cumulative payoffs revealed that subjects with superior (inferior) final payoffs (mean of 9675 vs. 7146) used aspiration thresholds more (less) frequently throughout the experiment (median of 10 vs. 6 periods in which aspirations were formed, $p = 0.019$, MWU-test).⁴³ The positive effect of regularly using aspirations on the subject's cumulative payoff is illustrated in figure 6.⁴⁴ According to our data, superior cumulative payoffs exclusively result from a higher search efficiency (mean of 0.860 vs. 0.645, $p < 0.001$, MWU-test) but not from lower search cost (mean of 49.06 vs. 55.52, $p = 0.320$, MWU-test).

Observation 7: *Subjects who regularly use an aspiration level in their search earn significantly more (than those who use a threshold-based search heuristic less often or not at all).*

⁴¹ Efficiency and payoffs in *CK* are significantly less than in *IK* ($p = 0.007$ and $p = 0.006$) and slightly less than in *IU* ($p = 0.059$ and $p = 0.129$, MWU-test).

⁴² Note that this is a hypothetical investigation, since a stochastic rather than a cumulative payoff rule was employed in the experiment. However, if possible income effects are considered as negligible, cumulative earnings are a suitable statistic to describe steady (or mean) subject performance.

⁴³ Replicating the above-stated test on the basis of the median split of observed final payoffs leads to the same conclusion (mean final payoff of €12.77 vs. €8.85, median of 10 vs. 6 periods of using aspirations, $p = 0.012$, MWU-test).

⁴⁴ The width of each box plot is determined by the log value of the relative frequency of observations it is based on.

5 Discussion

For the experimental study of search behavior in sequential search tasks and an explicit analysis the process of aspiration formation, satisficing, and aspiration adaptation, several search routines are provided which facilitate search. Relying on the classical “secretary problem,” we distinguish several experimental conditions by providing (im)perfect information about the sequence length n and inducing (no) competition. We also replace the standard payoff function by a cardinal one.⁴⁵

Overall, our results partly reinforce and partly question the findings of earlier studies. In our experiment, about half of all subjects adopt the offered search routines in the natural order of initial aspiration formation and subsequent satisficing. Unlike assumed by various heuristics discussed in the literature, we do not generally observe that subjects set their aspiration level at the maximal value observed in the experiment so far. Rather, the majority of subjects is satisfied with less (about 90% of that value). Moreover, our results support earlier claims that aspiration levels are rarely adjusted and, if so, commonly decreased after an unsuccessful search (cf., Simon, 1955; Tietz, 1997). Our data also reveal that a considerable share of subjects, who use an aspiration at all, spend much time on forming that aspiration before moving on to satisficing, even after gaining more experience.

On the contrary, subjects do not significantly modify their behavior when information about the length of the sequence is provided (*IK* and *CK*) or withheld (*IU*), since with the unknown sequence length (*IU*), the lower and upper bound of the uniform distribution from which n is drawn seem rather good substitutes for knowing n exactly. Coefficient estimates from two regression models suggest that subjects in *IU* tend to primarily rely on the lower bound \underline{n} as an indicator of the unknown sequence length n .

One central finding is that structured search, in the sense of satisficing behavior, significantly increases payoffs. Subjects who regularly use aspirations achieve a superior search efficiency, incur about the same search cost and consequently fare significantly better. Competition generally induces subjects to set lower aspiration levels and to reduce the time spent on forming an aspiration. This effect is particularly strong during the first two periods of the competition scenario (*CK*) after which it becomes weaker.

