

Uncertainty about Relative Ability in the Search Process

Armin Falk
IZA, University of Bonn
CEPR

David Huffman*
IZA, Bonn

Uwe Sunde
IZA, Bonn

October 13, 2006

Abstract

Standard search theory assumes that individuals know, with certainty, how they compare to competing searchers in terms of ability. In contrast, we hypothesize that searchers are uncertain about relative ability, with important implications for search behavior. We test our hypotheses in a laboratory experiment. The first main finding is that people are substantially uncertain about whether they are a type with a high or low probability of success, determined by being above or below the median in terms of ability. Self-confidence, defined as an individual's self-assessed probability of being a high type, is too high (above zero) for many low types, and too low (below 1) for many high types. Second, people update beliefs based on search outcomes. Self-confidence increases or decreases in the right direction, but is less sensitive to new information than predicted by Bayes' rule. Third, updating affects future search decisions: people are less likely to search as confidence about being a high type falls. Fourth, some search too little, and others search too much, due to wrong beliefs. Fifth, at the end of the experiment a substantial fraction turn down the chance to learn their exact rank. These are overwhelmingly those with low ability, suggesting an aversion to learning that one is one of the worst performers. Given that people are uncertain even in the simple setting of our experiment, our evidence strongly suggests that uncertainty about ability is relevant in more-complex, real-world search settings, including search for a job, search for a mate, or search for a journal that will publish one's paper. Focusing on the case of job search, we discuss how our findings can provide a new explanation for various important stylized facts from field evidence.

JEL-Classification: J64, D01, D83

Keywords: Job-search, Self-confidence, Discouraged workers, Subjective probabilities

*We thank the German Science Foundation (DFG) for financial support through SPP1169 (Potential for more Flexibility on Heterogeneous Labor Markets); Corresponding author: Armin Falk

1 Introduction

Standard search theory assumes that individuals are perfectly informed about abilities relative to competing searchers, and thus about their objective chances of success in search. If uncertainty is modeled at all, it is typically in the form of uncertainty about the shape of the wage distribution. In this paper we explore a simple alternative idea, namely that people are uncertain about whether they “have what it takes” when they are searching. We hypothesize that uncertainty about relative ability is the more appropriate behavioral assumption in many search settings. For example, uncertainty about ability is probably the rule rather than the exception in the case of a young researcher searching for journals that will publish her papers. In contrast to standard search theory, it seems likely that the researcher is not completely certain about her ability to meet the standards of top-tier journals, or to outperform the other researchers competing for space in these journals. Another example is a recently laid-off worker. The worker probably enters the new and unfamiliar labor market with substantial uncertainty about his ability to find a job. Searching for a mate may also involve uncertainty about ability, in this case own attractiveness relative to the competitors in the marriage market.

If searchers do have imperfect knowledge about relative ability, then we hypothesize that the search process unfolds in a way that is quite different from standard search theory. In particular, the process of search could change people: search outcomes are a signal about ability, and rational updating implies that people change their self-perceptions in response to this feedback. This updating would mean that unsuccessful search erodes confidence in relative ability, and thus that a searcher is increasingly likely to give up on search as the duration of unsuccessful search lengthens. For example, after a series of rejections, the young researcher may be sufficiently discouraged to switch to applying to lower tier journals, or choose another occupation altogether. Importantly, because search outcomes are likely to be only a noisy signal of ability, beliefs need not converge to the truth in the short run. Thus, the researcher might be overly discouraged by a run of bad luck, and stop search too early. Once she stops, she receives no new information, and thus never realizes her mistake. If there are systematic differences in confidence across different types of individuals, for example men and women, this could have important implications for explaining differences in the rates of discouragement, and economic outcomes, across

genders and other groups. With imperfect knowledge about relative ability, there could also be an additional, psychological cost to unsuccessful search, if people lose utility from receiving negative information about how they compare to others. This could make search a painful process, in a way that is missing from standard search theory.

We test our hypotheses using a laboratory experiment. An experiment is particularly useful in this case because the argument requires precise measurement of variables that are generally not available in field data, including subjective beliefs about relative ability, and the exact information that searchers receive in a given period of search. The experiment first measures a subject's ability on a very familiar task, and then informs the subject about own absolute performance on the task, but not the performance of others. Then the subject has a series of opportunities to engage in costly search. The subject knows that the probability of success in search depends on relative performance on the initial task: if she performed better than half of the other subjects, the probability is 0.6, and if she performed worse, the probability is only 0.3. The parameters of the experiment are such that a risk neutral individual who is certain of being above the median in ability (a high type) should always search, and someone who is certain of being below the median (a low type) should never search. We elicit subjective beliefs about relative performance before search begins, and at the end of each round of search, after the individual has learned the outcome. To our knowledge, this is the first search experiment to link success in search to a measure of real ability, and to directly elicit beliefs during the search process.¹

The first main result is that people are substantially uncertain about their relative abilities. If people were fully informed about their types, the distribution of beliefs observed in the experiment should be bi-modal, with half of the subjects correctly believing that they have a zero percent chance of being the high type and the other half believing they have a 100 percent chance of being the high type. Instead, the distribution of beliefs is more or less uniform, with many low types being overconfident, and many high types being underconfident. Given that people are uncertain even in the simple setting of our experiment, where there is only a single relevant dimension of ability, on a familiar task,

¹ There are previous job search experiments that study search behavior when the wage or price distribution is unknown (*E.g.*, Hey, 1982 and 1987), but this uncertainty is imposed rather than related to individual ability, and these experiments do not elicit beliefs directly, limiting the possibilities for verifying whether and how individuals update beliefs about wage distributions. For additional experimental evidence on search behavior, see *e.g.*, Schotter and Braunstein (1981); Cox and Oaxaca (1989) and (1996); Sonnemans (1998).

and the number of competing searchers is known, this is strong evidence that people must be at least as uncertain in the more complex setting of real-world search. We also find that women are significantly less confident in initial priors than men, despite performing slightly better on the test on average. This has the potential to lead to differences in search outcomes by gender, because as is discussed below, beliefs are the most important determinant of the decision of whether or not to continue search.

Second, we find that subjects update their beliefs about their relative ability depending on the outcomes of their search decisions. The subjective belief about the probability of being a high type significantly increases after a successful search episode and significantly decreases after an unsuccessful one. Although subjects update their beliefs in the right direction, the experimental evidence suggests that individuals are not fully consistent with Bayes' rule when updating about relative abilities. Instead, they are somewhat too conservative in changing self-perceptions, responding too little to new information provided by search outcomes.

Third, updating of beliefs has a systematic impact on future search decisions. The more confident subjects are of being a high type the more frequently they engage in costly search. The previous period belief is by far the most important predictor of the decision to search in the current period, dominating demographic factors and risk preferences. Thus, an individual's decision to search or not search is endogenous to the process, due to changing beliefs about his or her type.

Fourth, we find that subjects search too much or too little as compared to a situation in which they know their types with certainty. A substantial fraction of low types search, and some high types do not search. The low types who search are those who are overconfident, and believe that they have a high chance of being a high type. The high types who do not search are those who are underconfident, believing they have a low chance of being the high type. These findings illustrate the potential for people with the same objective type to endogenously arrive at different beliefs, and thereby make different search decisions.

The data also offer one piece of evidence that is suggestive of a psychological cost of unsuccessful search, above and beyond effort costs. At the end of the experiment we inform subjects whether they were a high or a low type, and then ask if they want to learn their exact rank on the test. Although this information is free, a substantial number of individuals turn it down. This information aversion is not random, but is systematically

related to relative ability; those who turn down the information are overwhelmingly those for whom the information might be most negative, i.e., those who scored in the bottom half on the test. This is consistent with subjects finding it psychologically painful to learn that they are one of the worst performers.

