



Schunk, Daniel und Winter, Joachim:

The Relationship Between Risk Attitudes and Heuristics in Search Tasks: A Laboratory Experiment

Munich Discussion Paper No. 2007-9

Department of Economics
University of Munich

Volkswirtschaftliche Fakultät
Ludwig-Maximilians-Universität München

Online at <https://doi.org/10.5282/ubm/epub.1377>

The Relationship Between Risk Attitudes and Heuristics in Search Tasks: A Laboratory Experiment*

Daniel Schunk**

University of Mannheim

Joachim Winter†

University of Munich

This version: February 22, 2007

Abstract: Experimental studies of search behavior suggest that individuals stop searching earlier than predicted by the optimal, risk-neutral stopping rule. Such behavior could be generated by two different classes of decision rules: rules that are optimal conditional on utility functions departing from risk neutrality, or heuristics derived from limited cognitive processing capacities and satisficing. To discriminate among these two possibilities, we conduct an experiment that consists of a standard search task as well as a lottery task designed to elicit utility functions. We find that search heuristics are not related to measures of risk aversion, but to measures of loss aversion.

Keywords: search; heuristics; utility function elicitation; risk attitudes; prospect theory

JEL classification: D83; C91

* The authors would like to thank Daniel Houser, Oliver Kirchkamp, Joerg Oechssler, Robert Sugden, Martin Weber, Matthias Weiss, and seminar participants at the University of Mannheim, George Mason University's ICES, the 2004 meetings of the Eastern Economic Association, the 2004 Annual Conference of the European Network for Training in Economics Research (ENTER), as well as the 2004 North American Meeting of the Economics Science Association (ESA) for their comments. Financial Support from the Deutsche Forschungsgemeinschaft (DFG) via the Sonderforschungsbereich 504 at the University of Mannheim is gratefully acknowledged.

** Corresponding author.

Address: Sonderforschungsbereich 504

University of Mannheim

D-68131 Mannheim, Germany

E-mail: dschunk@uni-mannheim.de

Web: www.mea.uni-mannheim.de/daniel-schunk

Fax: +49-621-181-3451

† *Address:* Department of Economics

University of Munich

D-80539 Munich, Germany

E-mail: winter@lmu.de

E-mail: www.vwl.uni-muenchen.de/lm_empwi/Personen/winter/winter.htm

Fax: +49-89-2180-3954

The Relationship Between Risk Attitudes and Heuristics in Search Tasks: A Laboratory Experiment

1 Introduction

Behavior in search situations receives much attention in various fields of economics, such as labor economics and marketing science.¹ But behavior in search tasks is interesting not only for the analysis of many substantive issues, it has also proven a useful object of behavioral research in psychology and economics. Search tasks are attractive for experimental studies because of their (superficially) simple structure that masks an underlying optimization problem that is quite complicated and in most cases cannot be solved in a human's mind but requires numerical methods and a computer. Conceptually, search tasks are representative of many situations in which one has to decide between committing resources to an attractive proposition or deferring the decision in the hope of receiving a better deal.

Price search situations and variations thereof – such as the well-known secretary problem, in which the decision whether to stop or to continue depends only on the relative ranks of the presented alternatives (e.g., Rapoport and Tversky, 1970; Seale and Rapoport, 1997) – have been investigated theoretically and empirically by numerous authors, e.g., Stigler (1961), Braunstein and Schotter (1982), Hey (1981, 1982, 1987), Kogut (1990), Harrison and Morgan (1990), Schotter and Braunstein (1981), Sonnemans (1998, 2000), and Houser and Winter (2004). Since individual price search behavior is difficult to examine in the field, research on price search is generally based on experimental studies. The existing experimental evidence suggests that people are very heterogeneous in their search behavior and that relatively simple heuristics describe observed behavior better than the optimal stopping rule. It has been found, though, that subjects' search behavior is nearly optimal in the sense that their actual earnings are close to the earnings they would have realized had they followed an optimal strategy. This observation, however, does not indicate that their stopping rule is necessarily close to the optimal rule – it could also be that the payoff

¹ See Eckstein and van den Berg (2007) and Zwick *et al.* (2003) for reviews of the literature in these fields.

to search tasks is not very sensitive to deviations from the optimal stopping strategy (see Harrison and Morgan, 1990; Seale and Rapoport, 1997, 2000). Overall, while people seem to behave as predicted by theory when parameters of the search environment change (e.g., Schotter and Braunstein, 1981), experimental findings in various search contexts suggest that individuals tend to search too little relative to the optimal strategy (Hey, 1987; Cox and Oaxaca, 1989; Houser and Winter, 2004; Seale and Rapoport, 2000; Sonnemans, 1998). Cox and Oaxaca suggest that this might be traced back to risk-averse behavior of the individuals (Cox and Oaxaca, 1989). Using an electronic information board method, Sonnemans (1998) finds that differences in learning behavior of the subjects might also be responsible for the observation of early stopping.

The existing experimental literature on search behavior is based on the assumption of risk neutrality. Under risk neutrality, optimal stopping rules can be derived, and experimental studies typically find that most subjects do not use such rules but rather follow some heuristic. These heuristics are often sophisticated in the sense that they allow subjects to get quite close to the payoffs they would have obtained using optimal rules. However, once one allows for heterogeneity with respect to the individual risk attitudes, the situation is more complicated: Decision rules that have been treated as heuristics in the literature could, in fact, be optimal conditional on the individual risk attitude. Consequently, search behavior that cannot be explained by the optimal stopping rule derived under risk neutrality could be generated by two entirely different classes of decision rules: (i) rules that are optimal conditional on the individual utility function or (ii) heuristics that derive, say, from satisficing or other cognitive processes. Distinguishing these two possibilities requires an independent measure of risk attitudes.

The contribution of our paper to the search literature is, therefore, to study the relation between properties of subjects' preferences (specifically, measures of risk attitude) and decision rules used in search tasks. We do this by presenting subjects not only with a search task that follows the standard in the literature, but also with a lottery task that serves to elicit subjects' individual utility functions. In addition, we use a questionnaire to obtain a psychometric measure of risk attitudes as an independent individual-level source of information on risk behavior.

In section 2, we present the design of our experiment. Section 3 describes our procedures to draw inferences on subjects' search behavior and risk attitudes. In section 4, we link these elements and discuss the results of our experiment. Section 5 concludes.

2 Design and Administration of the Experiment

Our experiment consists of three parts (A, B, and C) that were presented to the subjects in fixed order. Part A of the experiment serves to elicit features of subjects' preferences, namely, the shape of their utility functions in the gain and loss domains. Part B consists of a series of repeated price search tasks that is used to identify subjects' search heuristics. Part C is a survey instrument developed in the psychology literature to generate a measure of subjects' risk behavior. We describe these three parts in turn.

2.1 Parts A: Preferences

Part A builds on a method recently proposed by Abdellaoui (2000). A series of lottery tasks serve to elicit subjects' utility and probability weighting functions in a parameter-free way. In part A, we elicit each subject's utility function on the gain and loss domain, using a series of 64 lottery choice questions in total. Four of the lottery questions appear twice during the lottery elicitation process. This gives us the possibility to investigate whether subjects behave consistently during the utility elicitation questions, or whether preference reversals have occurred.²

The experiment used in Part A for the elicitation of subjects' utility function is based on the construction of "standard sequences of outcomes", i. e., monetary outcomes that are equally spaced in terms of utility. In our design, we use a 5-step bisection procedure to determine an outcome x_1 that makes the subject indifferent between two lotteries $A = (x_0, p; R, 1 - p)$ and $B = (x_1, p; r, 1 - p)$; where p is set to $2/3$ and we have $0 \leq r <$

² In our experiment, we also elicited each subject's probability weighting functions for gains and losses through a series of 72 lottery choice questions. Since subjects' probability weighting functions are not of interest in this study, we do not discuss results from these additional lottery tasks. The results from our probability weighting function elicitation are comparable to the results reported by Abdellaoui (2000); in particular, our estimates of the shape of the probability weighting function are similar to those obtained by Abdellaoui. The results of the probability weighting function elicitation part of the experiment can be obtained from the authors upon request.