One way to gain more insights into why subjects persistently feature considerable

⁴⁵ In this context, Miller and Todd (1998) and Gigerenzer and Todd (1999) showed that when searching with broader goals (e.g., “Pick one alternative among the top 10%!”) the optimal screening behavior significantly changes while the mean search efficiency decreases only marginally.

heterogeneity in their (idiosyncratic) search behavior and why they systematically deviate from “classical” benchmarks (e.g., *STR* and *MTR* heuristics), may be to directly elicit from subjects how they evaluate the quality of their own search and how satisfied they are with their search outcome. From their related study on sequential search, Bearden et al. (2004) conclude that subjects systematically overestimate the quality of their selected alternative and underestimate the potential of the so far unrevealed alternatives. This, however, does not exclude satisfaction with the own achieved outcome. It may also be of interest to explore the absorbability of several search heuristics (such as *STR*, *MTR*, and further heuristics with a varying parameterization and more or less related satisficing behavior). By offering a menu of alternative search routines to subjects, possibly after letting subjects test all of them, it could be observed whether certain specifications are systematically preferred to others and as such more closely resemble genuine human search behavior.

A Instructions (originally in German)

Thank you for participating in this experiment. For your punctual arrival you receive a show-up bonus of €3.00. Please read the following instructions carefully and from now on, do not speak to other participants anymore! If you have questions concerning the experiment, please raise your hand, and one of the experimenters will come to your place to answer your questions individually. During the experiment you have the possibility to earn money. Note that your earnings in the experiment will be stated in ECU (experimental currency units). At the end of the experiment, your accumulated ECU balance will be converted into euros at the rate of 200 ECU = €1.00 and will be disbursed to you in cash.

The experiment comprises two phases for each of which you will receive separate instructions. At the moment, you are reading the instructions for the first phase. You will receive the second part of the instructions once this phase is finished. Your final payoff in the experiment will be determined by the sum of your payoffs in three randomly drawn periods (two periods from phase 1 and one period from phase 2). Note that it is possible to realize a loss in the experiment. Should your final payoff at the end of the experiment be negative and exceed the show-up bonus, we will kindly ask you to “compensate” for that loss by staying in the room under supervision for five more minutes per €1.00 lost.

INSTRUCTIONS FOR PHASE 1

Your task

In each period of the experiment, you are presented with a series of alternatives which each hold a particular value. It is your task to choose one of these alternatives whereby a higher value of an alternative is associated with a higher payoff. The values of alternatives are randomly drawn from an interval whose bounds are, however, not known to you. Further, these values are uniformly distributed over the interval, i.e., each value has the same probability of being drawn. Moreover, the chronological order of alternatives is also randomly determined. Once you select an alternative, your search task and the period end. Should you not yet have chosen an alternative when reaching the end of the sequence, the last alternative will be automatically chosen for you.

For searching and selecting your preferred alternative, you are provided with

three distinct search routines which feature different properties. To commence search, you have to select one of the three routines. Thereafter, you are free to switch between search routines once or repeatedly (depending on the active routine, a constraint may apply). During each period, the currently active search routine is highlighted on your computer screen, and all other selectable search routines are indicated. If you have both the interest and the possibility to switch to a different search routine, you may do so by selecting that routine by mouse click. Depending on the chosen routine, you may have to provide additional parameters to start it. The search routines have in common that, in each routine, alternatives and their values are inspected sequentially, i.e., only one alternative is shown at a time.

Single observation routine

By selecting this routine, you are informed about the value of the next alternative in the sequence. If you wish, you may then select that alternative. If you decide not to select it, you may not return to it at any later point of time. In this case, you continue the search by again choosing one of the three search routines. By inspecting a single alternative, you will incur a cost of 20 ECU.

Group observation routine

By selecting this routine, you define an arbitrary number of alternatives to be inspected which may, however, not be selected. The specified number of alternatives and their respective values are then shown to you sequentially whereby it is *not* possible to abort the search in progress. You move on to the next alternative in line by clicking on the “continue” button. Generally, the specified number of alternatives can only be viewed but none of them can be selected. The only exception to this rule concerns the last of the inspected alternatives which you may choose if you wish. By inspecting a defined number of alternatives, you will incur a cost of 20 ECU once.