In summary, our findings provide a strong empirical basis for adopting imperfect knowledge about relative ability as a more realistic behavioral assumption for search. Furthermore, we identify systematic ways in which this uncertainty affects the search process. Although our findings are potentially relevant for any search setting where individuals are uncertain about ability, in a companion paper (Falk, Huffman, and Sunde, 2006b) we focus on job search as a particularly prominent example, and illustrate how our findings can be used to motivate a new general equilibrium model of the labor market in which workers are uncertain about their types. We describe this model briefly in the concluding remarks.

We believe that recognizing a role for uncertainty about ability in the search process opens up a valuable new perspective on a variety of important issues surrounding search. The empirical literature on job search is a prime example. For instance, our findings offer a new explanation for an important cost of long-term unemployment: the tendency for people to become increasingly unlikely to find a job as unemployment duration lengthens (see, *e.g.*, Frijters and van der Klaauw, 2006). Leading explanations in the literature focus on human capital depreciation or stigma. Although we believe that these are likely to be contributing factors, there is little direct empirical evidence for either (Machin and Manning, 1999). On the other hand, our experiment shows that falling self-confidence can generate this pattern, in a controlled setting where human capital depreciation and stigma can be ruled out by design. This explanation for negative duration dependence has not been discussed in the literature, presumably due to the maintained assumption of perfect knowledge about ability. Our findings also point to a different interpretation of what it means to be a nonparticipant, suggesting that some of those who are not searching for a job may in fact be high ability types who have wrong beliefs. These individuals are properly included in the social problem of unemployment, in contrast to definitions of unemployment based on criteria of active search. In fact, labor force surveys in the US identify a distinct state of non-participation called “discouragement”, in which a worker reports wanting a job, but says she is not searching because the probability of finding one

is too low (see Jones and Riddell, 1999.) Our experiment provides empirical evidence on the process that leads to discouragement, and suggests that some people can be overly discouraged.

A role for low confidence in sustaining long-term inactivity also raises new questions, and new challenges, for active labor market policy. For example, economists and policy makers are often skeptical about the effectiveness of job-training programs (see LaLonde, 2003). What happens if participants are similarly skeptical? To the extent that a searcher's effort is a decisive factor for finding a job, it is crucial that the searcher believes search is worthwhile. Unless a policy intervention creates a substantial improvement in the moral of participants, it may take only a short period of unsuccessful search to erode self-confidence again. Employment agencies could play an especially important role in this context. Agencies can draw on a large database of the search experiences of previous clients, and can potentially provide more reliable predictions regarding an individual's job-finding chances than the individual himself. However, it may then be necessary to design institutions that overcome various obstacles to sending credible signals about ability to searchers, for example, a perception by searchers that the agency has an incentive to maximize search effort among its clients.

The importance of confidence for search also points to a new explanation for differences in economic outcomes across different types of individuals, or groups, in society. For example, our finding of a gender difference in priors implies faster discouragement, and greater rates of non-participation among women, consistent with field evidence on gender differences in job search (see Bowlus, 1997; Frijters and van der Klaauw, 2006). Lower confidence could also play a role in explaining the gender wage gap. As illustrated in the general equilibrium model developed in our companion paper, confidence about ability to find a job affects wages, because it changes the worker's perceived outside option in wage bargaining with the firm. Thus, wages fall with unemployment duration, due to declining self-confidence. Lower confidence among women would lead, *ceteris paribus*, to lower wages at any point along the unemployment history.

Finally, evidence of a psychological cost of unsuccessful search provides empirical support for recent theoretical efforts to model search with "ego utility", i.e., a preference for positive beliefs about the self (for a discussion of ego utility, see Köszegi, 2006). For example, Andolfatto *et al.* (2004) incorporate ego utility into a two-period model of job

search. Ego costs of unsuccessful search imply that people may give up on search even more quickly, especially if they become increasingly sensitive to new information as they grow more pessimistic about ability. Incorporating ego utility into search models also has the potential to provide a link between search theory and evidence from psychology showing that subjective well-being, and mental health, tend to drop substantially for the unemployed, even controlling for the loss of income (see, *e.g.*, Winkelmann and Winkelmann, 1998; Gerdtham and Johannesson, 2003; Bjorklund and Eriksson, 1998; Mathers and Schofield, 1998). By assuming perfect knowledge about ability, standard search theory rules out this channel for search to affect well being, and has little to say about these important social costs of unemployment.

The rest of the paper is organized as follows. Section 2 describes the empirical design and the hypotheses tested in the experiment. Section 3 presents our results. Section 4 illustrates how the findings can be incorporated into a model of job search, and discusses implications. Section 5 concludes.

2 An Experiment on Search with Imperfect Knowledge about Ability

A laboratory experiment is a useful alternative to field data for answering our research questions, for several reasons. First, in existing field data it is typically not possible to find information on an individual's subjective beliefs about relative ability and about chances of success in search. Even more difficult is finding information on an individual's certainty about these beliefs, which is relevant for the present study. Second, an experiment allows us to study the impact of information about search outcomes on subjective beliefs about relative ability, in a tightly controlled way that rules out other confounding factors that might be present in the field. In the labor market, *e.g.*, objective job-finding chances may decline with unemployment duration, due to deterioration of objective skills (human capital depreciation) or employer stigma. Our experiment rules out these confounding factors by design and thus achieves a more clean study of the role of imperfect self-knowledge than is likely to be possible using field data. Third, the experiment controls precisely for other key variables in the search decision, including risk preferences, and effort costs, which are seldom directly observable in the field.

The experiment consists of three main stages. In the first stage we elicit subjects' risk attitudes with a simple lottery procedure. In the second stage it is (endogenously) determined whether a subject is of a high or low type. The third stage is the actual search phase. Here subjects can engage in costly search where the chances of success depend on the subject's type.

Stage 1: In this stage we measured subjects' attitudes towards risk. This is important since risk attitudes may affect search decisions, as has been pointed out by, e.g., Cox and Oaxaca (1989). The procedure we used is similar to the lottery procedure in Holt and Laury (2002). To our knowledge our search experiment is the first to control for risk attitudes in an incentive compatible manner. The lotteries in our measure of risk preference incorporated the same range of payoffs and probabilities that are associated with searching or not searching in the experiment, as will become evident when we discuss stage 3 below. Subjects were presented with a table that had 11 rows. In each row they had to decide between two lotteries. One lottery was the same in all 11 rows. This lottery was relatively risk free and paid out either 85 or 75 points with equal probability. The other lottery involved a chance of winning 200 points or zero points. The chances of winning 200 points changed from row to row. In the first row the chance of winning 200 points was relatively high: 60 percent. In the second row the probability dropped to 57 percent, in the third it was 54 percent, down to 30 percent by row eleven. If subjects have monotonous preferences, they prefer the more risky lottery up to a certain point, and then switch to preferring the relatively safe lottery in all subsequent rows. The switching row informs us about a subject's risk preference: a more risk-loving subject is willing to switch later, when the probability of winning the big prize has dropped relatively far. It is also possible of course that subjects are very risk averse, and prefer the first lottery immediately, in which case they are assigned a switching row of 1. Relatively risk loving subjects may prefer the second lottery even in the final row, and thus are assigned a switching row of 12. Subjects knew that after they had made a decision for each row, it would be randomly determined which row would be relevant for payment. This procedure guaranteed that each decision was incentive compatible. Subjects were not informed about the outcome of the lottery in stage 1 until the very end of the experiment.