$R < x_0 < x_1$. The parameters r , R , and x_0 are held fixed during the whole experiment. The first 5 presented lottery-pairs let us determine the desired x_1 that makes the subject indifferent between the lotteries A and B, see the Appendix for an example of the sequence of lotteries. The next step of this procedure is to present another 5 pairs of lotteries in order to determine a value x_2 that makes the subject indifferent between the lotteries $(x_1, p; R, 1 - p)$ and $(x_2, p; r, 1 - p)$. This procedure continues until we have determined an x_6 . In our experiment, we set (in the gain domain) R to €100, r to €0, and x_0 to €200. In the loss domain, we use the negative of these values.

Now, assume that preferences can be represented by cumulative prospect theory (CPT).³ Let $u(\cdot)$ denote the utility function on the gain or the loss domain and let $w(\cdot)$ denote the probability weighting function for the respective domain. Then indifference between two lotteries implies pairs of equations of the following type:

$$w(p)u(x_i) + (1 - w(p))u(R) = w(p)u(x_{i+1}) + (1 - w(p))u(r) \quad (1)$$

$$w(p)u(x_{i+1}) + (1 - w(p))u(R) = w(p)u(x_{i+2}) + (1 - w(p))u(r) \quad (2)$$

From these two equations follows:

$$u(x_{i+1}) - u(x_i) = u(x_{i+2}) - u(x_{i+1}) \quad (3)$$

That is, in terms of utility, the trade-off of x_i for x_{i+1} is equivalent to the trade-off of x_{i+1} for x_{i+2} . This method yields a standard sequence of outcomes, $\{x_0, x_1, \dots, x_6\}$, which is – by construction – increasing for gains and decreasing for losses.⁴ Note that the range of monetary outcomes in the elicitation procedure is specific for each subject, since it depends on individual decisions.

2.2 Part B: Search Behavior

In part B of the experiment, subjects perform a sequence of search tasks. Each subject's goal is to purchase an object which they value at €500. This article is sold at infinitely

³ The elicited utility function on gains is, indeed, a von-Neumann-Morgenstern utility function. Equation (3) holds also under Expected Utility Theory, as can be found by substituting p for $w(p)$ in equations (1) and (2).

⁴ A standard sequence $\{x_0, x_1, \dots, x_n\}$ requires the construction of n such indifferences $(x_i, p; R, 1 - p)$ and $(x_{i+1}, p; r, 1 - p)$.

many locations, and visiting a new location costs €1. At each location, a price is randomly drawn from a known distribution. On the instruction sheet, subjects are informed graphically and verbally that the price at each location is drawn independently from a truncated normal distribution with a mean of €500, a standard deviation of €10, and truncation at €460 and €540. The distribution is discretized such that only integer prices are realized.

After each new price draw (that is, at each location they visit), subjects are allowed to recall previously rejected price offers. That is, after each price draw, subjects can stop and choose any price (location) encountered so far, or they can continue their search at the incremental cost of another euro. The outcome of each search task is calculated as the evaluation of the object (€500) minus the price at the chosen location minus the accumulated search cost.

Note that we allow for recall in order to be closer to situations such as price-search in the internet: Indeed, in real-world situations, individuals can often perform their search and compare offers as long as they want; at a certain moment, they decide to stop their search and choose one of the offers that they have come across during their search.

Conceptually, the search problem presented in Sonnemans (1998, 2000) and Schunk (2006) is similar to our search task: The number of searches is unlimited, recall is accepted, the costs of one search action is constant and the price at each location is drawn independently from a distribution that is known to the searcher. In contrast to our setup, however, the price offers are drawn from a discrete uniform distribution in Sonnemans' experiments on search. In Hey (1982), the subjects also face an identical situation; however they do not know that the distribution of prices is normal (without truncation).

To ensure that subjects were experienced with the task and comfortable with the computer interface, and to minimize the impact of learning, subjects were allowed to perform an unlimited number of practice search tasks before performing a sequence of 10 or 11 tasks that determined their payoff for part B of the experiment.⁵ Finally, after the experiment was completed, one of these rounds was selected randomly to determine the part B pay-off.

⁵ 35 subjects played ten search rounds, and the second half, another 33 subjects, played 11 payment-relevant search rounds.

2.3 Part C: Risk Attitudes

The experiment ends with a short computerized questionnaire (part C). This survey instrument for assessing risk-taking was developed by Weber *et al.* (2002). Subjects rate their behavior with respect to 4 risky activities in the behavioral risk domain of gambling. Specifically, subjects report how likely it is that they engage in a certain gambling-related activity on a five-point rating scale ranging from 1 (“Extremely likely”) to 5 (“Extremely unlikely”)⁶.

Risk attitude is generally considered to be domain-specific in recent psychological literature (e. g., Bromiley and Curley, 1992). Based on our questions, we have a psychometric measure for individual risk attitude in the gambling domain. In our subsequent analysis, we correlate these measures with measures of risk attitudes obtained using the lottery tasks of part A and with behavior in the search tasks observed in part B of the experiment.

2.4 Administration

The study was conducted in the fall of 2003 in the experimental laboratory of Sonderforschungsbereich 504, a research center at the University of Mannheim. In four sessions, a total of 68 subjects participated in the main study.⁷ These subjects were recruited from the general student population. All experiments were run entirely on computers using software written by the authors.

All payments were made after subjects had completed all parts of the experiment. For each subject, the outcome of one of the 10 or 11 payment-relevant search tasks in part B was selected randomly, and added to or subtracted from a flat €8 show-up fee, depending on whether it was a gain or a loss. Subjects were told that their total payoff was truncated at €0.⁸ That is subjects would not suffer a loss from the experiment, they would at least

⁶ Based on subjects’ ratings in the risk domains (i) financial, (ii) recreational, (iii) social, (iv) health/safety, and (v) ethical., Weber *et al.* (2002) construct domain-specific scales of subjects’ risk attitudes and evaluate the construct validity and the consistency of these scales using standard approaches. However, for our purpose, only the domain of gambling risk is of interest.

⁷ A separate group of 5 subjects participated in a pilot study which allowed for fine-tuning of the parameters of the lottery and search tasks, the adjustment of the software, and optimization of the experimental protocol.

⁸ The lowest payoff that was paid in all sessions was €4, so no subject was forgiven any losses.

earn €0 from the experiment. Finally, one of the (on average) 17 subjects participating in each experimental session was randomly selected to play for a real monetary pay-off based on his or her choices made in one of the lottery tasks in parts A of the experiment; answers were collected as binary choices between two prospects, i.e. only the preferred lottery was played for real pay-off. Since the outcomes of the lotteries were up to €6000, we informed the subjects that the randomly selected person played for only 1% of the positive outcomes (i. e., the gains) presented in the lotteries.

3 Inference on Search Heuristics and Risk Attitudes

In this section, we discuss how we use the data from our experiment to draw inferences on subjects' preferences (the shape of their utility functions, as revealed in the lottery tasks) and behavior (the heuristics they use in solving the search task). The last subsection briefly explains how the psychometric measures of subjects' risk attitudes are obtained.

3.1 Estimation of the Shape of the Utility Function

As mentioned in section 2.1, the lottery tasks presented in part A of our experiment are based on those developed by Abdellaoui (2000). He uses his experimental data to estimate utility functions in the gain and loss domain as well as the corresponding probability weighting functions nonparametrically. For the purpose of our study, we need to order subjects according to their risk attitudes. We therefore use a parametric approach and specify the subjects' risk attitude based on the functional specification of a utility function with constant absolute risk aversion form (CARA). We estimate the utility function in the gain and loss domains separately using nonlinear least squares and the data from part A of the experiment.

We should point out that the procedure we use to elicit the shape of the utility function (Part A) operates on a monetary range of gains and losses that is different from the range considered in the search experiments (Part B). We made this decision on purpose, and we digress here for a brief discussion of the rationale for this decision. As pointed out by Wakker and Deneffe (1996), the curvature of the utility function is more pronounced if a

sufficiently wide interval of outcomes is investigated. Accordingly, our adaptive method elicits individuals' utility functions for monetary outcomes in a wide interval (the size of which depends on the subjects decisions, see Abdellaoui (2000)) below €-200 or above €200, respectively. In the search game, where actual payments were made, we had to reduce the outcome scale between €-40 and €40 because of budget limitations. It may well be the case that individual risk attitudes are different for high and low monetary outcomes. However, all we need for our empirical analysis is that the rank order of individuals by the measures of risk attitudes is preserved between the high-outcome range for which it is elicited and the low-outcome range that is relevant for the analysis of behavior in the search game.⁹ This is, in our view, a reasonable assumption and, in fact, a corollary of using a CARA-utility specification. Furthermore, using data from high and low outcome risk elicitation tasks by Holt and Laury (2002), this assumption can be investigated. The results are supportive, and are presented in section 5 of this paper.

