SCISPACE
formerly Typeset

Open access • Journal Article • DOI:10.1016/J.JEBO.2008.12.010

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

Daniel Schunk, Joachim Winter

Institutions: University of Zurich, Ludwig Maximilian University of Munich

Published on: 01 Aug 2009 - Journal of Economic Behavior and Organization (North-Holland)

Topics: Heuristics, Prospect theory, Satisficing, Loss aversion and Decision rule

Related papers:

- Strategies of search
- Laboratory experiments with a finite-horizon job-search model
- Search intensity in experiments
- Behavioral heterogeneity in dynamic search situations: Theory and experimental evidence
- Search for rules for search

Share this paper:

View more about this paper here: https://typeset.io/papers/the-relationship-between-risk-attitudes-and-heuristics-in-51ukpfo3cv

University of Zurich
Zurich Open Repository and Archive

Winterthurerstr. 190
CH-8057 Zurich
http://www.zora.uzh.ch

Year: 2009

The relationship between risk attitudes and heuristics in search tasks: A laboratory experiment

Schunk, D; Winter, J

Schunk, D; Winter, J (2009). The relationship between risk attitudes and heuristics in search tasks: A laboratory experiment. Journal of Economic Behavior & Organization, 71(2):347-360.
Postprint available at:
http://www.zora.uzh.ch

Posted at the Zurich Open Repository and Archive, University of Zurich.
http://www.zora.uzh.ch

Originally published at:
Journal of Economic Behavior & Organization 2009, 71(2):347-360.

The relationship between risk attitudes and heuristics in search tasks: A laboratory experiment

Abstract

Experimental studies of search behavior suggest that individuals stop searching earlier than the optimal, risk-neutral stopping rule predicts. Two different classes of decision rules could generate this behavior: rules that are optimal conditional on utility functions departing from risk neutrality, or heuristics derived from limited cognitive processing capacities and satisficing. To discriminate between these possibilities, we conduct an experiment that consists of a 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.

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

Daniel Schunk**

University of Zurich

Joachim Winter†

University of Munich

Abstract: Experimental studies of search behavior suggest that individuals stop searching earlier than the optimal, risk-neutral stopping rule predicts. Two different classes of decision rules could generate this behavior: rules that are optimal conditional on utility functions departing from risk neutrality, or heuristics derived from limited cognitive processing capacities and satisficing. To discriminate between these possibilities, we conduct an experiment that consists of a 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 Catherine Eckel (the associate editor) and the referees for their guidance. Helpful comments were also provided by Ernst Fehr, 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). Financial Support from the Deutsche Forschungsgemeinschaft (DFG) via the Sonderforschungsbereich 504 at the University of Mannheim is gratefully acknowledged.

** Corresponding author.
Address: Institute for Empirical Research in Economics
University of Zurich
CH-8006 Zurich, Switzerland
E-mail: dschunk@iew.uzh.ch
Phone: +41-(0)44-634-3704
Fax: +41-(0)44-634-4907

† *Address*: Department of Economics
University of Munich
D-80539 Munich, Germany
E-mail: winter@lmu.de

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. But behavior in search tasks is not only interesting in the analysis of many substantive issues; it has also proven to be a useful object of behavioral research in psychology and economics. Search tasks are attractive for experimental studies because of their (superficially) simple structure masking an underlying optimization problem that is quite complicated and that usually cannot be solved in the human mind, but that instead requires numerical methods and a computer. Conceptually, search tasks are representative of many situations where the decision must be made between either committing resources to an attractive proposition or deferring the decision in hope of receiving a better deal.

Numerous authors such as Stigler (1961), Braunstein and Schotter (1982), Hey (1981, 1982, 1987), Kogut (1990), Harrison and Morgan (1990), Houser and Winter (2004), Schotter and Braunstein (1981), Schunk (2008), and Sonnemans (1998, 2000), have investigated price search situations and their variations, such as the well known secretary problem, where the decision to stop or to continue depends only on the ranks of the presented alternatives.[1] 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 individuals are very heterogeneous in their search behavior and that relatively simple heuristics describe observed behavior better than the optimal stopping rule. Subjects' search behavior, however, has been found to be nearly optimal in the sense that their actual earnings approximate those they would have real-

[1] In the secretary problem, the decision-maker is faced with a set of items (e.g. job applications) which are presented to her one at a time in random order. The decision-maker can rank order all observed items in terms of their desirability, and in each period she must either accept the presented item (in which case the search process terminates) or reject it (in which case the next item in the randomly determined order is observed). See Rapoport and Tversky (1970) and Seale and Rapoport (1997, 2000), for example, for discussions of the secretary problem.

ized 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 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 individuals seem to behave as theory predicts 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 the individuals' risk-averse behavior (Cox and Oaxaca, 1989). Using an electronic information board method, Sonnemans (1998) finds that differences in the subjects' learning behavior 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. Experimental studies typically find, however, that most subjects do not use rules derived under the assumption of risk neutrality, but instead follow some heuristic. These heuristics are often sophisticated in the sense that they allow subjects to closely approximate the payoffs they would have obtained using optimal rules. If, however, we allow for heterogeneity with respect to the individual risk attitudes, the situation becomes more complicated: decision rules 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 between these two possibilities requires an independent measure of risk attitudes. We therefore not only present subjects with a search task that follows the standard in the literature, but also with a lottery task that serves to elicit their individual utility functions. Furthermore, we use a questionnaire to obtain a psychometric measure of risk attitudes as an independent individual-level source of information on risk behavior.

The contribution of our paper to the search literature is the study of the relation between properties of subjects' preferences and their decision heuristics used in search tasks. We find that while the individual risk attitude does not seem to be related to the decision heuristics subjects use, loss aversion is related to search behavior. Our findings are of inherent interest for the field, as economic science assumes that preferences are the key determinants of behavior. More specifically, our findings are of interest to researchers who investigate the determinants of labor and consumer market search behavior (e. g., Eckstein and van den Berg, 2007; Zwick et al., 2003); here, for example, much evidence stems from detailed econometric analyses of field and survey data, and our findings serve as a guide for novel econometric specifications that allow for heterogeneity with respect to search behavior. Furthermore, the findings are relevant for researchers who are generally interested in the determinants of behavioral heterogeneity in dynamic choice situations (e. g., Boswijk et al., 2007). In section 2, we present the design of our experiment. Section 3 describes our procedures for drawing 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 a 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 from the psychology literature in order 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 Abdellaoui (2000) recently proposed that elicits the whole individual preference functional without imposing any prior restrictions, i.e. it is more flexible than price list methods. In contrast to price list methods, the method we use is also robust against probability distortions, since only one pair of probabilities, p and

$(1-p)$, is necessary for the whole procedure. Moreover, it only involves choices between pairs of gambles that are presented simultaneously to the subjects. The choice between two gambles is more similar to the stop-or-go decision that subjects must make in the search task than is the decision using a price list method or a method asking subjects to state a certainty equivalent.[2]

The Abdellaoui (2000) method elicits each subject's utility function on the gain and loss domain, using a series of a total of 64 lottery choice questions.[3] The subjects are presented with two lotteries, lottery A and lottery B, in each of the questions. Subjects were then instructed to designate which lottery in each displayed pair of prospects they would prefer by means of a mouse click on the appropriate field ("lottery A" or "lottery B", see the instructions for a screenshot). After 30 lottery choice questions, we identified seven points, a so-called standard sequence of outcomes $\{x_0, x_1, ..., x_6\}$, on the individual utility function in the gain domain. Another 30 lottery choice questions serve to identify seven points on the individual utility function in the loss domain. The two standard sequences of outcomes can then be used for estimating the curvature of the individual utility function, i.e. a measure for individual risk attitude, separately in the gain and loss domains. The complete elicitation method is described in detail in the appendix.