Stage 2: The purpose of this stage was to determine subjects' types, which depended on how well they were able to perform simple calculation tasks. All subjects were asked to

solve a series of multiplication problems involving the product of a one-digit with a two-digit number, i.e., a problem of the form $x*yz$. They had to solve as many of these questions as possible in five minutes, without pencil or paper or any other helping device. For each correctly solved problem they received 25 points. All problems were presented to subjects on the computer screen. They could type in their answer in a box on the screen and confirm the answer by clicking an “OK”-button with their mouse. Upon having entered the answer in that way, a subject was informed whether the answer was correct. If it was correct, a new problem appeared automatically on the screen. If the answer was wrong, a subject had to work on the same problem again until the correct answer was entered. We forced subjects to solve a problem before a new question appeared on the screen in order to avoid subjects using a strategy of guessing, or searching for “easy” problems. Information about the current number of correct answers was presented on the screen throughout the test. Subjects were also told that the experiment would continue after stage 2 but were not given any details. They were, however, informed that the amount of correctly answered questions would, amongst other factors, determine “how easy it is for you to earn money later on in the experiment”. At the end of the test, subjects were told the total number of problems that they solved correctly, and the resulting payoff. Note that there was little scope for uncertainty about own absolute ability, given that performance on the task was easy for subjects to assess, and subjects were reminded explicitly about their performance at the end of the test. In this respect, the design provides a relatively hard test of the hypothesis that people are uncertain about relative ability. If subjects had not been informed about absolute performance, we are confident that this would have lead to even greater uncertainty about relative ability.

Stage 3: This was the actual search part of the experiment, consisting of eight periods. In each period subjects were endowed with 80 points. They had two options, to search or not search. If they decided to search they incurred search costs of 80 points and had a chance to receive either 200 points, in case of a successful search, or 0 points if they were unsuccessful. Importantly, subjects knew that the chance of being successful when searching depended on their type. A “high” type had a 60 percent chance of being successful, while a “low” type only had a 30 percent chance. Whether a subject was a high or a low type depended on relative performance on the math test in stage 2. High types were those subjects who solved more multiplication problems correctly in stage 2 than

half of all other participants in a session. Low types were defined analogously. In case a subject did not search in a given period this subject could either keep the endowment of 80 points or participate in a lottery, where with equal probability the outcome was either 75 or 85 points. We included this lottery option to help keep the potential “game value” of choosing not to search comparable to the case of choosing to search, i.e., to prevent the option of not searching from being relatively more “boring”. Importantly, the chance of winning in this lottery was independent of the individual’s type. Thus, subjects could only learn about their true types by searching.

Note that under risk neutrality, the parameter values of the experiment imply that someone who is certain of being a high type should always search, while someone who is certain of being a low type should never search. This follows straightforwardly from comparing the expected payoffs from searching in a given period to the payoff from keeping the 80 points and playing the simple lottery: $0.6 \cdot 200 > 0.5 \cdot 75 + 0.5 \cdot 85$ for high types; $0.3 \cdot 200 < 0.5 \cdot 75 + 0.5 \cdot 85$ for low types.²

Subjects were informed about the outcome, and their period income, in each period. In the case that they had searched they were informed about whether they had been successful or not, receiving either 200 or 0 points. In the case that they had not searched they were told that their period income was 80 or, if they had chosen to play the extra lottery, whether they had received 75 or 85 points. In addition subjects were presented with their search history, i.e., in each period they were reminded about the total number of times they had chosen to search, and the number of successful and unsuccessful search outcomes so far.

Before the first period of search, we asked subjects how confident they were of being a high type. Specifically, we asked each subject to indicate on an 11 point scale ranging from 0 percent to 100 percent, in steps of 10 percent, how likely they thought it was that they had answered more problems correctly than half of the other subjects in the room. For example, a subject who was completely confident of being a high type could indicate that the probability was 100 percent. A subject who was only moderately confident, and

² The value of search also includes a continuation value, above and beyond expected payoffs in the current period, once an individual is not completely certain about his or her type. This reflects the benefit of learning additional information. This is one channel through which type uncertainty can lead people to search too long, relative to the optimal decision under perfect certainty. We discuss the continuation value more formally in footnote 3 of the next section.

thought either type was equally likely, could instead indicate a probability of 50 percent. Even lower confidence could be expressed by choosing 20 percent, or 0 percent, etc. At the end of each subsequent round of search, after subjects had been informed about the outcome of that period, we asked again how likely they thought it was that they were a high type.

In summary, subjects decided in each of eight periods whether to search or not to search, were informed about the outcomes, and were then asked to indicate their subjective beliefs about the probability of being a high type. Asking beliefs in every period enabled us to study potential belief updating as a result of search outcomes. In addition we asked subjects about their prior of being a high type just before the first period. These answers inform us about subjects' initial priors without any information from search, as well as about the existence and amount of uncertainty about relative ability.

Investigating Psychological Costs of Search: After the last period subjects were informed about the outcome and the associated payoff from the lottery in stage 1, total payoffs from stages 2 and 3, and whether they had solved more problems than half of the subjects, i.e., whether they were a low or a high type. On a subsequent screen they were then asked if they wanted to learn the exact number of subjects in their session that had solved more problems in stage 2 correctly than they had. They were told that in case they wanted to learn their exact rank, the experimenter would tell them when they were paid (payment occurred in private, in a separate room). In case they wanted to learn their exact rank they could click on a "Yes-Box"; in case they did not want to know they could click on a "No-Box". Receiving this information was materially costless. However, there was a potential psychological cost in the sense that one could turn out to be among the worst performers, which could hurt self-esteem, or ego utility. We test whether those who turned down information were people who did poorly on the test, consistent with a desire to avoid finding out that one is one of the worst in terms of ability. The null hypothesis is that turning down the information was unrelated to performance on the test.

Procedural Details: The experiment was computerized using the software z-Tree (Fischbacher, 1999). Subjects were anonymous with respect to each other. The instructions were presented on the computer using a neutral frame, i.e., we did not use terms such as "search", "high types", etc.³ We ran two sessions with 22 subjects each for a total

³ Instructions for the experiment are presented in Appendix A.

of 44 independent observations. A session lasted, on average, about 40 minutes. Subjects were undergraduate students from different fields of study at the University of Bonn. Ten points in the experiment were exchanged for 0.08 Euro (1 Euro \sim 1.25 US Dollar). Average earnings for the 40 minutes, including a show-up fee of 2 Euro, were about 16 Euros.

3 Results

The experiment makes it possible to test a collection of hypotheses regarding the relevance of uncertainty about ability for search. The first hypothesis to be tested is whether subjects are in fact substantially uncertain about their relative ability, captured by their type. Denoting a subject's prior about the probability of being a high type by p^h , perfect information would imply zero confidence, $p^h = 0$, for low types and absolute confidence, $p^h = 1$, for high types, whereas limited information implies a substantial portion of both types with $0 < p^h < 1$. Next, we check whether subjects respond rationally to search outcomes by updating beliefs about themselves. If this is the case we should observe subjects updating their prior, p^h , in the direction predicted by Bayes' rule. It is also interesting to investigate whether updating is quantitatively correct according to Bayes' rule. Third, we test whether updating affects future search decisions. Beliefs are important for the search decision in the experiment, assuming subjects use a decision rule where search is only worthwhile if they are sufficiently confident that they are a high type, i.e., if their belief in period t is above a threshold probability for that period.⁴ This threshold may be higher or lower for different individuals, for instance due to differences in risk preferences, but if individuals compare their level of confidence to their personal threshold as hypothesized, we should observe that the likelihood of search decreases as confidence about being the