Based on, e. g., Currim and Sarin (1989) and Pennings and Smidts (2000), we assume the following exponential specification for our CARA-utility function on gains¹⁰:

$$u(x) = \frac{1 - e^{-\gamma(x-x_{min}^G)}}{1 - e^{-\gamma(x_{max}^G-x_{min}^G)}} \quad (5)$$

Here, x_{max}^G is the largest elicited value of x in the gain domain (in absolute values), i. e., x_6 ; x_{min}^G is the smallest elicited x -value on the gain domain, i. e., x_0 . For obtaining the utility function in the loss domain, we replace x_{max}^G and x_{min}^G by x_{max}^L and x_{min}^L , respectively, we use the absolute value of the denominator and the numerator and we take the negative of the right-hand side. For $\gamma = 0$ the function is defined to be linear, i. e., the subject is risk-neutral.

In our specification, the coefficients are estimated separately for gains and losses (γ and δ , respectively). These coefficients characterize each subject's risk attitude in the sense

⁹ Accordingly, our empirical analysis will only be based on rank correlations, that is "comparative risk aversion".

¹⁰ Note that another normalized version of the CARA-utility has the following form:

$$u(x) = \frac{1 - e^{-\gamma \frac{x-x_{min}^G}{x_{max}^G-x_{min}^G}}}{\gamma} \quad (4)$$

Fitting this function yields a significantly higher mean relative standard error of the coefficient estimate and a significantly lower coefficient of determination than fitting the functional form in equation (5). The substantive conclusions of our analysis remain unchanged when we use the form (4).

of an Arrow-Pratt-measure (Pratt, 1964) of risk attitude, that is $-u''(x)/u'(x) = \gamma$ for gains and $-u''(x)/u'(x) = \delta$ for losses. If $\gamma < 0$, the subject has a convex utility function and is risk-seeking on gains, if $\gamma > 0$, the subject is risk-averse, her utility function on gains is concave.

Furthermore, we calculate an individual-specific index for loss aversion from our data. Because subjects generally evaluate their choice options relative to salient reference points, Tversky and Kahneman propose that individuals process losses differently than gains (Tversky and Kahneman, 1992). That is, loss aversion can be considered a psychological factor, capturing the trade-off between gain- and loss-utility units. Generically, loss aversion is defined by $u(x) - u(y) \leq u(-y) - u(-x)$ for all $x > y \geq 0$ (Schmidt and Traub, 2002).

Based on work by Benartzi and Thaler (1995), Koebberling and Wakker (2005) propose an index of loss aversion that is – in contrast to other indices discussed in the literature – invariant to changes in the scale of the utility function, $u(\cdot)$, and it is invariant to scale transformations of the outcomes. This index is given by

$$\lambda = \frac{u' \uparrow(0)}{u' \downarrow(0)}. \quad (6)$$

Based on Koebberling and Wakker (2005) and our utility elicitation procedure, the index of loss aversion has the following form for $\gamma \neq 0$ and $\delta \neq 0$:

$$\lambda = \frac{\frac{\delta \cdot (e^{-\delta \cdot (|x^L_{max} - x^L_{min}|)})}{1 - e^{-\delta \cdot (|x^L_{max} - x^L_{min}|)}}}{\frac{\gamma \cdot e^{-\gamma \cdot (x^G - x^G_{min})}}{1 - e^{-\gamma \cdot (x^G_{max} - x^G_{min})}}} \quad (7)$$

For $\gamma = 0$, we have $u' \downarrow(0) = \frac{1}{x^G_{max} - x^G_{min}}$, for $\delta = 0$, we have $u' \uparrow(0) = \frac{1}{|x^L_{max} - x^L_{min}|}$ entering expression (7) in the denominator and numerator, respectively.

Note that our estimate of individual loss aversion is based on the assumption that the estimated form of the CARA-utility function of the individual is characteristic for her utility function over the whole domain, and identically scaled both on gains and on losses¹¹.

¹¹ Our estimates of loss aversion are based on the assumption that the combination of our utility elicitation method and Koebberling and Wakker's (2005) index for loss aversion yields a reasonable overall estimate of comparative individual loss aversion. Our findings on psychometric risk attitudes, reported later, support this claim. We acknowledge, however, that methods for the elicitation of an index of loss aversion that based on mixed lotteries (e. g., Schmidt and Traub, 2002), though also suffering from considerable uncertainty, could also be used in the context of a search experiment; Schunk (2006) uses

3.2 Classification of Decision Rules Used in the Search Task

The next step of our analysis is to determine, for each subject, the decision rule he or she uses in the search task. We specify a fixed set of candidate decision rules, comprised of the optimal decision rule and several simple heuristics that have been used in the earlier literature (e.g., Hey, 1982; Moon and Martin, 1990; Houser and Winter, 2004) to describe search behavior. For each subject and each candidate decision rule, we compute the number of stopping decisions that are correctly predicted. We assign to the subject the decision rule that generates the largest fraction of correct predictions, i.e., that fits observed behavior best. We start this subsection with the derivation of decision rules both under risk neutrality and without restrictions on individual risk attitude. The derivations are based on the assumption of a classical von-Neumann-Morgenstern utility function which is only defined on monetary gains since our experimental design implies that searchers cannot suffer a loss from the experiment. For the derivation of the decision rules, we consider two cases: In the first case, the cost of each completed search step are treated as sunk cost; in the second case, we derive the finite horizon optimal stopping rule assuming that subjects do not treat past search cost as sunk costs. Finally, we discuss the set of alternative heuristics, and describe our classification procedure more formally.

Stopping Rules in Search Tasks under Risk Neutrality

Assume that the searcher observes sequentially any number of realizations of a random variable X which has the distribution function $F(\cdot)$. In our case, $F(\cdot)$ is a discrete truncated normal distribution with mean €500 and standard deviation €10, the truncation is at €460 and €540. Let the cost of searching a new location be c . Assume that at some stage in the search process, the minimal value that the searcher has observed so far is m , and the searcher wonders whether to continue searching or whether to stop the search. Basic search theory assumes that individuals treat the cost of each search step,

mixed lotteries in an experimental study on search behavior that uses a different design to elicit individual preferences. We suggest that further experimental studies investigate the relationship between loss aversion indices derived from mixed lotteries (e.g., Schmidt and Traub, 2002) and indices derived from outward methods and pure lotteries, such as the method applied in the present paper.

once completed, as sunk costs (Lippman and McCall, 1976; Kogut, 1990) and compare the payoff of one additional search step with the payoff from stopping.¹²

Then, subjects solve the problem based on a one-step forward-induction strategy and the expected gain from searching once more before stopping in a search task such as ours, $G(m)$, is generally given by:¹³

$$G(m) = - \underbrace{[1 - F(m)]m}_{\otimes} - \underbrace{\int_{460}^m x dF(x)}_{\oplus} - c + m. \quad (8)$$

The term \otimes accounts for the case where a value larger than m is found with probability $(1 - F(m))$. In this case, m remains the minimum price. The term \oplus stands for the case where we find a lower value than m and calculates the expected value in this case. After some manipulation, we obtain the following condition for the parameter values of our search task,

$$G(460) = -c < 0. \quad (9)$$

That is, it does not make sense to continue searching if one draws the minimal value of €460. In our specification, the highest price that can be drawn is €540. In this case, the expected gain from searching at least one more time is always positive (since payoffs cannot become negative), so

$$G(540) > 0. \quad (10)$$

From these properties of $G(\cdot)$, it follows that there exists *a unique* value at which $G(\cdot) = 0$. We denote this value by m^* and solve equation (8) for m^* . Straightforward manipulation shows that the solution to this problem is identical to solving the following problem for m :

$$\pi(500 - m + 8) = (1 - F(m))\pi(500 - m - c + 8) + \int_{460}^m \pi(500 - x - c + 8)dF(x) \quad (11)$$

¹² Kogut's (1990) findings show that a certain proportion of subjects does not treat sunk costs as sunk.

¹³ Note that the one-step forward induction strategy is identical with the optimal solution of the infinite horizon problem if the searcher is risk-neutral.