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 they value at €500. This article is sold at infinitely many locations, and visiting a new location costs €1. A price is randomly drawn from a known distribution at each location. The instruction sheet informs subjects 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.

[2] Note further that Bostic *et al.* (1990) report experimental evidence that methods involving decisions between two lotteries yield more consistent answers than methods involving matching.

[3] Four of the 64 lottery questions appear twice during the lottery elicitation process, giving us the possibility of investigating whether subjects behave consistently during the utility elicitation questions, or whether preference reversals have occurred.

After each new price draw (that is, at each location they visit), subjects are allowed to recall previously rejected price offers. That is, subjects can stop after each price draw 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 accrued search cost.

There are several reasons why we chose to allow for recall and to provide subjects with full knowledge about the distribution of the prices. First, allowing for recall makes the task closer to real-world situations, such as price search in the internet: In these situations, individuals can typically 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 came across during their search. Second, providing subjects with information about the price distribution enables us to derive simple optimal decision rules that will serve as a benchmark for observed behavior. If subjects did not know the price distribution, they would update their – potentially different – priors about the price distribution after each search decision. While this can also be incorporated in a theoretical model, we would then have a less controlled environment, and the observation of learning behavior would distort our inference about heuristics in dynamic choice tasks.

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.[4] Finally, after the experiment was completed, one of these rounds was selected randomly to determine the part B payoff.

2.3 Part C: Risk Attitudes

The experiment ends with a short computerized questionnaire (part C). Weber *et al.* (2002) developed this survey instrument for assessing risk taking. They construct a psychometric scale, the so-called DOSPERT scale, for subjects' risk attitudes in five

[4] 35 subjects played ten search rounds, and the second half, another 33 subjects, played 11 payment-relevant search rounds.

content domains: (i) financial (further subdivided into investment and gambling), (ii) health/safety, (iii) recreational, (iv) ethical, and (v) social. Weber *et al.* (2002) evaluated the scale for a sample of more than 900 American undergraduate students and found that respondents' degree of risk taking was highly domain-specific. They reported test-retest reliability estimates and provided evidence for the convergent/discriminant and factorial validity with respect to constructs such as sensation seeking and dispositional risk taking. Furthermore, construct validity of the scale was also assessed based on correlations with the results of a risky gambling task. To the best of our knowledge, the scale is the only domain-specific risk attitude scale that has been officially translated into German language (DOSPERT-G) and validated in this language (Johnson *et al.* , 2004).

The DOSPERT-scale has been widely used in psychology, decision science, and economics, and its factor structure has been replicated in a wide range of settings and populations (see, e. g., Carpenter *et al.* , 2007; Hanoch *et al.* , 2006; Harris *et al.* , 2006; Wilke *et al.* , 2006). The scale allows the assessment of both conventional risk attitudes (defined as the reported level of risk taking) and perceived-risk attitudes (defined as the willingness to engage in a risky activity as a function of its perceived riskiness). We are only interested in risk attitudes for predictive purposes in our experiment, and we do not intend to influence the level of individual risk attitude. Therefore, only the questions on conventional risk attitudes are of interest for our purpose (see, e. g., Weber *et al.* , 2002, p. 282). Furthermore, we focus on the domain of gambling because of the nature of our experimental task.

In the risk attitude questions, subjects are asked to rate their behavior with respect to four risky activities in the behavioral risk domain of gambling.[5] Specifically, subjects report how likely they are to engage in a certain gambling-related activity on a five-point rating scale ranging from 1 ("Extremely likely") to 5 ("Extremely unlikely"). These questions provide us with 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.

[5] The corresponding questions are reported in the online-version of the appendix.

2.4 Administration

The study was conducted in fall 2003 in the experimental laboratory of Sonderforschungsbereich 504, a research center at the University of Mannheim. A total of 68 students participated in the main study in four sessions.[6] The 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 would be truncated at €0, i.e. subjects would not incur a loss from the experiment even if their payoff were negative (which is theoretically possible in our experiment).[7] Finally, one of the subjects participating in each experimental session (average number 17) was randomly selected to play for a real monetary payoff 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 payoff. 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 infer the 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.

[6] 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.

[7] The lowest payoff made in any session was €4, meaning no subject would have suffered a loss, even without truncation.

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 for a nonparametric estimation of utility functions in the gain and loss domain as well as the corresponding probability weighting functions. For the purpose of our study, we need to order subjects according to their risk attitudes and therefore use a parametric approach, specifying the subjects' risk attitudes 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 for eliciting the shape of the utility function (Part A) operates on a monetary range of gains and losses that differs from that considered in the search experiments (Part B). We purposely made this decision, and digress here for a brief discussion of its rationale. As Wakker and Deneffe (1996) point out, 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 over a wide interval below €-200 or above €200, respectively (with the size dependent on the subjects' decisions, see Abdellaoui (2000)). Budget limitations, however, forced us to reduce the outcome scale from €-40 to €40 in the search game. Even though individual risk attitudes may vary for high and low monetary outcomes, all that our empirical analysis requires is that the rank order of individuals by the measures of risk attitudes be preserved between the high-outcome range for which it is elicited and the low-outcome range relevant for the analysis of behavior in the search game.[8] This is, in our view, a reasonable assumption and, in fact, a corollary of using a CARA-utility specification. Furthermore, this assumption can be investigated using data from high and low outcome risk elicitation tasks by Holt and Laury (2002). The results are supportive, and are presented in section 5 of this paper.

[8] Accordingly, our empirical analysis will only be based on rank correlations, i.e. on "comparative risk aversion".