⁴ A subject's present discounted value of earnings in period t of the search phase can be written

$$\begin{cases} U_t = \sum_t^8 80 & \text{if the subject does not search} \\ S_t = \tilde{\lambda}_t \cdot 200 + \max \{S_{t+1}, U_{t+1}\} & \text{if the subject searches} \end{cases} \quad (1)$$

where $\tilde{\lambda} = .6 \cdot p^h + .3 \cdot (1 - p^h)$ is the individual's subjective probability of winning the prize, conditional on subjective beliefs about the probability of being the high type. The term in curly brackets is the continuation value of search, which in this case is the value of receiving an additional signal about one's true type. In the eighth period of the search phase, there is no continuation value, and thus it is clear that there is a critical probability of winning such that the individual is indifferent between searching and not. This in turn implies a critical level of confidence about being the high type, $p^h = 0.33$, such that a risk neutral individual searches only if $p^h > 0.33$. In earlier periods, there is a continuation value from search, so the threshold probability of success needed to induce search is lower. The threshold is lowest in the first period, when the searcher is most uncertain and the value of information is highest.

high type decreases. In a fourth step we explore the consequences that arise from the fact that people base their search decisions not on their true type, but on subjective beliefs about their type. We check the extent to which high types do not search, due to low confidence, while low types search, due to falsely inflated beliefs about their chances of being the high type. Finally, we assess support for the hypothesis that individuals who did poorly on the test will be systematically less willing to learn their exact rank, consistent with a psychological cost associated with learning that one is among the worst performers.

3.1 Uncertainty About Relative Ability

The first result concerns our basic hypothesis regarding imperfect knowledge about ability. Is there type uncertainty, or are people relatively accurate in their assessments of relative ability? Figure 1 provides an answer. The upper panel of the figure shows for both types, high and low, their subjective prior of being a high type. Without type uncertainty we would see a bimodal distribution, i.e., the subjects who are certain they are high types indicate a probability of 1 while those who have zero confidence of being a high type, and are sure that they are a low type, indicate a probability of zero. Instead the actual distribution is almost uniform. Only 18 percent of the subjects chose 0 or 1, i.e., the vast majority of subjects were at least to some extent uncertain about their type. The mean subjective prior is 0.57 with a standard deviation of 0.29. Thus we find a considerable degree of uncertainty about one's type, even in a relatively simple and well-defined environment. Interestingly, we also find evidence of a gender difference in confidence. Women performed slightly better on the test, answering on average 25.1 questions correctly compared to 24 correct on average among men. However, when we regress the subjective belief about the probability of being the high type on a dummy variable for female, and the number correctly answered on the test, the coefficient on female is negative and significant at the five percent level.⁵ This shows that women were less confident about being a high type before search began, for the same objective score on the test. This finding is consistent with previous evidence from psychology showing a similar gender difference in confidence (e.g., Deux and Farris, 1977; Lenney, 1977; Lundeberg et al., 1994). As we show below, lower beliefs about the probability of being a high type have important implications for search decisions, and thus could contribute to a gender

⁵ Results available upon request.

difference in search behavior.

The fact that the distribution of priors is almost uniform does not imply, however, that individual priors were just random or irrational. As the middle and lower panels of Figure 1 show, low and high types hold quite different priors. While the average prior of being a high type is 0.39 for the low types, it is 0.72 for the high types. The differences in priors between these two groups are highly significant (Mann Whitney test, $p < .001$). Another indication that priors were not just random comes from the fact that those with extreme outcomes on the test, relative to the median, were more certain of being a high or low type, respectively. The absolute difference between the median outcome and an individual's score is positively and significantly correlated with the individual's deviation from a prior of 0.5 (Spearman's $\rho = 0.57$, $p < .0001$). In other words, uncertainty about whether performance was above the median was most pronounced for those whose performance in stage 2 was close to the median, and lowest for those who were much better or much worse than the median.

3.2 Belief Updating

Given the existence of type uncertainty we now investigate whether beliefs about relative ability change with new information. Initial support for the hypothesis that subjects learned about ability in the process of search can be obtained by comparing initial beliefs and beliefs held in the final period. At the end of the last period of search 34 percent of the subjects were certain about their type, which is almost twice as much as before the first period of search.

Table 1 provides an overview of all changes in subjective beliefs in the experiment. Given that we have 44 subjects who were asked about their subjective beliefs at the end of eight sequential periods, there are 352 potential occasions for belief changes. Column (1) shows that in 224 cases beliefs did not change (64 percent). In 64 cases subjects thought it had become more likely that they were a high type. Equally often subjects thought it had become more likely that they were a low type. The outcome of search is informative for subjects, in the sense that the conditional probability of being a high type decreases given an unsuccessful search outcome. Conversely, a successful search outcome increases the conditional probability. Thus, we hypothesized that search success would systematically affect the direction of belief change. Columns (2) and (3) of Table

1 show that this is indeed the case. Of those who changed their beliefs after a successful search outcome, 86 percent thought it had become more likely that they were a high type. Likewise, 81 percent of all subjects who changed their belief after an unsuccessful search outcome became less convinced of being a high type. Subjects who did not search did not change beliefs in 79 percent of the cases, and if they did, change beliefs these changes were relatively symmetric, which is more suggestive of measurement error than systematic updating (see Column (4)).

Regression results in Table 2 show that unsuccessful search has a statistically significant effect on beliefs. The dependent variable is the subjective belief at the end of period t . The regressions use only those beliefs formed after periods in which a subject actually searched. Since beliefs are elicited in steps of 10 percent we use an interval regression procedure to correct for censoring.⁶ All significance tests are based on robust standard errors adjusted for clustering on individual. In Column (1) we regress the dependent variable on the dummy variable “Unsuccessful search in t ”, which takes value one if the search outcome was unsuccessful and zero otherwise. The respective coefficient is negative and significant at any conventional level. This means that updated beliefs of being a high type are lower for subjects that had an unsuccessful search outcome compared to those who were successful. In the second column we add controls for outcomes previous to period t : the number of previous unsuccessful search outcomes as of period $t-1$, the number of previous search decisions as of period $t-1$, a linear time trend (period) and gender. The size and significance of the main coefficient of interest remains essentially unchanged. In addition the number of previous search attempts enters positively and significantly. This finding reflects the fact that subjects who invested frequently had relatively high beliefs.

In Column (3) we take a slightly different perspective by focusing on changes in beliefs. To do so we regress the subjective belief of being a high type in period t on the respective belief in the previous period ($t-1$), in addition to the unsuccessful search dummy. Controlling for the previous period’s belief, the coefficient on unsuccessful search is negative and significant. This indicates that subjects changed their beliefs relative to previously held beliefs based on search outcomes, i.e., they became more pessimistic about being a high type after an unsuccessful search episode. Note that if subjects used previous outcomes for their belief updating, the coefficient on previous beliefs (belief in $t-1$) should

⁶ We find qualitatively similar results if we use OLS instead.

also capture information about previous unsuccessful search outcomes and the number of search decisions. We test this conjecture in Column (4) where we regress current beliefs on previous beliefs and previous outcomes simultaneously. It turns out that previous beliefs are positive and highly significant, while the other previous outcome coefficients are now insignificant. This suggests that previous beliefs actually comprised information from previous outcomes, i.e., subjects updated according to outcomes. As before, the dummy variable “Unsuccessful search in t ” in Column (4) is negative and significant. In summary, these findings establish that people do respond to search outcomes by updating beliefs about relative ability, in a way that is qualitatively consistent with rational updating.