Here, $\pi(\cdot)$ is the payoff-function from the search game and the show-up fee of €8 is included in this equation, since subjects' payoff from the search game is directly linked to the show-up fee. $\pi(\cdot)$ has the following form:

$$\pi(x) = \max\{0, x\} \quad (12)$$

In equation (11), the left-hand side of the equation is the payoff from stopping and the right-hand side denotes the payoff from continuing search. We find that the optimal strategy is to keep searching until a value of X less than, or equal to, the optimal value m^* has been observed. In our problem, we find that $m^* = 490$. That is, we have the following optimal decision rule for a risk-neutral searcher: Stop searching as soon as a price less than or equal to €490 is found.

Now, consider that subjects do not treat search costs as sunk costs. That is, for their decision whether to stop or to continue the search, they consider the total benefits and costs of search; the agent stops searching only if the stopping value is higher than the continuation value. In this case, subjects would not search for more than 48 steps since after 48 search steps the continuation value from the experiment would definitely be zero. It follows that the problem is treated as a finite horizon problem that is solved backwards. Define $S_t = \{t, m\}$ as the agents' state vector after making t search steps.

After the agent has stopped searching, she will buy the item and receive a total payoff of:

$$\Pi(S_t) = \max\{0, 500 - m - t \cdot c + 8\}. \quad (13)$$

Now, the agent stops searching only if the continuation value of search is lower than the stopping value. The recursive formulation of the decision problem is therefore:

$$J_t(S_t) = \max\{\Pi(S_t), E[J_{t+1}(S_{t+1})|S_t]\}. \quad (14)$$

$E(\cdot)$ represents the mathematical expectations operator, and the expectation is taken with respect to the distribution of $S_{t+1}|S_t$. Again, this problem has, at every t , the reservation price property. The reservation price begins at 490, then starts decaying slowly, reaches 483 in the 24th round and then decays at a rate of about one per round from that point forward.

Stopping Rules in Search Tasks Without Restrictions on Risk Attitudes

The derivations above are based on the assumption of a risk-neutral searcher. Sonnemans (1998), for example, refers to a model of the form (8) as an optimal stopping rule. Houser and Winter (2004) refer to a model of the form (14) as an optimal stopping rule. Note, however, that it is individually rational to use the risk-neutral optimal stopping rule only for risk-neutral subjects. Put differently, observing a subject that does not follow the optimal stopping rule derived under risk neutrality does not necessarily imply that his or her search is not rational.

As a more general case, we therefore consider a searcher with an arbitrary, monotone utility function $u(\cdot)$. If the searcher ignores sunk cost and takes her decisions based on a one-step forward-looking strategy, the equation that determines her reservation price m^* has the following form, which is an immediate extension of equation (11)¹⁴:

$$u(500 - m + 8) = (1 - F(m))u(500 - m - c + 8) + \int_{460}^m u(500 - x - c + 8)dF(x) \quad (15)$$

Equation (15) can be solved numerically for the reservation price $m^*(\eta)$, given a specific price distribution, search costs, and a utility function on gains that is characterized entirely by a parameter η . The problem has the constant reservation price property, which is reported as a search heuristic that is consistent with the behavior of a reasonable number of subjects in other studies (e.g., Hey, 1987). Figure 1 shows the constant reservation price as a function of the risk-parameter γ in the exponential utility function (5). Note that the reservation price $m^*(\eta)$ is invariant to changes of scale of the utility function. Henceforth, we will refer to rules of this type as forward optimal rules, keeping in mind that this rule is only optimal conditional on the individual utility function and on the assumption of a one-step forward strategy that ignores sunk costs.

Analogous to our derivation of the optimal search rule in the risk-neutral case, we now consider the case in which subjects do not treat search costs as sunk costs. Again, we

¹⁴ Note that this equation does *not* characterize the optimal solution to the search problem. It gives, however, the optimal strategy for a searcher with arbitrary risk-attitude who ignores sunk costs and who uses a one-step forward induction strategy.

have a finite-horizon problem that is solved using backward induction. After the agent has stopped searching, she will buy the item and receive a total payoff of:

$$\Pi^u(S_t) = \max\{0, u(500 - m - t \cdot c + 8)\}. \quad (16)$$

The agent stops searching only if the utility of continuing the search is lower than the utility from stopping. The recursive formulation of the decision problem is:

$$J_t^u = \max\{\Pi^u(S_t), E[J_{t+1}^u(S_{t+1})|S_t]\}. \quad (17)$$

Again, this problem has, at every t , the reservation price property. The monotonically falling reservation price for all arbitrary values of γ implies that the agent should not exercise recall. Figure 2 plots the path of reservation prices, calculated by solving the dynamic discrete choice problem implied by equation (17) for various risk attitudes γ of the individual. Henceforth, we will refer to rules of this type as backward optimal rules, which are optimal conditional on the individual utility function. From our theoretical deliberations so far we can conclude that – regardless of what type of optimal rule subjects use, forward or backward optimal rules – risk averse subjects should stop their search earlier, i. e., they have higher reservation prices on average, and risk-seeking subjects should stop their search later, that is they use lower reservation prices.

Alternative Search Rules

As has been pointed out in the search literature before, and as should have become clear in the previous sections, computation of the optimal search rule (either under risk neutrality or without restrictions on the risk attitude) is a demanding task, and it is unlikely that subjects can perform this task during a search experiment (or in real-life search situations, for that matter). Most papers in search literature therefore argue that subjects use heuristics rather than the optimal stopping rule, and there is some evidence that certain heuristics get subjects close to the pay-offs they could have obtained using the optimal rule.

We now specify our set of candidate search rules that are used in this paper to characterize behavior in experimental search tasks. In addition to the search rules that have been derived in the section above, we specify a set of heuristics that have been used in the

search literature to characterize behavior in experimental search tasks. These heuristics are based on experimental work by Hey (1982) and Moon and Martin (1990).

The first class of these decision rules comprises several “sophisticated” heuristics. These heuristics share the constant reservation price property. Each rule says that the subject uses an arbitrary, but constant reservation value $r \in \{480, \dots, 500\}$. Subjects behaving according to this heuristic search until a price quote lower than or equal to the reservation price is found. We refer to this constant reservation type of heuristic as *type 1 heuristics*. Note that this heuristic is identical to the forward optimal search rule, see above. Based on this rule, we attribute to every individual the constant reservation price value that explains most of her observed search decisions.¹⁵

The second class of decision rules that we consider are based on the finite horizon search model, i. e., the backward optimal search rules, as specified above. According to these search rules, subjects use a reservation price that is a function of the search step t and of the individual risk attitude γ that characterizes the utility function for which the search rule has been derived. Here, we consider that $\gamma \in \{-1.0, -0.95, -0.9, \dots, +0.95, +1.0\}$. We refer to this class of decision rules as *type 2 rules*. Based on this rule, we attribute to every individual a value γ_i^{search} , the risk-attitude coefficient that explains best the observed search behavior.

A third class of heuristics is also based on reservation prices that vary over the search time. Subjects using one of these heuristics stop searching as soon as their payment exceeds a certain individual threshold (or satisfaction-) level $t \in \{1, \dots, 20\}$. Given our parametrization of the problem, this results in a reservation price that linearly falls over time. For obvious reasons, this heuristic is sometimes called the “satisficer heuristic” and we refer to it as *type 3 heuristics*.

As *type 4 heuristics*, we consider the so-called “bounce rules”, suggested by Moon and Martin (1990) based on earlier work by Hey (1982). Subjects following the “one-bounce rule” (heuristic 4a) have at least 2 searches and they stop if a price quote is received larger

¹⁵ As is clear from the solution to equation (15), each constant reservation price used in the price search problem is consistent with a certain (interval of) value(s) of the individual utility risk coefficient γ in the gain domain. Instead of attributing a constant reservation price to the people, we could as well attribute the value of γ that corresponds to this constant reservation price.

than the previous quote. The “modified one-bounce rule” (heuristic 4b) is similar to the one-bounce rule, but an agent following this rule stops only if a price quote is received larger than the previous quote less the search cost.

Finally we consider rules that are based on winning streaks (*type 5 heuristics*). Subjects who follow this type of heuristics stop searching if they receive two (heuristic 5a) or three (heuristic 5b) consecutive price draws that are below some fixed threshold level $p \in \{485, \dots, 500\}$. That subjects might use these streak-based rules in search situations can be motivated by results from psychological research on behavior in uncertain environments, see Rabin (2002).

We should note that the type 4 and 5 heuristics have also been used to describe behavior in search environments in which the distribution of prices is not known. In our environment, where subjects know the expected value and variance of the price distribution, using these rules makes less sense. *A priori*, we would therefore not expect that these heuristics are used frequently by our subjects.