Automated search routine

When selecting this routine, you first define a minimal value that an alternative must satisfy such that you are willing to choose one. Once you have stated a minimal value, the automated search starts. In this process, the next (and subsequent) alternative(s) is (are) screened, and the value of each successive alternative is compared with the minimal value. If the alternative’s value is equal to or greater than the minimal value, it is automatically selected, its value is shown to you, and the search ends. Otherwise, you are only informed that the value of the inspected alternative

was less than your stated minimal value. Note that you are nevertheless free to select the alternative whose precise value is unknown to you. You can terminate or modify the automated search at any time by choosing a different search routine or respectively, by entering a new minimal value for the automated search. Note that you are not allowed to return to an alternative that has already been rejected by the automated search routine, whereas you may choose the active alternative if you decide to terminate the automated search. By using the automated search routine, you will incur a cost of 20 ECU once. Whenever you modify the minimal value, this is considered as another start of the routine which again entails a cost of 20 ECU.

Miscellaneous

The first phase of the experiment consists of eight periods whereby one search task corresponds to one period. In four of the eight periods, you will be informed about the exact number of alternatives in the sequence prior to starting your search. In the other four periods, you are only imprecisely informed about the exact sequence length. More specifically, you are told how long the sequence is at least (lower bound) and at most (upper bound).

To facilitate your search you can inform yourself at any time about the number and values (if known) of observed alternatives, which are displayed on your computer screen.

At the end of each period, you are not informed about the actual number of alternatives in the sequence and the sequence's minimal and maximal values, neither will you learn about the success of your search and the associated payoff in the respective period. A comprehensive summary of all periods will be provided to you at the end of the experiment.

How to determine your period payoff

Your period payoff is the difference between your income from the selected alternative and the search cost.

$$\text{Period payoff} = \text{Income from alternative} - \text{Search cost}$$

Your income from an alternative is calculated as follows:

$$\text{Period payoff} = \frac{\text{selected value} - \text{minimal value}}{\text{maximal value} - \text{minimal value}} * 1000$$

Your income thus depends on the value of your selected alternative in relation to the maximal and minimal value of all alternatives in the sequence. The relationship between the three above-mentioned values is illustrated in the figure below. The closer the selected value is to the maximal sequence value, the closer your income is to 1000 ECU. For instance, if you happen to select the maximal (minimal) sequence value, your income equals 1000 (0). In all other cases, your income equals an amount in between $[0, 1000]$.

Whenever you start a new search routine, switch between routines, or apply a




Figure 7: Income from alternative

routine repeatedly, you incur a search cost of 20 ECU.

$$\text{Search cost} = \text{count of employed search routines} * 20$$

An example:

The computer sequentially presents you with randomly drawn values from the interval $[2139, 2478]$. The value of the alternative you choose equals 2359. Given these three values, you can now compute your period payoff which equals:

$$\text{Income} = \frac{2359 - 2139}{2478 - 2139} * 1000 = 650$$

$$\text{Period payoff} = 650 - \text{Search cost}$$

Note that you may incur a loss in a period if your search cost exceeds your income from the chosen alternative!

Your payoff in phase 1

At the end of the first phase, two out of the eight periods are randomly selected whereby it is ensured that one search task from periods 1 to 4 and another task from periods 5 to 8 is chosen. Thus, your actual payoff in phase 1 is the sum of your period payoffs in these two periods. Should this amount be negative, then your loss in phase 1 may be compensated by a gain in phase 2 or your show-up bonus. If this should still be insufficient to cover your loss, we will kindly ask you to “compensate” for that loss by staying in the room under supervision for five more minutes per €1.00 lost.

Important notice

You are only allowed to participate in this experiment if you accept the above-stated rules. Otherwise, we kindly ask you to raise your hand now. We will then ask you to leave the room, and another participant will join the experiment in your place.

INSTRUCTIONS FOR PHASE 2

Your task

Again, in the second phase of the experiment, your task is to choose one alternative out of a series of alternatives. Unlike in phase 1, you are, however, searching a sequence of alternatives simultaneously with three other participants. All participants in your group search the identical sequence of alternatives whereby the chronological order of alternatives is the same for all participants. Moreover, the second phase only consists of four periods. Prior to starting search, you will be informed about the exact length of the sequence. To effect your search, you are again provided with the three search routines that were introduced in phase 1. Despite the competitive search environment, you still make your choices individually.