Although changes in beliefs are qualitatively consistent with rational updating, an interesting question is how close subjects came in *quantitative* terms to the benchmark of fully rational, Bayesian updating. Previous studies have tested whether Bayes’ rule is a good quantitative description of updating, and have typically concluded that people are too sensitive to new information compared to Bayes’ rule (see, e.g., Grether, 1980; El-Gamal and Grether, 1995; Kahneman and Tversky, 1972; Tversky and Kahneman, 1971 and 1973).⁷ To our knowledge, however, no previous study has undertaken a quantitative analysis of updating when the topic is an individual’s abilities relative to peers. It could be that updating works differently in this case, because updating about relative ability may be less psychologically neutral than updating about, e.g., the probability that a taxi cab is a certain color. Although there are several caveats, discussed below, the data from this experiment is suitable for providing an initial exploration of how close subjects come to using Bayes’ rule when updating beliefs about relative ability.

In order to test whether individuals adhered strictly to Bayes’ rule, aside from some random calculation error, we adopt the approach of Grether (1980) and estimate a structural model of belief formation, which uses the odds-ratio form of Bayes’ rule:

$$\frac{p_t^h(loss)}{1 - p_t^h(loss)} = e^\alpha \cdot \left(\frac{p_{t-1}^h(loss)}{1 - p_{t-1}^h(loss)} \right)^{\beta_1} \cdot \left(\frac{1 - \lambda^h}{1 - \lambda^l} \right)^{\beta_2} \cdot e^{u_{it}}. \quad (2)$$

Where $p_t^h(loss)$ is the subjective probability of being the high type, conditional on having been unsuccessful in period t , and the ratio of $p_t^h(loss)$ to $1 - p_t^h(loss)$ is the odds of being

⁷ El-Gamal and Grether (1995) find that the relative majority are too sensitive. However, there are also sub-samples of individuals who are consistent with Bayes’ rule, or are not sensitive enough to new information.

a high type conditional on a loss.⁸ The variables λ^h and λ^l denote the objective chances of being successful in search, conditional on being a high or low type, respectively. The advantage of the odds-ratio form is that taking logs decomposes the right hand side into separate components that are meaningful:

$$\ln \left(\frac{p_t^h(\text{loss})}{1 - p_t^h(\text{loss})} \right) = \alpha + \beta_1 \cdot \ln \left(\frac{p_{t-1}^h(\text{loss})}{1 - p_{t-1}^h(\text{loss})} \right) + \beta_2 \cdot \ln \left(\frac{1 - \lambda^h}{1 - \lambda^l} \right) + u_{it} . \quad (3)$$

The first logged term on the right hand side is the individual's prior subjective odds of being the high type, held at the beginning of period t . The second term is the likelihood ratio, reflecting the impact of new information, *i.e.*, the search outcome, in period t .⁹ If a subject forms her posterior odds according to Bayes' rule, the parameter α on the right hand side should be equal to zero, β_1 and β_2 should be equal to 1, and calculation error, captured by u_{it} should be equal to zero. We estimate the coefficients of the model using OLS, with robust standard errors adjusted for clustering on individual.¹⁰

Consistent with Bayesian updating, we find $\alpha = -0.006$, which is not significantly different from zero ($p < 0.946$), and $\beta_1 = 0.944$, which is not significantly different from 1 ($p < 0.16$). However, for the coefficient reflecting the impact of new information, we find $\beta_2 = 0.347$, which is significantly less than 1 at any conventional significance level ($p < 0.001$). This suggests that subjects in the experiment were somewhat too conservative in updating their beliefs in response to new information.

One caveat in interpreting these results is that we did not elicit beliefs about the probability of being a high type in terms of a probability *distribution*. This would have been a more complete way to describe an individual's beliefs, but we opted for the simpler approach of asking for a single probability (which can be thought of as the mode of the distribution), because subjects typically find it very difficult to understand questions about probability distributions. Ignoring the weight that subjects put on other probabilities besides the one that they state in the experiment could bias our estimation results, although the direction of the bias is not clear. A second caveat is that we might under-

⁸ The odds of an event are the probability of that event divided by the probability of its negation.

⁹ In the case of successful search the likelihood ratio is the reciprocal of the one shown above.

¹⁰ We restrict the analysis to occasions where individuals chose to search. Because beliefs were elicited in intervals, prior and posterior odds used in the estimation are constructed using midpoints of belief intervals. For example, a stated belief of 0 or 1 is set equal to 0.025 or 0.975, respectively. Beliefs 0.1 to 0.9 are interpreted as midpoints of corresponding belief ranges and used directly, e.g., 0.1 is the midpoint of the range 0.05 to 0.15, 0.2 is the midpoint between 0.15 to 0.25, and so on.

estimate β_2 , and exaggerate the conservatism of subjects, because posterior beliefs are elicited in intervals. For example, even if a subject calculates beliefs correctly according to Bayes' rule, small changes in beliefs are censored, due to the intervals, which could lead the coefficient for β_2 to be less than 1. One way to estimate of the size of this bias is to calculate the correct Bayesian beliefs based on a subject's previous period beliefs and most recent search outcome, and then categorize these according to the same intervals used for eliciting beliefs in our data. Using these correct but censored beliefs as the dependent variable, we estimated the same regression as described above. The resulting coefficients show how much this bias affects our estimates. We find that β_2 is approximately equal to 0.753, indicating that censoring does bias the estimate of responsiveness to information downward. However, above we reported that subjects' actual beliefs implied a β_2 of 0.347. Thus, subjects are still too conservative, even after taking interval elicitation into account.

Keeping in mind the potential caveats, the evidence we find is suggestive of overly conservative updating, in which people place too little weight on new information. Interestingly, our findings contrast with previous studies of updating about other topics, which conclude that people are on average overly sensitive to new information in Bayesian updating problems. A possible explanation could be that people have a tendency to overestimate the extent of their knowledge about own and relative ability, which leads them to ignore new information. The observed conservatism could also reflect the fact that people tend to neglect information about their relative abilities if this information is unfavorable (see also Section 3.5, on aversion to negative information).

3.3 Is Search Affected by Subjective Beliefs?

Given that beliefs are updated in the process of search, we next investigate whether these updated beliefs affect future search decisions. We hypothesize that search activity should be positively correlated with confidence that one is a high type. Figure 2 confirms this prediction. It shows the fraction of subjects who searched in period t , depending on beliefs held in period $t-1$. This relation is strong and positive: the higher the belief of being a high type the more frequent was the search activity. Probit regressions (unreported) confirm that this relation is highly significant. When we regress a binary variable that is equal to 1 if the individual decides to search in period t on previous beliefs, the coefficient is

positive and significant at any conventional level ($p < .001$).¹¹ This result remains strong and significant if we control for risk preferences,¹² the number of times the individual searched as of $t-1$, the number of previous unsuccessful search outcomes as of $t-1$, the current period of the experiment and gender. Interestingly, none of the other variables turn out to be significant, although the sign of the coefficient for risk preferences indicates that those who are more risk averse are less likely to search, and women are less likely to search.¹³

3.4 Searching Too Much or Too Little Due to Type Uncertainty

In the experiment, risk neutral subjects who were certain of being the high type should always have searched. Similarly, subjects who knew they were low types should never have searched.¹⁴ Of course if subjects were extremely risk averse, they may not have wanted to search even if they were certain of being a high type. Similarly, very risk loving subjects may have preferred searching over receiving the safe option even if they know they were a low type. The elicited risk attitudes, however, help rule out both possibilities. First, only 3 out of the 44 subjects refused to play the lottery when the chance of winning was 60 percent, the probability that prevails if the individual is certain of being the high type. Second, only 2 out of 44 preferred to play the lottery when the chance of winning was only 30 percent, i.e., in conditions corresponding to certainty of being the low type. Thus for the great majority of subjects we can be relatively certain that risk preferences are such that they prefer not to search when they know that they are a high type, and prefer to search even if they know they are the low type.