Table 1 presents a summary of the 116 candidate decision rules (optimal stopping rule and heuristics) that we specify for the subsequent analysis.

Classification Procedure

Our approach to drawing inferences about search behavior is to determine, for each subject, the proportion of choices consistent with each decision rule and then to maximize this proportion over the set of all candidate decision rules. We assume that each subject follows exactly one of the decision rules in our universe of candidate rules and that he or she uses the same heuristic in each of the 10 or 11 pay-off tasks. This latter assumption seems reasonable in view of the fact that all subjects are experienced when they begin the pay-off tasks.

Formally, our classification procedure can be described as follows.¹⁶ Each heuristic $c_i \in \mathcal{C}$, where \mathcal{C} is the set of all search rules described above, is a unique map from subject i 's information set S_{it} to her continuation decision $d_{it} \in \{0, 1\} : d_{it}^{c_i}(S_{it}) \rightarrow \{0, 1\}$. Now, let

¹⁶ Houser and Winter (2004) implement a similar classification procedure in a completely specified maximum-likelihood framework.

d_{it}^* denote the observed decision of subject i in period t . Then, we can define the indicator function:

$$X_{it}^{c_i}(S_{it}) = 1(d_{it}^* = d_{it}^{c_i}(S_{it})) \quad (18)$$

Let T_i be the number of decisions that we observe for subject i . We attribute to each subject the heuristic that maximizes the likelihood of being used by that subject:

$$\hat{c}_i = \arg \max_{c_i \in \mathcal{C}} \sum_{t=1}^{T_i} X_{it}^{c_i}(S_{it}) \quad (19)$$

As we have motivated by reference to the existing literature, all relevant search heuristics should be included in our universe of 116 candidate decision rules. Based on our classification procedure, we attribute a decision rule to each subject, i. e., we can classify the subjects by the decision rules that they use. We can then investigate for each subgroup and for the whole sample the relationship between the observed search behavior and the risk preferences of the individuals.

3.3 Psychometric Measures

The questionnaire was constructed so that respondents evaluate their likelihood of engaging in an activity of the gambling-domain on a five-point rating scale ranging from 1 (“Extremely likely”) to 5 (“Extremely unlikely”). For each subject, we calculate a measure of risk attitude as the arithmetic mean score of the response to the four questions.

4 Results

This section starts with self-contained descriptions of both the results of the utility function elicitation (Part A) and the classification of the search behavior (Part B). We continue with a comparative analysis of our results on preferences and behavior (also including the psychometric measure of risk attitude).

In our experiment, 68 subjects participated in total. Of these 68 subjects, we delete four subjects from the sample.¹⁷ These 4 subjects apparently did not take the utility elicitation part of the experiment seriously.

The 64 subjects that we keep in the sample show a preference reversal rate of 21.9% on gains and 23.4% on losses in the utility function elicitation part of the experiment.¹⁸

4.1 Part A: Preferences

In Table 2, we report the standard errors of the nonlinear least squares estimates for the risk coefficients γ and δ . Furthermore, we report the sum of the squared residuals (SSR) and the coefficient of determination R^2 . We see that the standard errors are reasonably low and that the coefficients of determination are close to 1 for our nonlinear regressions. The estimation results suggest that the risk coefficients are reliable measures that allow for a rank-ordering of individuals according to their risk-attitude. Our results support the hypothesis of diminishing sensitivity for gains and losses if we consider the whole sample. Similar to Abdellaoui (2000), who uses a different measure for the classification of subjects' risk attitudes, we see a preponderance of risk-averse subjects in the gain domain, and a preponderance of risk-seeking subjects in the loss domain. Overall, our results on individual preferences are consistent with the predictions of prospect theory (Tversky and Kahneman, 1992) and subsequent experimental work based on prospect theory.

¹⁷ Two of these subjects are outliers in terms of the time needed for the completion of the lottery questions: They needed less than 60 seconds for either the 32 lottery questions on gains or the 32 questions on losses – considerably less than the other participants in the experiment who needed at least 1 minute 41 seconds. The two other subjects are outliers in terms of the standard error of the coefficient estimates of the utility function: Their standard errors of the coefficient estimate is more than one standard deviation larger than the standard errors of coefficient estimates for all the other subjects, i.e. preference parameters are measured imprecisely. Additionally, these two subjects are the only ones in the sample that revealed preference reversals on *all* four consistency check questions (see the section 2.1).

¹⁸ The reversal rate is a measure for how consistent subjects behave in a certain utility elicitation mechanism. Our reversal rate is somewhat higher than the rate in Abdellaoui (2000), who finds an error rate of 17.9% on gains and of 13.7% on losses. Abdellaoui's overall error rate, including the probability weighting function elicitation part of the experiment, is 19%. However, our reversal rate is lower than that of Camerer (1989), who reports that 26.5% of the subjects reversed preferences.

4.2 Part B: Search Behavior

A natural starting point for the investigation of search behavior is to assume that all subjects use a heuristic of the constant reservation price type, i. e., a type 1 heuristic. The reservation value that has been attributed to each subject can be considered a proxy for whether subjects tend to be early stoppers or late stoppers: The higher the attributed reservation price, the earlier subjects stop.

Figure 3 shows the distribution of reservation prices in the sample of 64 subjects, obtained under the assumption that each subject follows a constant reservation price decision rule. We find that 55% of the subjects are classified as “early stoppers”, i. e., their attributed reservation price is higher than the risk-neutral optimal reservation price of €490. 3% use the risk-neutral optimal reservation value and 42% are “late stoppers” with a reservation price lower than €490. Furthermore, note that if subjects use the risk-neutral optimal reservation stopping rule with a reservation price of €490, they should stop, on average, after having seen 5.85 prices. We find that the mean number of observed price draws per round is 5.07. The preponderance of early stoppers relative to the risk neutral constant reservation price stopping rule confirms results from earlier experimental studies of search behavior (Hey, 1987; Cox and Oaxaca, 1989; Sonnemans, 1998).

Next, we classify subjects according to the decision rule they use in the search tasks (see Table 3). Figure 4 shows the number of subjects for whom a certain heuristic is a “best” heuristic (numbers in parentheses indicate the fraction of correctly explained choices for the particular subjects). We find that for the 13% of the subjects, a constant reservation price heuristic explains behavior better than all other heuristics, for 3% a type 2 rule (the optimal finite horizon rule) is better than all others, and for 16% a satisficer rule (type 3) explains more observations than all other rules. For 84% of all subjects one of the conditionally optimal rules (type 1 or type 2) is a best decision rule, for 63% we find that they use one of the optimal rules (type 1 or type 2) and do *not* use the satisficer-heuristic (type 3); in contrast, 37% of the subjects can be termed satisficers – this result is similar to Sonnemans (1998), who finds that about one third of the subjects’ behavior is most

consistent with a satisficer rule. However, for 47% of the subjects, we cannot distinguish between the use of a forward or a backward optimal search rule (type 1 or type 2).¹⁹

Compared to these figures, it may be somewhat astonishing that the bounce-rules (the type 4 heuristics) and the streak-heuristics (type 5 heuristics) perform rather poorly: In total, only 35.9% of the observed decisions are consistent with the one-bounce rule, 33.6% are consistent with the modified one-bounce rule; 38.5% of the decisions are consistent with a type 5a heuristic, and 39.4% with a type 5b heuristic. However, Hey (1982), who has proposed the one-bounce rules following individual tape recordings of the subjects, finds equally low levels of consistency in a search environment where the price distribution was unknown.

In summary, heuristics of type 1, type 2, and type 3 do reasonably well in describing observed behavior. However, for a certain proportion of the subjects, our data do not discriminate between the usage of type 1 or type 2 or type 3 decision rules.²⁰ As a result of these findings, we classify the 64 subjects into 4 categories, labelled C1, C2, C3, and C4, respectively:

C1 All subjects whose observed behavior is explained best by a type 1 heuristic (49 subjects).

C2 All subjects whose observed behavior is explained best by a type 2 heuristic (45 subjects).

C3 All subjects whose observed behavior is explained best by a type 3 heuristic (24 subjects).