Based on Currim and Sarin (1989) and Pennings and Smidts (2000), for example, we assume the following exponential specification for our CARA-utility function on gains[9]:

$$u(x) = \frac{1 - e^{-\gamma(x - x^G_{min})}}{1 - e^{-\gamma(x^G_{max} - x^G_{min})}} \tag{2}$$

Here, x^G_{max} is the largest elicited value of x in the gain domain (in absolute values), i. e., x_6; x^G_{min} 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^G_{max} and x^G_{min} by x^L_{max} and x^L_{min}, respectively; we use the absolute values 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 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 – both

[9] Note that another normalized version of the CARA-utility has the following form:

$$u(x) = \frac{1 - e^{-\gamma \frac{x - x^G_{min}}{x^G_{max} - x^G_{min}}}}{\gamma} \tag{1}$$

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 (2). The substantive conclusions of our analysis remain unchanged when we use the form (1).

invariant to changes in the scale of the utility function, $u(\cdot)$, and to scale transformations of the outcomes. This index is given by

$$\lambda = \frac{u'\uparrow(0)}{u'\downarrow(0)}. \tag{3}$$

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 $\delta \neq 0$:

$$\lambda = \frac{\frac{\delta\cdot(e^{-\delta\cdot(|x^L - x^L_{min}|)})}{1-e^{-\delta\cdot(|x^L_{max} - x^L_{min}|)}}}{\frac{\gamma\cdot e^{-\gamma(x^G - x^G_{min})}}{1-e^{-\gamma\cdot(x^G_{max} - x^G_{min})}}} \tag{4}$$

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 (4) in the denominator and numerator, respectively.

Note that our estimate of individual loss aversion is based on the assumption that the estimated form of an individual's CARA-utility function is characteristic for her utility function over the whole domain, and identically scaled both for gains and for losses[10].

3.2 Classification of Decision Rules Used in the Search Task

The next step of our analysis is to determine the decision rule each subject 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 were used in the earlier literature (e. g., Hey, 1982; Moon and Martin, 1990; Houser and Winter, 2004) to describe search behavior. We compute the number of stopping decisions that are correctly predicted for each subject and each candidate decision rule. We then assign to the subject the decision rule that generates the largest fraction of correct predictions, i.e. that which best fits observed behavior.

[10] 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 yield 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 based on mixed lotteries (e. g., Schmidt and Traub, 2002) could also be used in the context of a search experiment, even though they are also subject to considerable uncertainty. Schunk (2008) uses mixed lotteries in an experimental study on search behavior that uses a different design to elicit individual preferences. We suggest that further experimental studies should 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.

We start this subsection with the presentation of optimal decision rules conditional on individual risk attitude which is based on the assumption of a classical von-Neumann-Morgenstern utility function.[11] Then, we discuss the set of alternative heuristics, and we end this section with a formal description of our classification procedure.

Decision Rules in Search Tasks Conditional on Risk Attitude

We consider two cases: In the first, we treat the cost of each completed search step as a sunk cost and assume an infinite horizon; in the second, we derive the finite horizon optimal stopping rule assuming that subjects do not treat past search costs as sunk.

In the first case, we assume in line with basic search theory that individuals treat the cost of each search step, 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. Subjects must then solve the problem based on a one-step forward-induction strategy, and the corresponding decision rule has the constant reservation price property, which has been reported as a search heuristic consistent with the behavior of a reasonable number of subjects in other studies (e. g., Hey, 1987). The optimal strategy for a risk neutral searcher is to keep searching until a price less than or equal to €490 is found. Figure 1 shows the constant reservation price as a function of the risk-parameter γ in the exponential utility function (2). 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.

** Include figure 1 about here **

In the second case, we assume that subjects do not treat search costs as sunk costs, meaning they consider the total benefits and costs of search in their decision whether to stop or to continue the 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 the continuation value from the experiment would definitely be zero after 48 search steps. It follows that the problem is treated as a finite horizon problem that is

[11] Details of the derivation are presented in the appendix of this paper.

solved backwards. For a risk neutral searcher, the corresponding reservation price path 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. Figure 2 plots the path of reservation prices for various risk attitudes γ of the searcher. Henceforth, we will refer to rules of this type as backward optimal rules.

** Include figure 2 about here **

We can conclude from our previous theoretical deliberations 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, while risk-seeking subjects should stop their search later, i. e., they use lower reservation prices.

Alternative Search Rules

As the search literature has previously pointed out, and as we described in the previous section and in the appendix, 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 the search literature therefore argue that subjects use heuristics rather than the optimal stopping rule, and there is evidence that certain heuristics result in subjects receiving close to the payoffs they could have obtained using the optimal rule.

We now specify our set of candidate search rules we use in this paper for characterizing behavior in experimental search tasks. In addition to the search rules presented in the section above, we specify a set of heuristics that were used in the search literature for describing 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 property of a constant reservation price. Each rule says that the subject uses an arbitrary, but constant reservation value $r \in \{480, .., 500\}$. 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 discussed above. Based on this rule, we attribute the constant reservation price value to every individual that explains most of her observed search decisions.

The second class of decision rules 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 $\gamma \in \{-1.0, -0.95, -0.9, ..., +0.95, +1.0\}$. We refer to this class of decision rules as *type 2 rules.* Based on this rule, we attribute a value γ_i^{search} to every individual, i.e. 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 $t \in \{1, .., 20\}$. Given our parametrization of the problem, this results in a reservation price that falls linearly over time. For obvious reasons, this heuristic is sometimes called the "satisficer heuristic" and we refer to it as *type 3 heuristics*.

We consider the so-called "bounce rules", suggested by Moon and Martin (1990) based on earlier work by Hey (1982) as *type 4 heuristics*. Subjects following the "one-bounce rule" (heuristic 4a) have at least 2 searches and they stop after receiving a price quote larger 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 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 $p \in \{485, .., 500\}$. Results from psychological research on behavior in uncertain environments, such as Rabin (2002), can provide motivation as to why subjects use these streak-based rules in search situations.

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

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

** Include table 1 about here **

Classification Procedure

Our approach to drawing inferences about search behavior is to determine the proportion of choices consistent with each decision rule for each subject 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 payoff tasks. This latter assumption seems reasonable in view of the fact that all subjects are experienced when they begin the payoff tasks.

Formally, our classification procedure can be described as follows.[12] 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 d_{it}^{*} denote the observed decision of subject i in period t. Then, we can define the indicator function:

$$X_{it}^{c_i}(S_{it}) = 1(d_{it}^{*} = d_{it}^{c_i}(S_{it})) \tag{5}$$

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:

$$\hat{c}_i = \arg\max_{c_i \in \mathcal{C}} \sum_{t=1}^{T_i} X_{it}^{c_i}(S_{it}) \tag{6}$$

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

As the existing literature indicated, 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 the relationship between the observed search behavior and the individuals' risk preferences for each subgroup and for the whole sample.

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"). We calculate a measure of risk attitude as the arithmetic mean score of the response to the four questions for each subject.

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).

A total of 68 subjects participated in our experiment; we were forced to delete four subjects from the sample as they apparently did not take the utility elicitation part of the experiment seriously.[13]

The 64 subjects that we retain 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.[14]

[13] 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 to complete 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 are more than one standard deviation larger than those for all the other subjects, i. e., their preference parameters are measured imprecisely. Furthermore, these two subjects are the only ones in the sample with revealed preference reversals on *all* four consistency check questions (see the section 2.1).