Table 4 shows search decisions by type. While our lottery elicitation indicates that almost all subjects would prefer not to search given perfect certainty of being a low type, low types in the experiment searched in 50 percent of all possible cases. Put differently there was substantial excess search activity among low types, which would not exist if

¹¹ Significance results are based on robust standard errors adjusted for clustering on individuals. Regression results are available on request.

¹² To control for risk preferences we include for each subject the switching point in the lottery of stage 1. Two subjects did not have monotonous preferences and were therefore excluded from the analysis. If someone never played the more risky lottery they were assigned a switching point of one. Likewise if someone always played the risky lottery was assigned a switching point of 12.

¹³ Risk preferences become significant at the 10 percent level if we do not use robust standard errors.

¹⁴ See Section 2.

low types were perfectly informed about their true relative abilities. But there was also insufficient search on the side of high types. In 18 percent of all possible cases they did not search even though they had a high chance of success.¹⁵

Table 4 also displays the subjective beliefs of those who searched and those who did not search. Low types who searched had a substantially higher median belief of being a high type than those who did not (0.5 vs. 0.2). Among high types the differences are quite substantial as well. High types who did not search in particular periods had a median belief that places the probability of being a high type at only 0.4. In contrast, the ones that searched had a median subjective belief of 0.8. Moreover, low types who searched had a higher subjective belief than high types who did not (0.5 vs. 0.4). In summary, these results show that low types who search, and high types who do not search, have systematically overconfident or underconfident beliefs, respectively. This illustrates the potential for identical people come to think of themselves differently as a consequence of their individual search history, and exhibit different search behavior as a result.

3.5 Psychological Costs of Search

We conclude the results section with an investigation of potential psychological costs of search. To fix ideas, imagine that someone who searches for a job or a mate is often refused and told that he or she is unable or unattractive. This experience probably does more than just affect subjective beliefs about the probability of success. It may have a direct effect on utility. The psychological cost of receiving negative information about the self has been discussed in terms of “ego utility” by Köszegi (2006), who models an individual choosing between an ambitious task that is informative about ability and a less ambitious task that is less informative but also less remunerative. The model illustrates the potential for an individual to be less ambitious than is in his or her material self-interest, in order to avoid a loss in ego utility arising from negative information about the self. As noted by Andolfatto (2004) this type of information neglect could be relevant for search decisions as well. However, direct evidence on this type of psychological costs is, to the best of our knowledge, rare. We therefore investigate whether there is evidence for such psychological costs in our experiment.

¹⁵ This number is virtually unchanged if we exclude subjects who strictly prefer not to play the risky lottery in stage 1, even if the chance of winning is exactly 60 percent.

Recall that at the very end of the experiment subjects were asked whether they want to learn their exact rank. At this point in time, all they knew for sure was whether their performance was above or below the median in their session. Assume that most subjects are curious to know their exact ranking but at the same time that it is particularly harmful to learn that one is among the worst performers. If this is true, subjects who knew that they were above the median faced no risk of hurting their ego utility but those below did. If ego utility matters, information aversion should therefore have been more pronounced among low types compared to high types. Table 5 shows that this is in fact the case. Of the 44 subjects 15 did not want to learn their exact rank. 84 percent of these information averse subjects were low types (Column (2)). In contrast, among those who wanted to learn, only 31 percent were low types. Columns (3) to (5) further show that subjects who were information averse held lower final beliefs, had a worse expected rank, and a worse actual rank, than those who were information seeking. In other words, a higher chance of learning that one is the worst in the room reduces the willingness to learn about one's objective performance or ability.

Table 6 explores the determinants of information aversion using Probit models, where the dependent variable is equal to 1 if the subject turns down information and zero otherwise. The regressions confirm that weak performers were more likely to be information averse. Explanatory variables include a subject's belief about the probability of being the high type after the final period of search, the subject's average subjective belief over all search periods, and a more objective measure of performance, the number of correctly solved problems. The negative and significant coefficients reveal that subjects with lower beliefs and fewer correct answers were significantly more likely to be information averse. In Columns (4) to (6) we find similar results when we control for gender with a dummy variable that equals 1 for females and 0 for males.

4 Conclusion

Standard search theory assumes that individuals are perfectly informed about their relative abilities. In this paper we have explored a simple alternative idea: what happens if searchers have imperfect knowledge about ability. We ran a laboratory experiment where search takes place in a relatively simple environment. Nevertheless we found that subjects

were quite uncertain about their relative abilities, categorized as being a high or low type. Given the initial uncertainty, subjects updated beliefs about ability based on search outcomes, in the right direction but too conservatively compared to Bayes' rule. Based on updated beliefs subjects decided whether or not to continue searching. Since beliefs did not coincide with true types, some high ability types searched too little and some low ability types searched too long. Finally we found that a substantial proportion of individuals were averse to receiving information about ability if it was likely to be negative, i.e., reveal that they were one of the worst performers.

These findings provide a strong case for adopting imperfect knowledge about ability as a more realistic behavioral assumption for search. In a companion paper, we illustrate how our findings can be formalized, focusing on the case of job search (Falk, Huffman, and Sunde, 2006). In that model, we assume two types of workers, high and low, which only differ in one dimension: the probability of receiving a job offer. High types have a relatively high probability of receiving a job offer if they search, and low types have a low probability. The key assumption is that individuals do not know their own type with certainty. Hence, unemployed workers choose between costly search, or giving up on search and becoming a nonparticipant, on the basis of their *subjective* job finding probability. This probability is endogenous, depending on aggregate labor market conditions, but also on the individual's subjective belief about his or her own type. In line with our experimental findings, individuals are assumed to use information from past search outcomes to update their beliefs about their type. The longer an unemployed individual has searched unsuccessfully, the more certain the individual is of being a low type.

The optimal search strategy of individuals in this model is characterized by a threshold unemployment duration, after which beliefs about the probability of being a high type are sufficiently pessimistic that it is optimal to abandon costly search and become a nonparticipant. The threshold duration of unemployment in the model depends, among other things, on labor market conditions and the potential wages that the unemployed worker could earn conditional on finding a job. Notably, because search outcomes are only a noisy signal of an individual's type, some high types stop search and never realize their mistake.

Adding firms and moving to a general equilibrium version of the model generates additional insights, including a novel channel through which unemployment duration can affect worker's potential starting wages. We assume that successful matches between un-

employed workers and employers are generated in a frictional labor market with free entry, where firms try to fill vacancies. Wages are determined as the outcome of a bargaining game about the surplus, once unemployed and firms have been matched. The wage a worker with a given unemployment history can negotiate depends on his bargaining position. Workers with a long unemployment duration have a relatively low subjective job finding probability, and therefore a lower threat point than unemployed workers with a short unemployment record. Hence, the model predicts that increasing duration of unemployment leads to lower starting wages for workers, through the channel of falling self-confidence. Another novel implication of the general equilibrium model is greater volatility of unemployment in response to productivity shocks, due to an additional feedback mechanism that affects the threshold unemployment duration for workers.

Adopting type uncertainty in a model of the labor market therefore has several important implications for understanding labor market outcomes, such as an alternative explanation for why unemployed are increasingly unlikely to find a job, and are increasingly likely to choose nonparticipation, as unemployment duration increases; an explanation for a fall in starting wages as unemployment duration increases; and a mechanism that generates greater volatility of unemployment in response to productivity shocks than standard equilibrium search models. Potentially useful extensions include allowing for a gender difference in priors in the labor market, or a role for ego utility in the search decisions of workers.