¹⁹ Both, forward and backward optimal rules, have very similar reservation price paths that only differ after a considerable number of search steps, see Figures 1 and 2. Therefore, the reported weak discrimination between both types of rules does not come unexpectedly. Changes in the experimental design will not improve the discrimination between these two types of rules: (i) A decrease in the standard deviation of the price distribution decreases the number of search steps in which forward and backward rules are identical (for identical parameter γ). However, a decrease in the price distribution also leads to fewer search steps per individual (Hey, 1987), which then complicates discrimination. (ii) An increase in the search costs per step decreases the number of search steps in which forward and backward rules are identical (for identical parameter γ). However, an increase in the search costs also leads to fewer search steps per individual (Hey, 1987), which, again, complicates discrimination.

²⁰ Technically, the likelihood function is rather flat, although the different decision rules are asymptotically identified; see the discussion in Houser and Winter (2004).

C4 Subjects whose observed behavior is explained best by a type 1 or a type 2 heuristic, but not by a type 3 heuristic (40 subjects).

4.3 The Relationship Between Preference Parameters, Search Behavior, and Risk Attitudes

The first question we investigate is whether there is a relationship between the observed search behavior and the elicited individual preferences (i. e., the coefficient of risk attitude). From the theoretical considerations above, the hypothesis follows that (at least) for those subjects that are classified as users of one of the conditionally optimal search rules (type 1 or type 2 rules), there exists an association between their risk attitude observed in the utility function elicitation part of the experiment and their behavior in the search experiment. We also extend this type of analysis to the whole sample. That is, we implicitly assume that all subjects behave according to just one search rule, either a type 1, a type 2 or a type 3 rule. We should find for all subjects in the sample that risk averse subjects generally use a higher reservation price or have a higher value of γ^{search} ; subjects classified as risk seeking should be attributed a lower reservation price or a lower value of γ^{search} .

Since the assumption of a normal distribution of the observed individual parameters γ , δ and λ across subjects is clearly rejected, we base part of our analysis on Spearman rank correlation coefficients. The significance of the Spearman correlations is tested using the null hypothesis that the two variables under question are independent.

We focus on the key parameters that characterize individual search behavior, the attributed constant reservation price level (RP), the average number of search steps per search round (AS) and the search coefficient γ^{search} . According to the basic search model (15), we hypothesize that – at least for subgroup C1 – γ is positively correlated with RP and negatively correlated with AS. We further hypothesize that at least for subgroup C2, γ is positively correlated with γ^{search} and negatively correlated with AS. Furthermore, due to being derived from the same underlying utility functional, the attributed constant reservation price (RP) and the attributed γ^{search} are strongly positively correlated (Spearman- ρ : 0.946, p -value: 0.00); we should therefore expect the hypothesized correla-

tions also for the subgroup C4. Table 3 reports the corresponding Spearman correlation coefficients for all subgroups C1 through C4 and the whole sample.

Our data do not reject the hypothesis of independence between γ and the search parameters RP, AS, and γ^{search} for all subgroups and the whole sample. The finding from our correlation analysis in Table 3 is that the utility function based measures for risk attitude on gains and losses, do not exhibit any significant relationship with individual behavior in search problems. This holds true regardless of whether we impose the usage of one specific type of search rule (e. g., the one-step forward-optimal search rule) to all subjects, or whether we attribute to each subject the type of rule that describes best her behavior and then, consequently, only consider the respective subgroups of the sample. To further investigate this point, we classify the subjects according to their risk attitude γ as measured in the utility function elicitation part. t -tests under the assumption of different variances show that our hypothesis motivated above – that risk averse ($\gamma > 0$) subjects generally use higher reservation price levels (RP) than risk-seeking ($\gamma < 0$) subjects – cannot be confirmed: The null hypothesis of equal mean reservation price levels is clearly not rejected across all subgroups considered. Even stronger: The mean reservation price of risk seeking subjects is higher than the mean reservation price of risk-averse subjects across all subgroups and the whole sample.

We now consider the correlation between the *psychometric measure for risk attitude in the gambling domain* and *search behavior*. There is some evidence that people who dislike taking risks in the gambling domain tend to search less: For C2-subjects, we have a Spearman- ρ of 0.26 (p -value 0.087) and for C4-subjects a Spearman- ρ of 0.29 (p -value of 0.07) for the correlation between the measure for risk on gambling and the average number of search steps per round (AS).²¹

With respect to the relationship between the *utility function based risk measures* and the *psychometric risk measures*, we find that apart from the subgroup C4, the loss aversion parameter does correlate at least marginally with the psychometric measure for risk on gambling. If we consider the complete sample, we find a Spearman ρ of -0.32 and a p -

²¹ The corresponding Spearman- ρ and p -values for the C1- and C3-group and for the whole sample are 0.14 (0.337), 0.06 (0.76) and 0.16 (0.21).

value of 0.009 for the correlation between the loss aversion parameter and the psychometric measure for risk on gambling.²²

In summary, our data do not confirm our hypotheses on the relationship between utility function based measures for risk aversion and search behavior. However, in Table 3 we do find significant relationships between the attributed constant reservation price level (RP) and the loss aversion index λ derived from the utility function, as well as between the average number of search steps (AS) and the loss aversion index. These correlations are significant or at least marginally significant across all subgroups considered. For the whole sample, we find significant correlations between the loss aversion index and both, RP and AS. Across all subgroups, subjects with a higher degree of loss aversion tend to have a higher attributed reservation price and stop their search earlier. Additionally, subjects' reported attitude towards risky gambles is related to their loss aversion and to the average number of search steps that they perform: People who avoid gambles tend to have a higher degree of loss aversion and they tend to stop their search earlier.

5 Discussion and Conclusions

This study combines elements from different literatures in experimental and behavioral economics – a lottery-based experiment designed to elicit subjects' individual utility functions (in particular, to estimate an index of risk attitude) and a search experiment designed to reveal subjects' decision rules in a search task. These experiments are augmented with a psychometric survey instrument that generates domain-specific measures of risk attitudes. We should first point out that the results of each of these components are broadly in line with earlier results in the literature. In particular, the data from our search experiment confirm that subjects tend to search less often than predicted by the optimal decision rule derived under the assumption of risk neutrality. Also, relatively simple heuristics, such as the constant reservation price heuristic and the satisficer heuristic, describe observed search behavior very well.

²² The corresponding Spearman- ρ and p -values for the subgroups C1, C2, C3, and C4 are -0.31 (0.032), -0.37 (0.01), -0.38 (0.068), and -0.23 (0.153), respectively.

The key question raised in this paper is whether the decision rules we observe in our data correspond to optimal behavior of risk-averse subjects (even though they are not optimal in the standard search model under risk neutrality). We therefore relax the assumption of risk neutrality made in the standard search models. We allow for departures from risk neutrality and develop optimal decision rules for such preferences. These decision rules (type 1 and type 2 rules) classify the observed behavior of the largest part of our sample. However, even the specifications of the generalized search models with risk aversion do not seem to be able to describe search behavior observed in our experiment fully. Our analysis rejects the hypothesized relationship between the individual preference parameter γ (the measure for risk aversion) and various parameters that characterize the observed search behavior over various subgroups under consideration.

This result may be disappointing. Since the search problem formally corresponds to a generalized lottery task, and since both the lottery-based utility elicitation tasks and the search tasks were performed in one experimental session, we should expect some correlation between the parameters of the lottery-based utility function elicitation task and characteristics of behavior in the search task at the subject level. However, while the individual risk parameter γ does not correlate with individual search parameters, we find that the loss aversion parameter λ does correlate with observed search behavior across all subgroups considered. This latter parameter accounts for the fact that individuals process losses differently than gains, and is related to the influential work on individual preferences by Kahneman and Tversky that led to the development of prospect theory. Conceptually, our results support other studies (e.g., Camerer, 2005; Kahneman *et al.*, 1991; Rabin and Thaler, 2001) that have suggested that loss aversion might be a major factor in observed attitudes towards risk, at least for modest scales.

We conclude this section with a discussion of some restrictions of our experimental design and of our analysis. First, a drawback of the procedure we used to elicit the shape of the utility function is that it operates on a monetary range of gains and losses that is higher than the range considered in the search experiment. While this separation is helpful for experimental design and parameter identification purposes, it may be the case that individual risk attitudes are different for high and low monetary outcomes. To allow for

this possibility, we analyzed our data under the weak assumption that the rank order of individuals by the relevant measure of risk attitude is preserved between the high-outcome range for which it is elicited and the low-outcome range that is relevant for the analysis of behavior in the search game.²³ A less restrictive, but also much more costly, experimental design would implement both the utility function elicitation procedure and the search game on the same high payment scale, or on the same low payment scale. The latter has been implemented in Schunk (2006), using a different utility function elicitation procedure. The findings support all conclusions drawn in this paper.