All participants in a group are simultaneously informed about the active alternative, e.g., the next alternative is only shown once the slowest participant has stated his choice of the current alternative. If just one participant in the group selects an alternative, he is assigned that alternative and exits the search, while the remaining active participants move on to the next alternative. If two or more participants in a group claim the same alternative, it is randomly assigned to one of them. The unsuccessful participant(s) continue(s) his (their) search.

Should the number of remaining alternatives equal the number of active participants in the group, the following allocation procedure applies:

1. All active participants in the group state their choice.
2. If one participant or several of the active participants select(s) the current alternative, it is assigned to one participant as described above.
3. If none of the active participants chooses the alternative, it is randomly (and without consent) assigned to one of them. The remaining participant(s) continue(s) his (their) search.

An example:

Three participants in a group have not yet chosen an alternative while there are only three remaining alternatives in the sequence. Since nobody is willing to accept the current alternative, this implies that if the search was to be continued, only two alternatives would be left for three participants. To avoid the consequential shortage

of alternatives, one of the participants is automatically assigned the alternative. The two remaining participants then continue their search.

Similar to phase 1, at the end of each period you are not informed about the value range of alternatives, the choices of co-players, and your period payoff.

Your payoff in phase 2

Once the second phase is concluded, you will receive a comprehensive summary of the interval bounds of each sequence, your choice in each period, and the resulting period payoffs. Your actual payoff in phase 2 is determined similarly to the payoff computation in phase 1: one of the four periods is selected at random whose payoff then determines your actual payoff in phase 2.

Your final payoff

At the end of the experiment, we will kindly ask you to remain seated and to answer a number of questions concerning the experiment. Subsequent to finishing the questionnaire, you will receive your payment for participating in this experiment. For this purpose, the payoffs from phases 1 and 2 are summed up and the resulting amount is converted into euros at the rate of $200 \text{ ECU} = \text{€}1.00$ and disbursed to you in cash.

References

- Bearden, J., Murphy, R., & Rapoport, A. (2004). *Sequential observation and selection with rank-dependent payoffs: An experimental test*. (Working Paper, University of Arizona)
- Bolle, F. (1979). Das Problem des optimalen Stoppens (Doctoral dissertation, University of Hamburg, Germany). *PhD thesis, University of Hamburg, Germany*.
- Camerer, C. (2003). *Behavioral game theory*. Princeton University Press.
- Chow, Y. S., Moriguti, S., Robbins, H., & Samuels, M. (1964). Optimal selection based on relative rank (the “secretary problem”). *Israel Journal of Mathematics*, 2, 81-90.
- Cohen, M., & Bacdayan, P. (1994). Organizational routines are stored as procedural memory: Evidence from a laboratory study. *Organization Science*, 554-568.
- Dudey, T., & Todd, P. (2001). Making good decisions with minimal information: Simultaneous and sequential choice. *Journal of Bioeconomics*, 3, 195-215.
- Egidi, M. (1996). The rational foundations of economic behaviour. In K. Arrow (Ed.), (chap. Routines, hierarchies of problems and procedural behaviour: Some evidence from experiments). MacMillan, London.
- Fellner, G., Güth, W., & Maciejovsky, B. (2005). *Satisficing in financial decision making - A theoretical and experimental attempt to explore bounded rationality*. (Working Papers on Strategic Interaction 23, Max Planck Institute of Economics, Jena, 2005)
- Fellner, G., Güth, W., & Martin, E. (2006). *Task transcending satisficing - An experimental study*. (Working Papers on Strategic Interaction 9, Max Planck Institute of Economics, Jena, 2006)
- Ferguson, T. (1989). Who solved the secretary problem? *Statistical Science*, 4(3), 282-296.
- Fischbacher, U. (1999). *Zurich toolbox for readymade economic experiments*. (Working Paper No. 21, University of Zurich, Switzerland)
- Freeman, P. (1983). The secretary problem and its extentions: A review. *International Statistical Review*, 51, 189-206.
- Gigerenzer, G., & Todd, P. (1999). *Simple heuristics that make us smart* (G. Gigerenzer & P. Todd, Eds.). Oxford University Press.
- Gilbert, J., & Mosteller, F. (1966). Recognizing the maximum of a sequence. *Journal of the American Statistical Association*, 61(313), 35-73.
- Greiner, B. (2004). *The online recruitment system ORSEE 2.0 - A guide for the orga-*