References

- ANDOLFATTO, D., S. MONGRAIN, AND G. MYERS (2004): "Self-Esteem and Labour Market Choices," *mimeo*, Simon Fraser University.
- BJORKLUND, A., AND T. ERIKSSON (1998): "Unemployment and Mental Health: Evidence from research in the Nordic Countries," *Scandinavian Journal of Social Welfare*, 7, 219–235.
- BOWLUS, A. J. (1997): "A Search Interpretation of Male-Female Wage Differentials," *Journal of Labor Economics*, 15(4), 625–657.
- 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.
- (1996): "Testing Job Search Models: The Laboratory Approach," *Research in Labor Economics*, 15, 171–201.
- DEAUX, K., AND E. FARRIS (1977): "Attributing Causes for Ones Own Performance: The Effects of Sex, Norms, and Outcome," *Journal of Research in Personality*, 11, 59–72.
- EL-GAMAL, M. A., AND D. M. GREETHER (1995): "Are People Bayesian? Uncovering Behavioral Strategies," *Journal of the American Statistical Association*, 90, 1137–1145.
- FALK, A., D. HUFFMAN, AND U. SUNDE (2006): "Do I Have What it Takes? Equilibrium Search with Uncertainty About the Self," *mimeo*, IZA, Bonn.
- FISCHBACHER, U. (1999): "A Toolbox for Readymade Economic Experiments," *University of Zurich Working Paper*, 21.
- FRIJTERS, P., AND B. VAN DER KLAUW (2006): "Job Search with Non-Participation," *Economic Journal*, 116(508), 45–83.
- GERDTHAM, U.-G., AND M. JOHANNESSON (2003): "A Note on the Effect of Unemployment on Mortality," *Journal of Health Economics*, 22(3), 505–518.
- GREETHER, D. M. (1980): "Bayes Rule as a Descriptive Model: The Representativeness Heuristic," *The Quarterly Journal of Economics*, 95(3), 537–557.
- HEY, J. D. (1982): "Search for Rules for Search," *Journal of Economic Behavior and Organization*, 3(1), 65–81.
- (1987): "Still Searching: Search for Rules for Search," *Journal of Economic Behavior and Organization*, 8(1), 137–144.
- HOLT, C. A., AND S. K. LAURY (2002): "Risk Aversion and Incentive Effects," *American Economic Review*, 92(5), 1644–1655.
- JONES, S. R., AND C. W. RIDDELL (1999): "The Measurement of Unemployment: An Empirical Approach," *Econometrica*, 67(1), 147–162.
- KAHNEMAN, D., AND A. TVERSKY (1972): "Subjective Probability: A Judgement of Representativeness," *Cognitive Psychology*, 3, 430–454.

- KÖSZEGI, B. (2006): “Ego Utility, Overconfidence, and Task Choice,” *Journal of the European Economic Association*, forthcoming.
- LALONDE, R. (2003): “Employment and Training Programs,” in *Means Tested Transfer Programs in the U.S.*, ed. by M. Feldstein, and R. Moffitt.
- LENNEY, E. (1977): “Womens Self-Confidence in Achievement Settings,” *Psychological Bulletin*, 84, 1–13.
- LUNDEBERG, MARY, P. F., AND J. PUNCCOHAR (1994): “Highly Confident but Wrong: Gender Differences and Similarities in Confidence Judgments,” *Journal of Educational Psychology*, 86, 114–121.
- MACHIN, S., AND A. MANNING (1999): “The Causes and Consequences of Longterm Unemployment in Europe,” in *Handbook of Labor Economics*, ed. by O. C. Ashenfelter, and R. Layard, vol. 3.
- MATHERS, C., AND D. SCHOFIELD (1998): “The Health Consequences of Unemployment: The Evidence,” *Medical Journal of Australia*, 168, 178–182.
- SCHOTTER, A., AND Y. BRAUNSTEIN (1981): “Economic Search: An Experimental Study,” *Economic Inquiry*, 19, 1–25.
- SONNEMANS, J. (1998): “Strategies for Search,” *Journal of Economic Behavior and Organization*, 35, 309–332.
- TVERSKY, A., AND D. KAHNEMAN (1971): “Belief in the Law of Small Numbers,” *Psychological Bulletin*, pp. 105–110.
- (1973): “Availability: A Heuristic for Judging Frequency and Probability,” *Cognitive Psychology*, pp. 207–232.
- WINKELMANN, L., AND R. WINKELMANN (1998): “Why are the Unemployed so Unhappy? Evidence from Panel Data,” *Economica*, 65(257), 1–15.

Tables

Table 1: Changes in Subjective Beliefs: Probability of Being the High Type

	Overall (frequency)	Search was successful (fraction)	Search was unsuccessful (fraction)	Did not search (fraction)
No change	224	0.63	0.48	0.79
High type more probable	64	0.32	0.10	0.13
High type less probable	64	0.05	0.42	0.08

Table 2: Impact of Search Outcomes on Beliefs About Relative Ability

Dependent Variable: Subjective Probability of Being the High Type, End of Period t				
	(1)	(2)	(3)	(4)
Unsuccessful search in t	-0.109*** [0.037]	-0.100*** [0.033]	-0.077*** [0.022]	-0.076*** [0.022]
Belief in $t-1$			0.899*** [0.042]	0.876*** [0.055]
Number of losses as of $t-1$		-0.039 [0.040]		-0.013 [0.010]
Number of search attempts as of $t-1$		0.122*** [0.046]		0.015 [0.013]
Period		-0.080* [0.042]		-0.005 [0.009]
Female		0.065 [0.067]		0.017 [0.023]
Constant	0.672*** [0.039]	0.655*** [0.073]	0.047* [0.028]	0.044 [0.037]
log sigma	-1.373*** [0.067]	-1.456*** [0.069]	-2.103*** [0.168]	-2.113*** [0.163]
Observations	232	232	232	232

Notes: Interval regression coefficient estimates. The dependent variable is an individual's subjective belief about the probability of being the high type. Sample restricted to occasions where an individual searched. Robust standard errors are in brackets, adjusted for clustering at the individual level; ***, **, * indicate significance at 1-, 5-, and 10-percent level, respectively.

Table 3: Impact of Beliefs About Relative Ability on Search Behavior

	Dependent Variable: Search in Period t		
	(1)	(2)	(3)
Belief in $t-1$	1.139*** [0.157]	1.124*** [0.142]	1.010*** [0.126]
Risk Measure (Switching Value in Lottery)		-0.019 [0.020]	-0.020 [0.021]
Number of losses as of $t-1$			-0.083 [0.061]
Number of search attempts as of $t-1$			0.059 [0.041]
Period			0.139 [0.089]
Pseudo R ²	0.354	0.349	0.372
Log Likelihood	-145.82	-134.39	-129.55
Observations	352	328	328

Notes: Binary Probit marginal effects estimates. The dependent variable is 1 if an individual searches in period t and zero otherwise. Robust standard errors are in brackets, adjusted for clustering at the individual level; ***, **, * indicate significance at 1-, 5-, and 10-percent level, respectively.

Table 4: Searching Too Much or Too Little, Due to Wrong Beliefs

	Fraction of all occasions	Median belief: probability of being the high type
Low types searching	0.5	0.5
Low types not searching	0.5	0.2
High types searching	0.82	0.8
High types not searching	0.18	0.4

Table 5: Characteristics of the Information Averse

	Frequency	Fraction low types	Final belief	Average Expected rank	Average rank
Information averse	15	0.84	0.39	15.67	16.13
Information seeking	29	0.31	0.60	7.67	9.10


Table 6: Determinants of Information Aversion

Dependent Variable: Information Averse (binary measure)						
	(1)	(2)	(3)	(4)	(5)	(6)
Final belief	-0.411*			-0.372*		
	[0.212]			[0.219]		
Average belief		-0.632**			-0.599**	
		[0.259]			[0.258]	
Correct answers			-0.024**			-0.026**
			[0.007]			[0.007]
Gender				-0.136	-0.125	-0.286*
				[0.159]	[0.160]	[0.170]
log sigma	-26.34	-24.89	-21.46	-25.96	-24.57	-19.96
Observations	44	44	44	44	44	44

Notes: Probit marginal effects estimates. The dependent variable is equal to 1 if an individual turned down free information about exact rank on the math test, having previously learned whether the score was above or below the median for the group. Final belief is the individual's subjective belief about the probability of being the high type after the final period of search, and average belief is the average subjective belief over all eight periods of search. Correct answers is the individual's objective score on the math test. Robust standard errors are in brackets; ***, **, * indicate significance at 1-, 5-, and 10-percent level, respectively.