[14] The reversal rate is a measure for how consistent subjects behave in a certain utility elicitation mechanism. Our reversal rate is somewhat higher than that 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

4.1 Part A: Preferences

In Table 2, we report the results of the nonlinear least squares estimation of the risk coefficients γ and δ. Note that the mean coefficient of determination, R^2, is close to 1 over all utility function estimations, both for gains and losses. This suggests that the risk coefficients are reliable measures of the curvature of the subject's utility function. Abdellaoui (2000), who developed the risk attitude elicitation method we use, focuses on the classification of subjects into risk averse, risk neutral, and risk seeking types based on a non-parametric statistic for the curvature implied by the elicited points of the utility function. Our parametric estimates also yield a classification of subjects into these three types, but – extending Abdellaoui's work – our estimates further allow for a rank-ordering of individuals according to the degree of curvature of their utility function.[15] Comparing our risk attitude classification with that obtained by Abdellaoui (2000), we find that 63% of the subjects are risk averse, 15% are risk neutral, and 22% are risk seeking in the gain domain, while Abdellaoui (2000) finds 58%, 20%, and 22%, respectively. We find 23% risk averse, 18% risk neutral, and 59% risk seeking subjects in the loss domain, compared to 22% risk averse, 29% risk neutral, and 49% in Abdellaoui (2000). Overall, our results on the distribution of risk attitudes in the population are consistent with the predictions of prospect theory (Tversky and Kahneman, 1992).

** Include table 2 about here **

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 attributed to each subject can be considered a proxy for whether

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.

[15] Abdellaoui does not provide results on the distribution of parametric estimates for the subjects' utility functions. He mentions, however, that median estimates from fitting a power utility function are 0.89 and 0.90 for γ (curvature on the gain domain) and δ (curvature on the loss domain), respectively. Fitting a power function to our data, we find median estimates of 0.89 for the gain domain, and 0.92 for the loss domain.

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 55% of the subjects to be "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 were to use the risk-neutral optimal reservation stopping rule with a reservation price of €490, they would 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).

** Include figure 3 about here **

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 a constant reservation price heuristic explains behavior better than all other heuristics for 13% of the subjects; a type 2 rule (the optimal finite horizon rule) is better than all others for 3% , while a satisficer rule (type 3) explains more observations than all other rules and for 16% . One of the conditionally optimal rules (type 1 or type 2) is a best decision rule for 84% of all subjects, while we find that 63% use one of the optimal rules (type 1 or type 2) without using 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. We cannot, however, distinguish between the use of a forward or a backward optimal search rule (type 1 or type 2) for 47% of the subjects.[16]

[16] 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. The reported weak discrimination between both types of rules thus 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

** Include figure 4 about here **

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, while 38.5% of the decisions are consistent with a type 5a heuristic, and 39.4% with a type 5b heuristic. Hey (1982), however, who 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. Our data, however, do not discriminate between the usage of type 1 or type 2 or type 3 decision rules for a certain proportion of the subjects.[17] Based on these findings, we classify the 64 subjects into 4 categories, labeled 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).
- 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).

---

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), thus complicating 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.

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

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. Our subjects' preferences are characterized by their coefficients of risk attitude in the gain and loss domains, γ and δ, respectively. To characterize search behavior, we focus on the following parameters: 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 our basic search model, we hypothesize that γ should be positively correlated with RP and negatively correlated with AS, at least for subgroup C1. We further hypothesize that γ is positively correlated with γ^{search} and negatively correlated with ASfor subgroup C2. Furthermore, note that the attributed constant reservation price (RP) and the attributed γ^{search} are strongly positively correlated (Spearman-ρ: 0.946, p-value: 0.00), due to being derived from the same underlying utility functional; we should therefore also expect the correlations hypothesized above for subgroup C4. Table 3 reports the corresponding Spearman correlation coefficients for all subgroups C1 through C4 and the whole sample.

** Include table 3 about here **

Table 3 reveals that we neither find a statistically significant relationship between the risk attitude measure γ and the search parameters RP, AS, and γ^{search} for any subgroup nor for the whole sample. This holds true, regardless of whether we impose the usage of only one specific type of search rule (e. g., the one-step forward-optimal search rule) on *all* subjects, or whether we attribute the type of rule that best describes behavior to each subject and then 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 unequal variances show that our hypothesis – 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. An even stronger result is that risk seeking subjects' mean reservation prices are higher than those of risk averse subjects across all subgroups and the whole sample.

We now move to the psychometric measures of risk attitude and investigate the correlation between the *psychometric measure for risk attitude in the gambling domain* and *search behavior.* There is some evidence that individuals who dislike taking risks in the gambling domain tend to search less. We have a Spearman-ρ of 0.26 (p-value 0.087) for C2-subjects and a Spearman-ρ of 0.29 (p-value of 0.07) for C4-subjects for the correlation between the measure for risk on gambling and the average number of search steps per round (AS).[18]

With respect to the relationship between the *utility function based risk measures* and the *psychometric risk measures*, we find that, apart from subgroup C4, the loss aversion parameter at least marginally correlates 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-value of 0.009 for the correlation between the loss aversion parameter and the psychometric measure for risk on gambling.[19]

In summary, our data do not confirm our hypotheses on the relationship between utility function based measures for risk aversion and search behavior. However, we find significant relationships between the loss aversion index λ derived from the utility function, and both, the average number of search steps (AS) and the attributed constant reservation price level RP in Table 3. These correlations are significant or at least marginally significant across all subgroups considered and for the whole sample. Subjects across all subgroups with a higher degree of loss aversion tend to have a higher attributed reservation price and to 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: individuals 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 func-

[18] The corresponding Spearman-ρ and p-values for the C1- and C3-group and for the whole sample are 0.14 (0.34), 0.06 (0.76) and 0.16 (0.21).

[19] 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.

tions and a search experiment designed to reveal their decision rules in a search task. A psychometric survey instrument that generates domain-specific measures of risk attitudes augments the experiment. Concerning these elicited measures, we first would like to point out that the results of each of these components broadly correspond with earlier results in the literature. In particular, the data from our search experiment confirm that subjects tend to search less often than the optimal decision rule derived under the assumption of risk neutrality predicts. Also, relatively simple heuristics, such as the constant reservation price heuristic and the satisficer heuristic, describe observed search behavior very well.

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 (denoted as 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 apparently cannot fully describe the search behavior observed in our experiment. 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 seem 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 a 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 it 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 see two potential limitations in our experimental design and in our analysis. First, the procedure we used to elicit the shape of the utility function has a drawback in 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, individual risk attitudes might differ between high and low monetary outcomes. To allow for this possibility, we analyzed our data under the weak assumption that individuals' rank order by the relevant measure of risk attitude remains preserved between the high-outcome range for which it is elicited and the low-outcome range relevant for the analysis of behavior in the search game.[20] 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 was implemented in Schunk (2008), 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 ultimately only assign subjects 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 related open issue in our analysis of

[20] 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 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 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. The Spearman correlation coefficient is 0.61 ($p < 0.0000$) for the second low-payoff risk attitude measure and the high-payoff risk attitude. Identical significance levels are found if we only use those subjects who 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.