- nization of experiments in economics. (Working Paper, University of Cologne)
- Holland, J., Holyoak, K., Nisbett, R., & Thagard, P. (1986). *Induction*. Cambridge, MA: MIT Press.
- Kogut, C. (1990). Consumer search behavior and sunk costs. *Journal of Economic Behavior and Organization*, 14(3), 381-392.
- Lindley, D. V. (1961). Dynamic programming and decision theory. *Applied Statistics*, 10, 39-51.
- Miller, G., & Todd, P. (1998). Mate choice turns cognitive. *Trends in Cognitive Sciences*, 2(5), 190-198.
- Pereira, M. (1996). *The emergence of cooperative playing routines: Optimality and learning*. (Laboratorio di Economia Sperimentale e Computazionale, Università di Trento)
- Rapoport, A., & Tversky, A. (1970). Choice behavior in an optimal stopping task. *Organizational Behavior and Human Preferences*, 5, 105-120.
- Samuels, S. M. (1991). Handbook of sequential analysis. In B. K. Ghosh & P. K. Sen (Eds.), (p. 381-405). Marcel Dekker, Boston.
- Sardelis, D. A., & Valahas, T. M. (1999). Decision making: A golden rule. *The American Mathematical Monthly*, 106(3), 215-226.
- Sauermann, H., & Selten, R. (1962). Anspruchsanpassungstheorie der Unternehmung. *Zeitschrift für die Gesamte Staatswissenschaft*, 118, 577-597.
- Schank, R., & Abelson, R. (1977). *Scripts, plans, goals, and understanding: An inquiry into human knowledge structures* (L. Erlbaum, Ed.). Hillsdale, N.J.
- Seale, D., & 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.
- Seale, D., & Rapoport, A. (2000). Optimal stopping behavior with relative ranks: The secretary problem with unknown population size. *Journal of Behavioral Decision Making*, 13, 391-411.
- Selten, R. (1998). Features of experimentally observed bounded rationality. *European Economic Review*, 42, 413-436.
- Simon, H. A. (1955). A behavioral model of rational choice. *Quarterly Journal of Economics*, 69(1), 99-118.
- Tietz, R. (1992). Semi-normative theories based on bounded rationality. *Journal of Economic Psychology*, 42, 297-314.
- Tietz, R. (1997). Understanding strategic interaction. Essays in honor of Reinhard Selten. In W. Albers, W. Güth, P. Hammerstein, B. Moldovanu, & E. van

- Damme (Eds.), (p. 345-364). Berlin and Heidelberg: Springer.
- Tietz, R., Weber, H., Vidmajer, U., & Wentzel, C. (1978). Bargaining behavior, Beiträge zur experimentellen Wirtschaftsforschung. In H. Sauermann (Ed.), (Vol. 7, p. 88-102).
- Todd, P. (1997). Simulating social phenomena. In R. Conte, R. Hegselmann, & P. Terna (Eds.), (p. 419-436). Springer-Verlag, Berlin.
- Todd, P., Rieskamp, J., & Gigerenzer, G. (in press). Handbook of experimental economics results. In C. Plott & V. Smith (Eds.), (chap. Social heuristics). Amsterdam: North Holland/Elsevier Press.
- Zwick, R., Rapoport, A., Lo, A., & Muthukrishnan, A. (2003). Consumer search: Not enough or too much? *Marketing Science*, 22(4), 503-519.

Vitae

Werner Güth

is presently the Director of the Strategic Interaction Group at the Max Planck Institute of Economics in Jena and an Honorary Professor at the University of Jena. His research is mainly in microeconomics, evolutionary theory, and experimental economics. His web page is <http://www.econ.mpg.de/esi/gueth/index.html>.