Second, the classification method used to assign decision rules to subjects may seem rather heuristic. For instance, depending on the set of candidate decision rules, this procedure may result in over-fitting. In our data, over-fitting is not an issue – we end up assigning subjects only to three classes of decision rules, and the variation within these classes (i. e., the constant reservation price assigned to each subject) is akin to estimating other preference parameters from experimental data. A final open issue of our analysis of search behavior is the role of errors in decision-making – in general, allowing for errors would tend to reduce the heterogeneity in preference parameters and decision rules. Using more sophisticated statistical methods for the classification of decision rules that allow for errors, as in Houser and Winter (2004) and Houser *et al.* (2004), is difficult given the nature of objective functions in search tasks and unlikely to produce substantively different results (Houser and Winter, 2004).

In summary, this study was motivated by the desire to understand search behavior and its relation to individual preferences, in particular risk attitudes. We have been able

²³ In order to investigate the appropriateness of using rank correlations, we conducted a secondary analysis of the data presented by Holt and Laury (2002). In their experimental study, Holt and Laury elicit three measures of risk aversion for each subject: two measures in a low-payoff condition as well as one measure in a high-payoff condition. The latter involves payoffs that are either 20, 50, or 90 times the amount of the low payoff condition. They also used both real and hypothetical payoffs.

When we re-analyze the data on those 187 subjects that were in a real payoff treatment (i.e. subjects that earned real money for lottery participation) we find a Spearman correlation coefficient of 0.49 ($p < 0.0000$) between the first low-payoff risk attitude measure and the high-payoff risk attitude measure. For the second low-payoff risk attitude measure and the high-payoff risk attitude, the Spearman correlation coefficient is 0.61 ($p < 0.0000$). Identical significance levels are found if we use only those subjects that were in a hypothetical treatment. We conclude that individual measures of risk attitudes elicited in low and high payoff situations exhibit a (stable) rank correlation. Further details of our re-analysis of the Holt and Laury (2002) data are available on request.

to replicate results from various previous studies on individual preferences and search behavior. Our main methodological contribution is to combine experiments on preferences and search so that correlations at the subject level could be analyzed. We find that there is considerable difference in the strategies that subjects use to solve the search task. These differences, however, do not seem to be systematically related to individuals' risk attitude elicited in lottery experiments. In contrast, we do find a relationship between the degree of loss aversion revealed in the lottery tasks and search behavior. In addition, our results suggest that a psychometric measure of their attitude towards risky gambles is also related to observed behavior in the experimental search task.

According to Kahneman's and Tversky's prospect theory, the finding of a correlation between individual loss aversion and search behavior suggests that reference point effects play a role when solving the search tasks; subjects apparently do not solve the search task only on the gain domain, as suggested by classical search theory. Schunk (2006) constructs and experimentally tests a descriptive model of search behavior that accounts for the observed reference points effects in search behavior and finds results that are in line with the findings in this paper. Overall, this model provides a better empirical fit than the standard model derived under risk neutrality or the extensions considered in the present paper. Testing such models experimentally as well as combining psychometric and decision-theoretic instruments for predicting behavior in sequential gambles should be the focus of future research on search behavior in particular and dynamic choice behavior in general. Furthermore, our findings are of interest for work in applied search theory, e.g. consumer and labor search: Here, results on individual search behavior and preferences might be helpful as a guide to econometric specifications that allow for heterogeneity, for example with respect to individual search duration.

References

- Abdellaoui, M. (2000):** Parameter-free elicitation of utility and probability weighting functions. *Management Science*, 46(11), 1497–1512.
- Benartzi, S. and R. H. Thaler (1995):** Myopic loss aversion and the equity premium puzzle. *Quarterly Journal of Economics*, 110, 73–92.
- Braunstein, Y. M. and A. Schotter (1982):** Labor market search: An experimental study. *Economic Inquiry*, 20, 133–144.
- Bromiley, P. and S. Curley (1992):** *Individual differences in risk-taking*. John Wiley: New York.
- Camerer, C. F. (1989):** An experimental test of several generalized utility theories. *Journal of Risk and Uncertainty*, 2, 61–104.
- Camerer, C. F. (2005):** Three cheers - psychological, theoretical, empirical - for loss aversion. *Journal of Marketing Research*, 42, 129–133.
- Cox, J. C. and R. L. Oaxaca (1989):** Laboratory experiments with a finite-horizon job search model. *Journal of Risk and Uncertainty*, 2, 301–330.
- Currim, M. and R. K. Sarin (1989):** Prospect versus utility. *Management Science*, 35, 22–41.
- Eckstein, Z. and G. J. van den Berg (2007):** Empirical labor search: A survey. *Journal of Econometrics*, 136, 531–564.
- Harrison, G. W. and P. Morgan (1990):** Search intensity in experiments. *Economic Journal*, 100, 478–486.
- Hey, J. D. (1981):** Are optimal search rules reasonable? And vice versa? (And does it matter anyway?). *Journal of Economic Behavior and Organization*, 2, 47–70.
- Hey, J. D. (1982):** Search for rules for search. *Journal of Economic Behavior and Organization*, 3, 65–81.
- Hey, J. D. (1987):** Still searching. *Journal of Economic Behavior and Organization*, 8, 137–144.
- Holt, C. A. and S. K. Laury (2002):** Risk aversion and incentive effects. *American Economic Review*, 92(5), 1644–1655.
- Houser, D., M. Keane, and K. McCabe (2004):** Behavior in a dynamic decision problem: An analysis of experimental evidence using a Bayesian type classification algorithm. *Econometrica*, 72(3), 781–822.
- Houser, D. and J. Winter (2004):** How do behavioral assumptions affect structural inference? Evidence from a laboratory experiment. *Journal of Business and Economic Statistics*, 22(1), 64–79.

- Kahneman, D., J. L. Knetsch, and R. H. Thaler (1991):** The endowment effect, loss aversion, and status quo bias: Anomalies. *Journal of Economic Perspectives*, 5, 193–206.
- Koebberling, V. and P. P. Wakker (2005):** An index of loss aversion. *Journal of Economic Theory*, forthcoming.
- Kogut, C. A. (1990):** Consumer search behavior and sunk costs. *Journal of Economic Behavior and Organization*, 14, 381–392.
- Lippman, S. A. and J. J. McCall (1976):** The economics of job search: A survey. *Economic Inquiry*, 14, 155–189.
- Moon, P. and A. Martin (1990):** Better heuristics for economic search: Experimental and simulation evidence. *Journal of Behavioral Decision Making*, 3, 175–193.
- Pennings, J. M. E. and A. Smidts (2000):** Assessing the construct validity of risk attitude. *Management Science*, 46(10), 1337–1348.
- Pratt, J. W. (1964):** Risk aversion in the small and in the large. *Econometrica*, 32, 122–136.
- Rabin, M. (2002):** Inference by believers in the law of small numbers. *Quarterly Journal of Economics*, 117, 775–816.
- Rabin, M. and R. H. Thaler (2001):** Anomalies: Risk aversion. *Journal of Economic Perspectives*, 15(1), 219–232.
- Rapoport, A. and A. Tversky (1970):** Choice behavior in an optimal stopping task. *Organizational Behavior and Human Performance*, 5, 105–120.
- Schmidt, U. and S. Traub (2002):** An experimental test of loss aversion. *Journal of Risk and Uncertainty*, 25(3), 233–249.
- Schotter, A. and Y. M. Braunstein (1981):** Economic search: An experimental study. *Economic Inquiry*, 19, 1–25.
- Schunk, D. (2006):** Sequential decision behavior with reference-point preferences: Theory and experimental evidence. *University of Mannheim, Working Papers*.
- Seale, D. A. and A. Rapoport (1997):** Sequential decision making with relative ranks: An experimental investigation of the “secretary problem”. *Organizational Behavior and Human Decision Processes*, 69(3), 221–236.
- Seale, D. A. and A. Rapoport (2000):** Optimal stopping behavior with relative ranks: The secretary problem with unknown population size. *Journal of Behavioral Decision Making*, 13, 391–411.
- Sonnemans, J. (1998):** Strategies of search. *Journal of Economic Behavior and Organization*, 35, 309–332.
- Sonnemans, J. (2000):** Decisions and strategies in a sequential search experiment. *Journal of Economic Psychology*, 21, 91–102.
- Stigler, G. J. (1961):** The economics of information. *Journal of Political Economy*, 69, 213–225.