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; they are unlikely to produce substantively different results (Houser and Winter, 2004).

In summary, the desire to understand search behavior and its relation to individual preferences, in particular risk attitudes, motivated this study. We were able 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 are considerable differences in the strategies subjects use to solve the search task. These differences, however, do not seem to be systematically related to their risk attitude elicited in lottery experiments. In contrast, we find that a psychometric measure of their attitude towards risky gambles is related to search behavior. Moreover, the degree of loss aversion revealed in the lottery tasks is also significantly related to search behavior: The more loss averse a subject is, the earlier she stops the search process.

How can loss aversion be related to search behavior? There are two explanations for this finding: First, it is possible that loss aversion can be measured better than risk aversion in our experimental task. This explanation would be in line with Rabin (2000) who provides theoretical arguments that loss aversion can account better for observed decision behavior over modest stakes than the standard notion of risk aversion. A second explanation is that individuals set reference points during their search. This explanation is in line with existing laboratory and field evidence on myopic loss aversion in dynamic decision tasks (e. g., Baucells *et al.* , 2007; Benartzi and Thaler, 1995; Odean, 1998).

To understand how reference point-updating can explain our results, assume that subjects in each search step take the current best alternative that they were offered as their reference point and compare their possible future offers from continuing the search with that best offer. They would then consider every alternative that is worse than the current best payoff (which they have for sure) as a loss relative to that sure payoff. Every alternative that is better than the sure payoff would be considered as a gain. Clearly, more loss averse individuals – who are overrepresented among the less educated and among the elderly, according to an empirical study by Johnson *et al.* (2006) – would then stop the search process earlier, i. e., these individuals would tend to deviate more from optimal search behavior than others.[21]

Psychologically, the observation that loss aversion is related to search behavior might be due to an endowment effect (Huck *et al.* , 2005; Kahneman *et al.* , 1991): The object that can be kept for sure might be considered as an endowment, and the future steps in the search process are only evaluated relative to this endowment. In evaluating the benefit of continuing the search, loss averse decision makers would then combine the evaluation of any potential future outcomes that are below this endowment with an emotional reaction of fear (Loewenstein *et al.* , 2001). We speculate that endowment effects should be stronger if there is a longer time between the search steps and if a person identifies more strongly with the object she is searching for. This should be tested in future experimental work on this issue. Furthermore, in contrast to the vast empirical literature on behavior in static decision situations, the foundations of dynamic choice behavior remain largely unexplored in experimental literature. Therefore, testing reference-point based models of dynamic choice behavior experimentally as well as combining psychometric and decision-theoretic instruments for predicting behavior in sequential gambles should be the focus of future research in this field.

[21] Schunk (2008) develops a descriptive model of search behavior that accounts for the observed reference points effects in this behavior. He tests this model experimentally and finds evidence in line with the findings in this paper. Overall, this new model provides a better empirical fit than the standard model derived under risk neutrality or the extensions considered in the present paper.

References

Abdellaoui, M., 2000. Parameter-free elicitation of utility and probability weighting functions. Management Science 46(11), 1497-1512.

Benartzi, S., Thaler, R. H., 1995. Myopic loss aversion and the equity premium puzzle. Quarterly Journal of Economics 110, 73-92.

Baucells, M., Weber, M., Welfens. F., 2007. Reference point formation over time: A weighting function approach. Working Paper, University of Mannheim. Available at: http://ssrn.com/abstract=1022930

Bostic, R., Herrnstein, R. J., Luce, R. D., 1990. The effect on the preference reversal phenomenon of using choice indifferences. Journal of Economic Behavior and Organization 13, 193-212.

Boswijk, H. P., C. Hommes, Manzan, S., 2007. Behavioral heterogeneity in stock prices. Journal of Economic Dynamics and Control 31, 1938-1970.

Braunstein, Y. M., Schotter, A., 1982. Labor market search: An experimental study. Economic Inquiry 20, 133-144.

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.

Carpenter, J., Matthews, P. H., Schirm, J., 2007. Tournaments and Office Politics: Evidence from a Real Effort Experiment. IZA Discussion Paper No. 2972. Institute for the Study of Labor, Bonn.

Cox, J. C., Oaxaca, R. L., 1989. Laboratory experiments with a finite-horizon job search model. Journal of Risk and Uncertainty 2, 301-330.

Currim, M., Sarin, R. K., 1989. Prospect versus utility. Management Science 35, 22-41.

Eckstein, Z., van den Berg, G. J., 2007. Empirical labor search: A survey. Journal of Econometrics 136, 531-564.

Hanoch, Y., Johnson, J. G., Wilke, A., 2006. Domain Specificity in Experimental Measures and Participant Recruitment. Psychological Science 17(4), 300-304.

Harris, C. R., Jenkins, M., Glaser, D., 2006. Gender Differences in Risk Assessment: Why do Women Take Fewer Risks than Men? Judgement and Decision Making 1(1), 48-63.

Harrison, G. W., Morgan, P., 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., Laury, S. K., 2002. Risk aversion and incentive effects. American Economic Review 92(5), 1644-1655.

Houser, D., Keane, M., McCabe, K., 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., Winter, J., 2004. How do behavioral assumptions affect structural inference? Evidence from a laboratory experiment. Journal of Business and Economic Statistics 22(1), 64-79.

Huck, S., G. Kirchsteiger, Oechssler, J., 2005. Learning to like what you have: Explaining the endowment effect. The Economic Journal 115, 689-702.

Johnson, E. J., Gaechter, S., Herrmann, A., 2006. Exploring the nature of loss aversion. IZA-Discussion Paper No. 2015, Institute for the Study of Labor, Bonn.

Johnson, J. G., Wilke, A., Weber, E. U., 2004. A domain-specific scale measuring risk perceptions, expected benefits, and perceived-risk attitude in German-speaking populations. Polish Psychological Bulletin 35, 153-172.

Kahneman, D., Knetsch, J. L., Thaler, R. H., 1991. The endowment effect, loss aversion, and status quo bias: Anomalies. Journal of Economic Perspectives 5, 193-206.

Koebberling, V., Wakker, P. P., 2005. An index of loss aversion. Journal of Economic Theory 122, 119-131.

Kogut, C. A., 1990. Consumer search behavior and sunk costs. Journal of Economic Behavior and Organization 14, 381-392.

Lippman, S. A., McCall, J. J., 1976. The economics of job search: A survey. Economic Inquiry 14, 155-189.

Loewenstein, G. F., Weber, E. U., Hsee, C. K., Welch, N., 2001. Risk as feelings. Psychological Bulletin 127(2), 267-286.

Moon, P., Martin, A., 1990. Better heuristics for economic search: Experimental and simulation evidence. Journal of Behavioral Decision Making 3, 175-193.

Odean, T., 1998: Are investors reluctant to realize their losses. The Journal of Finance 53, 1775-1798.