Ev Martin

is a research assistant at the Max-Planck-Institute of Economics in Jena. Her research interests are in behavioral game theory and experimental economics, with a particular emphasis on aspects of bounded rationality.

Torsten Weiland

is a research assistant at the Max-Planck-Institute of Economics in Jena. His research interests include individual and interactive decision making, behavioral game theory, and experimental economics.

Figures

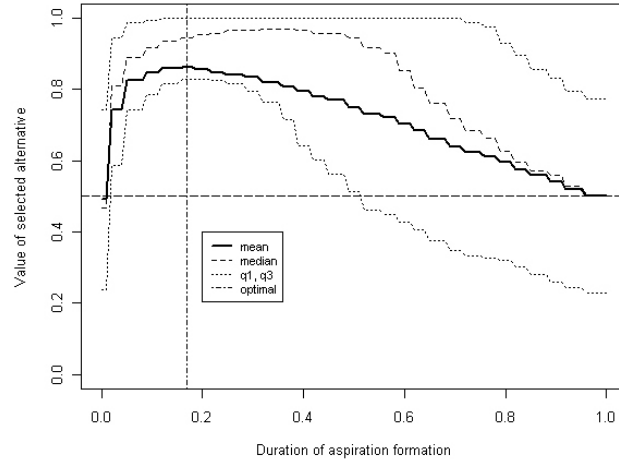


Figure 1: Optimal cutoff value (STR)




Figure 2: Search duration and variance of outcomes (STR & MTR)