- Tversky, A. and D. Kahneman (1992):** Advances in prospect theory: Cumulative representation of uncertainty. *Journal of Risk and Uncertainty*, 5, 297–323.
- Wakker, P. P. and D. Deneffe (1996):** Eliciting von Neumann-Morgenstern utilities when probabilities are distorted or unknown. *Management Science*, 42, 1131–1150.
- Weber, E. U., A.-R. Blais, and N. E. Betz (2002):** A domain-specific risk-attitude scale: Measuring risk perceptions and risk behaviors. *Journal of Behavioral Decision Making*, 15(4), 263–290.
- Zwick, R., A. Rapoport, A. K. C. Lo, and A. V. Muthukrishnan (2003):** Consumer sequential search: Not enough or too much? *Marketing Science*, 22(4), 503–519.

FIGURES AND TABLES

FIGURE 1

Optimal constant reservation price level depending on the individual risk coefficient γ




FIGURE 2

Optimal reservation price path depending on individual risk coefficient γ




FIGURE 3
Distribution of the constant reservation prices observed in the experiment.

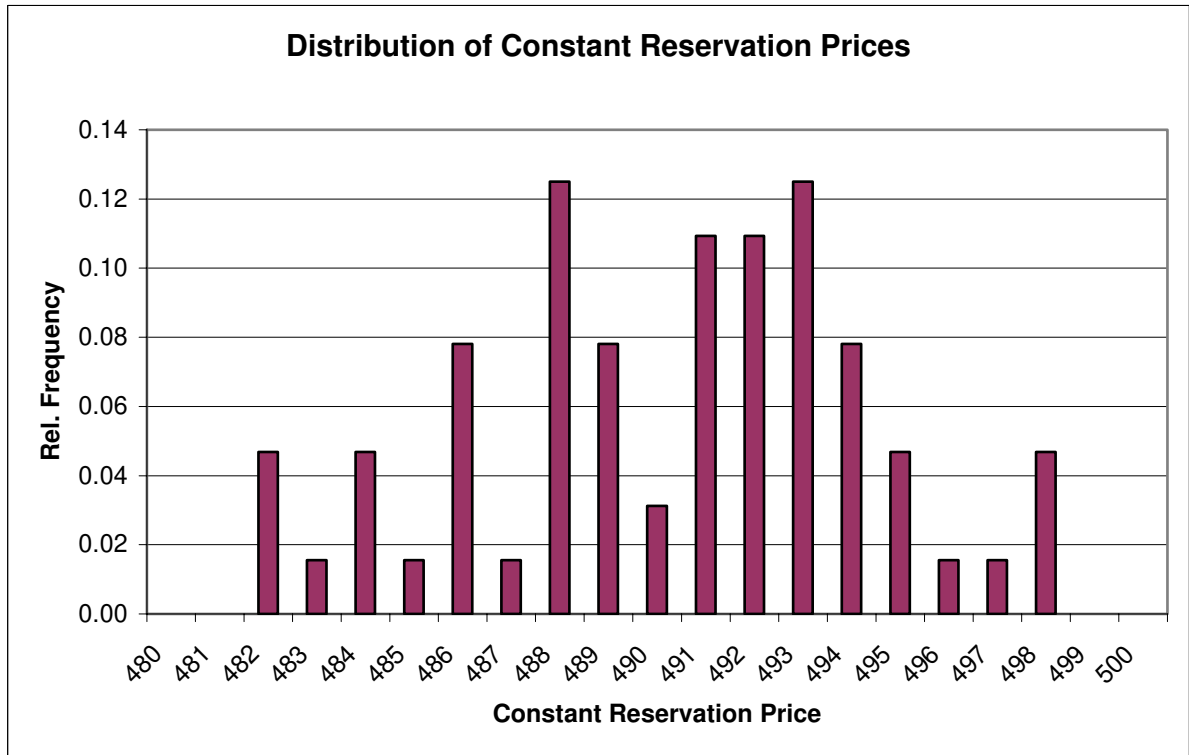


TABLE 1
Decision rules for the search problem

Number	Description	Parameter Values
1	<i>Constant reservation price heuristic</i> Stop searching as soon as a price below x € is found.	$x \in \{480, \dots, 500\}$
2	<i>Finite horizon optimal search</i> Stop searching in search step t as soon as a price below the reservation price $x_{t,\gamma}$ €, as specified by the finite horizon search model, is found.	$\gamma \in \{-1.0, -0.95, \dots, +0.95, +1.0\}$
3	<i>Satisficer heuristic</i> Stop searching as soon as the payoff from stopping exceeds a certain threshold level of x €	$x \in \{1, \dots, 20\}$
4a	<i>One-bounce rule</i> Have at least 2 searches and stop if a price quote is received larger than the previous quote.	
4b	<i>Modified one-bounce rule</i> Have at least 2 searches and stop if a price quote is received larger than the previous quote less the search cost.	
5a	<i>Streak-based rule</i> Stop searching as soon as 2 consecutive price draws that are below some fixed threshold level x € are received.	$x \in \{485, \dots, 500\}$
5b	<i>Streak-based rule</i> Stop searching as soon as 3 consecutive price draws that are below some fixed threshold level x € are received.	$x \in \{485, \dots, 500\}$

TABLE 2

Utility function estimation results and risk classification of the individuals.

	Utility function	
	Gains (γ)	Losses (δ)
Median estimate	2.003E-04	2.045E-04
Mean R²	0.9949	0.9948
Risk averse subjects	63%	23%
Risk neutral subjects	15%	18%
Risk seeking subjects	22%	59%

TABLE 3
 Correlations between the search parameters (reservation price, average number of searches, search coefficient γ^{search}) and the preference parameters (γ, δ, λ) by subgroup.

Group (N)	Preference Parameters	Constant Reservation Price		Average Number of Searches		Search coefficient γ^{search}	
		Spearman - ρ	p-value	Spearman - ρ	p-value	Spearman - ρ	p-value
C1 (49)	γ (Risk on Gains)	-0.03	0.82	0.02	0.90	-0.05	0.72
	δ (Risk on Losses)	-0.03	0.83	0.14	0.34	-0.06	0.66
	λ (Loss aversion)	0.23	0.12	-0.25	0.08	0.21	0.15
C2 (45)	γ (Risk on Gains)	0.02	0.90	0.00	0.99	0.06	0.72
	δ (Risk on Losses)	-0.03	0.87	0.11	0.48	-0.02	0.89
	λ (Loss aversion)	0.26	0.08	-0.30	0.04	0.21	0.16
C3 (24)	γ (Risk on Gains)	-0.19	0.39	0.34	0.10	-0.12	0.59
	δ (Risk on Losses)	-0.25	0.24	0.46	0.02	-0.17	0.43
	λ (Loss aversion)	0.35	0.09	-0.31	0.14	0.29	0.17
C4 (54)	γ (Risk on Gains)	0.03	0.86	-0.18	0.26	0.05	0.76
	δ (Risk on Losses)	0.05	0.77	-0.13	0.43	0.03	0.84
	λ (Loss aversion)	0.22	0.16	-0.33	0.04	0.14	0.39
All (64)	γ (Risk on Gains)	-0.07	0.56	0.09	0.47	-0.03	0.79
	δ (Risk on Losses)	-0.05	0.71	0.14	0.27	-0.44	0.73
	λ (Loss aversion)	0.28	0.02	-0.32	0.01	0.21	0.10

TABLE A1

Assessing x_1 through bisection. An example of the Abdellaoui (2000) procedure.

Question number	Alternatives	Outcomes (€) x_1 €	Choice
1	A = (200, 2/3; 100, 1/3) B = (700, 2/3; 0, 1/3)	[200, 1200]	A
2	A = (200, 2/3; 100, 1/3) B = (950, 2/3; 0, 1/3)	[700, 1200]	B
3	A = (200, 2/3; 100, 1/3) B = (820, 2/3; 0, 1/3)	[700, 950]	A
4	A = (200, 2/3; 100, 1/3) B = (880, 2/3; 0, 1/3)	[820, 950]	A
5	A = (200, 2/3; 100, 1/3) B = (910, 2/3; 0, 1/3)	[880, 950]	B
6			
End		[880, 910]	

$x_1 = \text{€ } 200$, $p = 2/3$, $r = 0$, $R = \text{€ } 100$