Pennings, J. M. E., Smidts, A., 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., 2000: Risk aversion and expected utility theory: A calibration theorem. Econometrica 68(5), 1281-1292.

Rabin, M., 2002. Inference by believers in the law of small numbers. Quarterly Journal of Economics 117, 775-816.

Rabin, M., Thaler, R. H., 2001. Anomalies: Risk aversion. Journal of Economic Perspectives 15(1), 219-232.

Rapoport, A., Tversky, A., 1970. Choice behavior in an optimal stopping task. Organizational Behavior and Human Performance 5, 105-120.

Schmidt, U., Traub, S., 2002. An experimental test of loss aversion. Journal of Risk and Uncertainty 25(3), 233-249.

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

Schunk, D., 2008. Behavioral Heterogeneity in Search Situations: Theory and Experimental Evidence. University of Zurich, Working Paper.

Seale, D. A., Rapoport, A., 1997. Sequential decision making with relative ranks: An experimental investigation of the secretary problem. Organizational Behavior and Human Decision Processes 69(3), 221-236.

Seale, D. A., Rapoport, A., 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., Kahneman D., 1992. Advances in prospect theory: Cumulative representation of uncertainty. Journal of Risk and Uncertainty 5, 297-323.

Wakker, P. P., Deneffe, D., 1996. Eliciting von Neumann-Morgenstern utilities when probabilities are distorted or unknown. Management Science 42, 1131-1150.

Weber, E. U., Blais, A.-R., Betz, N. E., 2002. A domain-specific risk-attitude scale: Measuring risk perceptions and risk behaviors. Journal of Behavioral Decision Making 15(4), 263-290.

Wilke, A., Hutchinson, J. M. C., Todd, P. M., Kruger, D. J., 2006. Is risk taking used as a cue in mate choice? Evolutionary Psychology 4, 367-393.

Zwick, R., Rapoport, A., Lo, A. K. C., Muthukrishnan, A. V., 2003. Consumer sequential search: Not enough or too much? Marketing Science 22(4), 503-519.

Appendix

1 Details of the preference elicitation procedure

Individuals' utility functions for the gain and loss domains are determined using a parameter free method which is taken from Abdellaoui (2000). The method uses a series of lottery tasks to elicit subjects' utility functions in a parameter-free way. In total, the method consists of a series of 64 individually adapted lottery choice questions which are presented by the computer. 32 questions concern the utility function for gains and contain only lotteries with positive monetary payoffs. The other 32 questions concern the utility function for losses and contain only lotteries with negative monetary payoffs. Of the 32 question for gains or losses, respectively, 2 questions are only consistency check questions. That is, only 30 questions are being used for the elicitation of the utility function on gains or losses, respectively.

In our description of the elicitation procedure, we focus here on the 30 lottery choice questions that are used for the elicitation of the utility function on gains. The elicitation procedure for the utility function on losses works identically, only with negative numbers. The elicitation procedure is based on the construction of so-called standard sequences of outcomes, $(x_0, ..., x_6)$. Standard sequences of outcomes are monetary outcomes that are equally spaced in terms of their utility, and they are useful for the identification of the curvature of the utility function because of the equal spacing of the monetary outcomes (see Figure A.1 for an example).

** Include figure A.1 about here **

The first step of the procedure is to normalize $u(x_0) = 0$ and $u(x_6) = 1$, and to fix x_0 at a starting value. Then, we use 30 lottery choice questions in order to determine the values x_1, ..., and x_6. The construction of the standard sequence x_1, ..., and x_6 is based on the following idea: Assume that a person is indifferent between the two lotteries $A = (x_0, p; R, 1-p)$ and $B = (x_1, p; r, 1-p)$, and she is as well indifferent between $C = (x_1, p; R, 1-p)$ and $D = (x_2, p; r, 1-p)$. Assume further that preferences can be represented by cumulative prospect theory (CPT).[1] 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 the lotteries implies pairs of equations of the following type:

$$w(p)u(x_0) + (1 - w(p))u(R) = w(p)u(x_1) + (1 - w(p))u(r) \tag{7}$$

$$w(p)u(x_1) + (1 - w(p))u(R) = w(p)u(x_2) + (1 - w(p))u(r) \tag{8}$$

From these two equations follows:

$$u(x_1) - u(x_0) = u(x_2) - u(x_1) \tag{9}$$

That is, in terms of utility, the trade-off of x_0 for x_1 is equivalent to the trade-off of x_1 for x_2. This is a standard sequence of outcomes, $\{x_0, x_1, x_2\}$, which is – by construction – increasing for gains and decreasing for losses.

From these derivations, it is clear that the goal of our lottery choice questions is to determine a value x_1 for each experimental participant, such that she is indifferent between the lotteries A and B. Repeating this procedure for x_2, ..., x_6 will then result in the desired standard sequence, which allows for the estimation of the curvature of the utility function. In our procedure, 5 lottery choice decisions are used for the determination of each $x_i, i = 1, ..., 6$, i. e., we first determine a value x_1^1, then x_1^2 and so on until we have found a value x_1^5 from which we can compute the value x_1 that we are looking for. That is, we need 30 ($= 5 * 6$) questions in total for eliciting x_1, ..., x_6. In the following, we describe the elicitation of x_1, for which we use x_0 as a starting value. The elicitation method for all other x_i is identical, with the only difference that x_{i-1} and not x_0 is the respective starting value. In our method we set all parameters according to Abdellaoui (2000) and choose the following parameters which are held fixed during the whole experiment:

- All lotteries involve only the probabilities 66.7% and 33.3%, i. e., $p = 2/3$.
- R= €100, r=€0, Δ=€500.

[1] The elicited utility function on gains is, indeed, a von-Neumann-Morgenstern utility function. Equation (9) holds also under Expected Utility Theory, as can be found by substituting p for $w(p)$ in equations (7) and (8).

Table A.1 presents an example of the five lotteries that are used for the elicitation of x_1. Based on the parameters shown above and the starting value x_0 =€200, we see that the first decision that a subject has to take is to choose between lottery $A = (x_0, p; R, 1-p) =$ (€200, 2/3; €100, 1/3) and lottery $B = (x_0 + \Delta = x_1^1, p; r, 1-p) =$ (€700, 2/3;€0, 1/3). At this point of the procedure, we know that our value of interest, x_1, is in the interval $[200, 1200]$ (see the column labeled "Outcomes"). Assume now that the subject decides in favor of lottery A (see Table A.1). This means that x_1 >€700, i. e., we know that $x_1 \in$ [€700, €1200]. Therefore, we determine that for the second iteration step x_1^2 =€950, i. e., x_1^2 is the midpoint of the interval in which x_1 lies. Accordingly, in the second iteration step, the subject now has to decide between lottery $A =$ (€200, 2/3;€100, 1/3) and lottery $B =$ (€950, 2/3;€0, 1/3) (see the second row in Table A.2). For our example, we now assume that the subject prefers lottery B (see Table A.2). It is clear that this choice indicates that $x_1 \in$ [€700,€950], and we therefore determine for the third iteration step that x_1^2 =€820. This procedure continues (see Table A.2) and further narrows down the interval in which x_1 lies. After the fifth iteration step, we know that $x_1 \in$ [€880,€910] (see the last row in Table A.1). We then choose x_1 to be the midpoint of this interval, i.e €895. This value is rounded up to €900, and we use this value as the starting value for the next five iteration steps that will determine x_2. If we repeat this procedure for 6 times in total, we obtain our standard sequence of outcomes, $(x_0, ..., x_6)$, which is used for the estimation of the curvature of the utility function.