Figure 3: Relative aspiration levels across periods

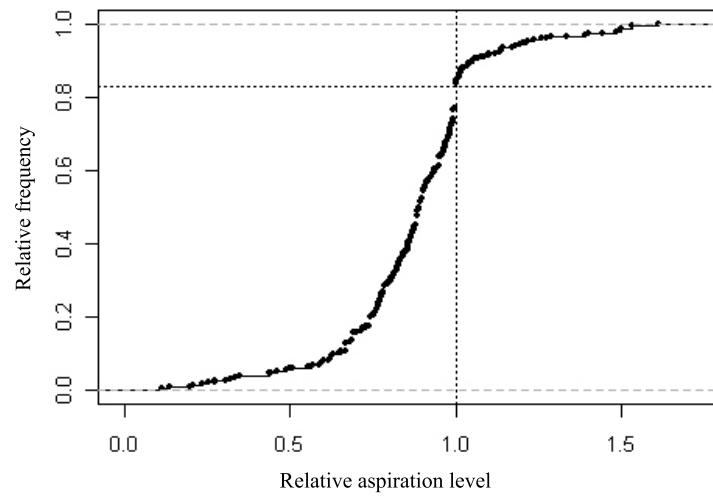


Figure 4: Cumulative distribution function of relative aspiration levels




Figure 5: Relative duration of aspiration formation phase

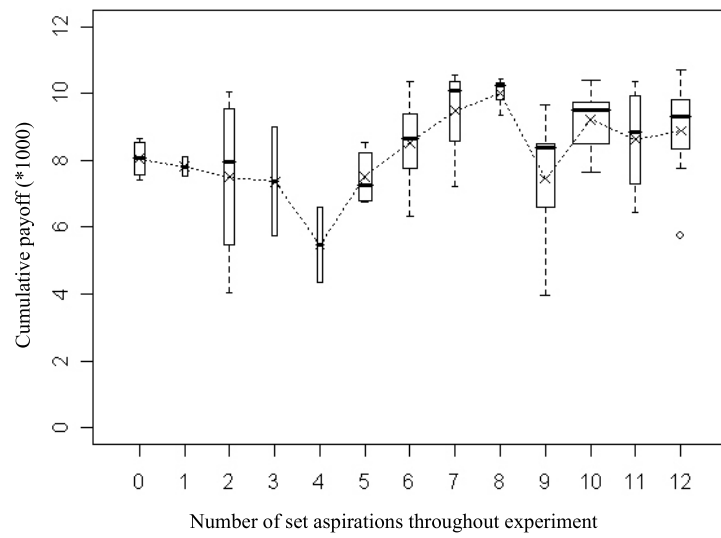


Figure 6: Relationship between use of aspirations and final payoff

Tables

Periods	Phase 1	Phase 2	Phase 3
Type 1	K	U	K
Type 2	U	K	K

Table 1: Experiment design

Search profile	Relative frequency	Search method duration			Total duration
		S	E	A	
EA	0.42	0	10.3	6.0	16.3
E	0.23	0	10.6	0	10.6
ES	0.08	2.8	10.1	0	12.9
SA	0.05	2.9	0	5.9	8.8
S	0.05	2.8	0	0	2.8
EAS	0.05	2.9	10.0	8.0	20.9
A	0.04	0.0	0	7.5	7.5
ESA	0.02	1.9	12.2	2.6	16.7
SE	0.01	1.4	9.4	0	10.8
Other	0.04	3.0	7.6	6.0	16.5

Table 2: Summary statistics, search profiles

<i>Coefficient</i>	<i>Estimate</i>	<i>Std.Error</i>	<i>p-value</i>
Intercept	-0.710	0.054	0.041
period	0.126	0.007	< 0.001
AIC: 815	BIC: 829	logL: -404	
$\gamma_{subject} \sim N(0, 4.937)$		$\varepsilon \sim N(0, 2.222)$	

Table 3: Determinants of forming an aspiration (mixed-effects logit model)

<i>Coefficient</i>	<i>Estimate</i>	<i>Std.Error</i>	<i>p-value</i>
Intercept	91.357	2.349	< 0.001
$\delta_{phase=2}$	-2.109	2.576	0.414
$\delta_{phase=3}$	-8.389	2.566	0.001
AIC: 3522	BIC: 3541	logL: -1756	
$\gamma_{subject} \sim N(0, 8.919)$		$\varepsilon \sim N(0, 19.747)$	

Table 4: Determinants of relative aspiration level (linear mixed-effects model)

<i>Coefficient</i>	<i>Estimate</i>	<i>Std.Error</i>	<i>p-value</i>
Intercept	1.275	0.865	0.141
\underline{n}	0.283	0.019	< 0.001
\bar{n}	0.121	0.021	< 0.001
$\delta_{phase=2}$	0.436	0.339	0.199
$\delta_{phase=3}$	-1.320	0.355	< 0.001
AIC: 2018	BIC: 2046	logL: -1002	
$\gamma_{subject} \sim N(0, 4.067)$		$\varepsilon \sim N(0, 2.495)$	

Table 5: Duration of aspiration formation phase (linear mixed-effects model)

Search profile	Rel. frequency in periods			(Median) results		
	1-4	5-8	9-12	Search cost	Payoff	Efficiency
EA	0.35	0.45	0.46	40	845	0.89
E	0.25	0.24	0.21	20	839	0.86
EAS	0.02	0.04	0.08	100	744	0.82
ES	0.12	0.09	0.05	60	779	0.86
SA	0.05	0.06	0.05	70	777	0.86
S	0.06	0.04	0.05	40	494	0.52
A	0.07	0.02	0.04	20	128	0.16
ESA	0.02	0.03	0.01	100	824	0.91
SE	0.02	0.01	0.00	50	690	0.75
Other	0.04	0.03	0.05	120	656	0.76

Table 6: Efficiency and relative frequency of search profiles

<i>Coefficient</i>	<i>Estimate (lme)</i>	<i>Std.Error</i>	<i>p-value</i>	<i>Estimate (qr. $\tau = 0.5$)</i>
Intercept	80.874	2.936	< 0.001	92.150
period	-0.929	0.310	0.003	-1.288
$\delta_{period=1}$	-23.827	3.867	< 0.001	-25.875
$\delta_{aspir=1}$	4.487	2.286	0.050	6.413
AIC: 7249	BIC: 7277	logL: -3619		
$\gamma_{subject} \sim N(0, 11.066)$		$\varepsilon \sim N(0, 25.886)$		

Table 7: Determinants of search efficiency (linear mixed-effects model)