** Include table A.1 about here **

2 Derivation of the Decision Rules Used in the Search Tasks

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. We will proceed as follows: The first subsection introduces the models for these two cases under the simplifying assumption of risk neutrality. In the second subsection, we generalize the models for the two cases by relaxing the assumption of risk neutrality.

2.1 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, 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.[2]

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:[3]

$$G(m) = -\underbrace{[1 - F(m)]m}_{\otimes} - \underbrace{\int_{460}^{m} xdF(x)}_{\oplus} - c + m. \tag{10}$$

The term $\bigotimes$ 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 $\bigoplus$ 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,

$$G(460) = -c < 0. \tag{11}$$

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

$$G(540) > 0. \tag{12}$$

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

[3] 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.

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 (10) for m^*. Straightforward manipulation shows that the solution to this problem is identical to solving the following problem for m:

$$\pi(500 - m + 8) = (1 - F(m))\pi(500 - m - c + 8) + \int_{460}^{m} \pi(500 - x - c + 8)dF(x) \quad (13)$$

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:

$$\pi(x) = max\{0, x\} \quad (14)$$

In equation (13), 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, as the second case, 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:

$$\Pi(S_t) = max\{0, 500 - m - t \cdot c + 8\}. \quad (15)$$

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:

$$J_t(S_t) = max\{\Pi(S_t), E[J_{t+1}(S_{t+1})|S_t]\}. \quad (16)$$

$E(\cdot)$ 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.

2.2 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 (10) as an optimal stopping rule. Houser and Winter (2004) refer to a model of the form (16) 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 (13)[4]:

$$u(500 - m + 8) = (1 - F(m))u(500 - m - c + 8) + \int_{460}^{m} u(500 - x - c + 8)dF(x) \tag{17}$$

Equation (17) 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 in the main body of the paper shows the constant reservation price as a function of the risk-parameter γ in the exponential utility function (2). 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

[4] 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.

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 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:

$$\Pi^u(S_t) = max\{0, u(500 - m - t \cdot c + 8)\}. \tag{18}$$

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:

$$J_t^u = max\{\Pi^u(S_t), E[J_{t+1}^u(S_{t+1})|S_t]\}. \tag{19}$$

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, shown in the main body of the paper, plots the path of reservation prices, calculated by solving the dynamic discrete choice problem implied by equation (19) 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.

TABLE 1
Decision rules for the search problem

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

TABLE 2
Utility function estimation results and individuals' risk classification.

	Utility function	
	Gains (γ)	**Losses (δ)**
Median estimate	$2.003 \cdot 10^{-4}$	2.045 · 10-4
Mean R^2	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.

Type of heuristic by which behavior is explained best	Group name (N)	Preference Parameters	***Search Parameters*** Constant Reservation Price		Average Number of Searches		Search coefficient γ^{search}	
			Spearman - ρ	p-value	Spearman - ρ	p-value	Spearman - ρ	p-value
Type 1	**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
Type 2	**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
Type 3	**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
Type 1 or 2	**C4 (40)**	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
Any heuristic (Type 1, 2, 3, 4, or 5)	**All subjects (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

FIGURES AND TABLES

FIGURE 1
Optimal constant reservation price level depending on the individual risk coefficient γ in a model where the searcher treats past search cost as sunk cost.

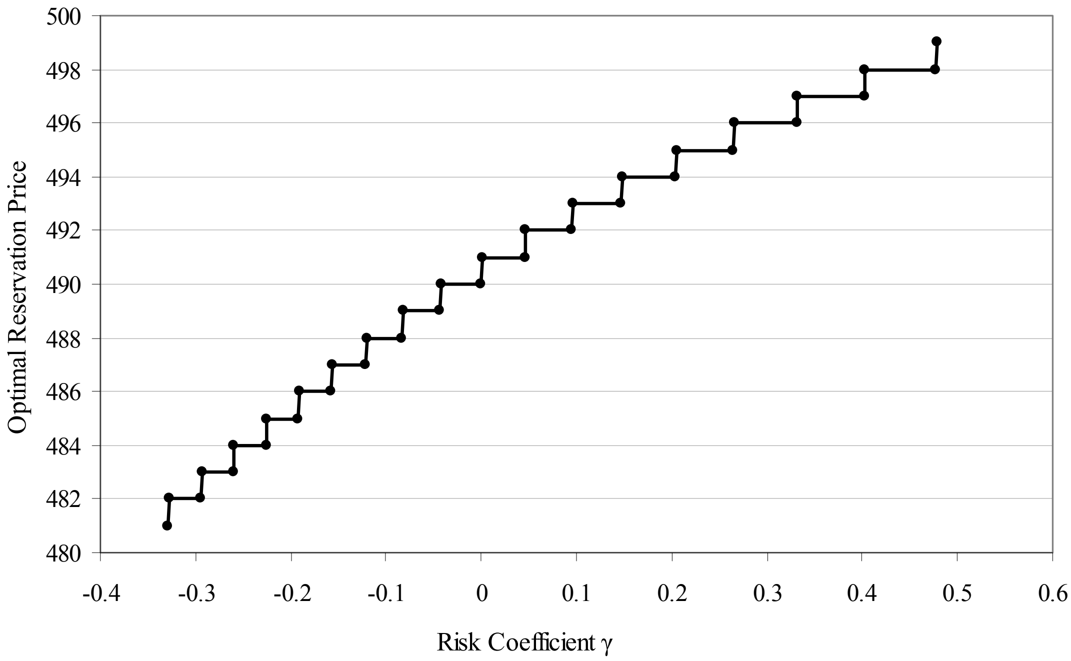

FIGURE 2
Optimal reservation price path depending on individual risk attitude coefficient γ in a finite horizon search model where the searcher does not treat past search cost as sunk cost.

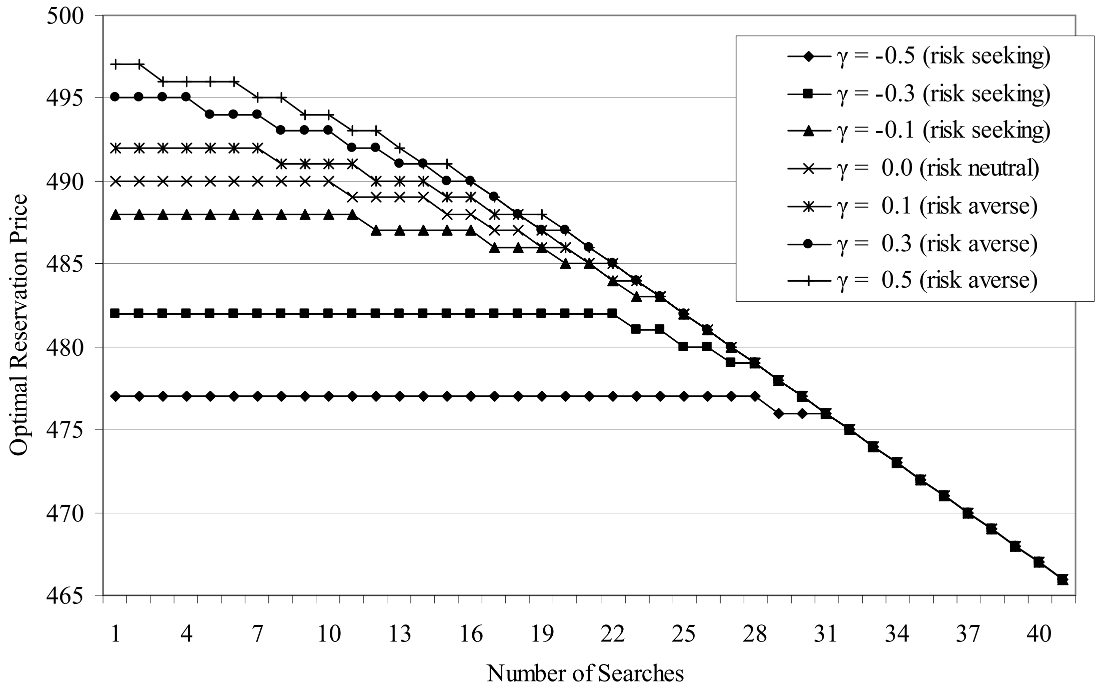

FIGURE 3
Distribution of the constant reservation prices observed in the experiment. The risk neutral reservation price is 490, lower reservation prices represent a risk seeking strategy, higher reservation prices show a risk averse strategy.

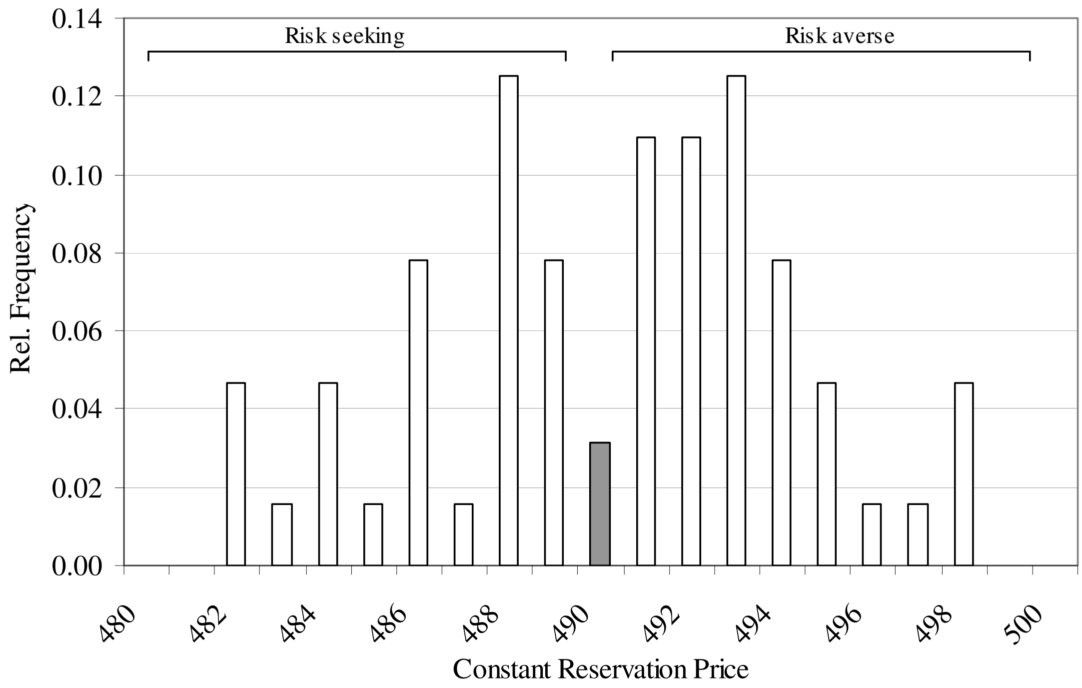

FIGURE 4

Number of subjects for whom a certain heuristic type is a best heuristic. The numbers in parentheses show the fractions of correctly explained choices for each heuristic type.

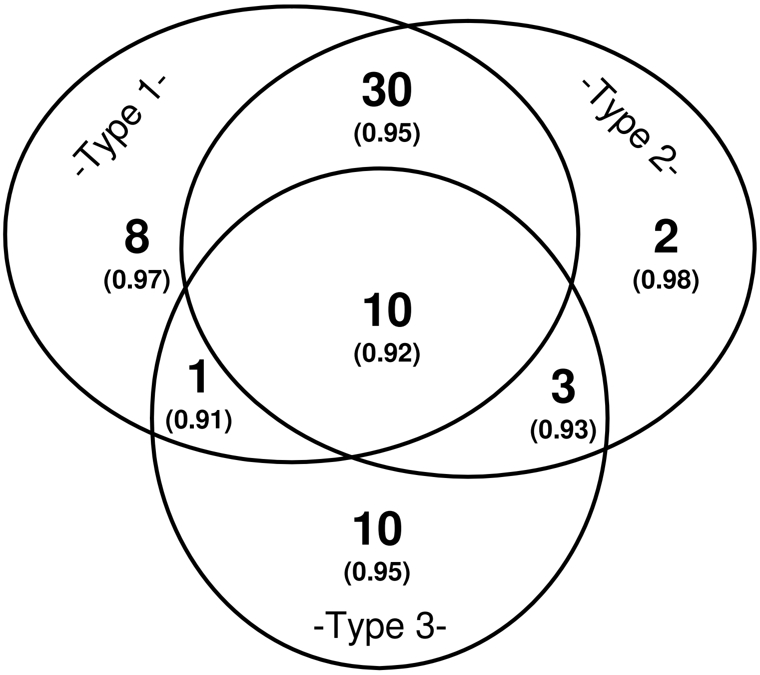

TABLE A.1
Assessing x_1 through bisection. An example of the Abdellaoui (2000) procedure.

Decision number (=Iteration step)	**Two alternative lotteries**	**Outcomes (€) $x_1 \in$**	**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 = € 200, p = 2/3, r = 0, R = € 100

FIGURE A.1
The x_i are a standard sequence of outcomes, since they are equally spaced in terms of their utility. This allows for the assessment of the curvature (i.e. the risk attitude) of the utility function.