Figures

Figure 1: Uncertainty About Relative Ability



Notes: Subjects were assigned a high job finding probability in the search experiment (high type) if they scored higher than the median on an initial math test. After being informed about their own test score, but not the scores of others, subjects were asked: how likely do you think it is, in percentage terms, that you answered more questions correctly than half of the other subjects in the room today?

Figure 2: Impact of Beliefs on Search Decisions



Notes: The figure shows the fraction of individuals searching, as a function of previous period belief about the probability of being a high type (high types had a relatively high probability of success in search, whereas low types had a low probability).

A Appendix: Instructions for the Experiment

Welcome Screen

Welcome to the Experiment

During the entire experiment communication is prohibited. Please use the computer only for the functions intended for the experiment. Communication or playing around with the computer will lead to exclusion from the experiment.

For participating in this experiment you receive a show-up fee of **2.00 Euros**.

Depending on your decisions during the experiment you can earn additional money.

During the experiment we will not refer to Euros, but instead to ‘Talers’. Your earnings will be computed in terms of Talers during the experiment. At the end of the experiment, we will convert your total earnings from Talers into Euros, where

$$\mathbf{10\ Talers = 8\ Euro-Cents.}$$

The experiment has 4 phases. At the beginning of each phase you will be informed in detail about how the phase will proceed.

Once you have read this, please continue by clicking “Continue”.

Instructions for Phase 1:

In the first phase you will be asked to decide between playing different lotteries.

In the following you will be presented with 11 situations. In each situation you can decide to play one of two lotteries. Lottery 1 is the same across all situations. If you choose to play lottery 1 you can receive, with equal probability, either 85 Talers or 75 Talers. In other words: You receive 85 Talers with 50 percent probability, or 75 Talers with 50 percent probability.

Lottery 2, on the other hand, is different across situations. If you choose to play lottery 2 you can receive either 200 Talers or 0 Talers, but the probability of receiving 200 Taler is different in each situation. In the first situation, for example, the probability of receiving 200 Talers is 60 percent, while the probability of receiving 0 Talers is 40 percent. In each of the subsequent situations up to situation 11, the probability of winning 200 Talers **decreases** by 3 percent, and the probability of winning 0 Talers **increases** by 3 percent.

On the following screen you will be presented with the 11 situations. Decide in each situation which lottery you would like to choose. Once you have decided in all situations, the computer will choose one situation randomly for you. According to your decision for that situation you participate in one of the two lotteries. Therefore, each decision in each situation can be relevant for your earnings.

Please click “Continue” to begin making your decisions.

After clicking continue, subjects saw a screen with a table. In each row of the table they saw Lottery 1, which was always the same, and a Lottery 2, which varied from row to row. In the first row Lottery 2 was the following: win 200 with probability 60 percent, or win 0 with probability 40 percent, corresponding to the gamble faced in the search part of the experiment, for someone who is certain of being the high type. In subsequent rows the probability of winning in Lottery 2 reflected decreasing certainty of being the high type, in steps of 10 percent. In the final row, Lottery 2 involved a 30 percent chance of winning 200 and a 70 percent chance of winning nothing, corresponding to the gamble faced by someone certain of being the low type. A subject’s switching point in the table is an indicator of how certain they need to be about being a high type, in order to be willing to search, i.e., the switching point is a measure of risk preference, using the parameter ranges that are relevant in the search part of the experiment. After making a choice in each row, and clicking continue, subjects saw a screen reminding them that they would learn the outcome of their lottery at the end of the experiment.

Instructions for Phase 2:

In this phase you can earn money by fast computations.

In the following you have 5 minutes to solve as many multiplication exercises as you can. You have to solve the exercises in your head without using pencil or paper.

The more exercises you solve correctly the more money you can earn. For each correctly solved exercise you earn 25 Talers.

In addition, the number of exercises that you solve correctly will, among other factors, determine how easy it is for you to earn money later on in the experiment.

At the end of this phase you will be informed how many exercises you solved correctly.

After clicking "Continue" subjects saw a screen with a 10 second countdown before the beginning of the computations. At the end of the 10 seconds, subjects saw a screen with the first math problem. They could enter their answer in a box and continue by clicking 'OK'. After clicking 'OK', subjects were either told that their answer was incorrect and asked to try again, or were presented with the next math problem if their answer was correct. At all times, the remaining time for the math test was visible at the upper right corner of the screen. Also the screen always reported how many exercises the subject had solved correctly up to that point.

At the end of 5 minutes, subjects saw a screen that summarized how many exercises they solved correctly and their earnings from the math test. Subjects were not informed about the performance of other subjects.

Instructions for Phase 3:

This phase has 8 periods. Every period has three stages.

Stage 1:

In the first stage of every period you receive an endowment of 80 Talers. You then have to decide whether or not you want to “invest” the 80 Talers.

In case that you decide to invest, you can earn one of two possible amounts – which one depends on chance – the two possible outcomes are

200 Talers or 0 Talers.

Important: Your chance of earning 200 depends on how many exercises you solved correctly in phase 2 compared to the other participants.

In case you solved more exercises correctly in phase 2 than half of the other participants, you have a **60 percent chance** of winning 200. In case you did not solve more exercises correctly than half of the other participants, you have a **30 percent chance** of winning 200. In total, 22 people are participating in today’s experiment.

In other words, if you were better in phase 2 than at least 11 other participants you have a relatively high chance of winning (60 percent). If, on the other hand, you did not solve more exercises correctly than at least 11 other participants, you have a relatively low chance of winning (30 percent).

In case several participants have solved the same number of exercises correctly, the computer will determine randomly the ranking of these participants.

Only at the end of the third phase, after you have made all decisions, you will be informed about your relative performance in phase 2. Consequently, you will have to make all investment decisions based on your personal assessment of your relative performance.

In case you decide not to invest, you have to make another decision. In this case you can either keep the endowment of 80 Talers, or “**flip a coin**” with which you can earn either 5 more or 5 less than 80 Talers, with equal probability. In other words, you can choose between

80 Talers for sure or 85 Talers with 50 percent chance, 75 Talers with 50 percent chance.

Stage 2:

In this stage you are informed about the outcomes of your decisions on stage 1. In case you have invested, you are told whether you have earned 200 Talers or 0 Talers in this

period. In case you have not invested and “flipped the coin”, you are told whether you have earned 85 Talers or 75 Talers in this period. In case you have neither invested nor “flipped the coin” you earn 80 Talers.

Stage 3:

In this stage you are asked about your assessment of how likely it is that you solved more exercises correctly in phase 2 than at least 11 other participants.

After stage 3, the next period begins.

Before period 1 actually began, subjects saw the following screen, which asked for their assessment of the probability that they solved more multiplication exercises correctly than at least half of the participants in the experiment.

Before the first period begins, we would like to ask you the following:

How likely is it in your estimation that you solved more multiplication exercises correctly in the second phase than half of the other participants (in the experiment there are 22 participants, so we are asking how likely it is that you solved more exercises correctly than 11 other participants)?

Recall: You answered xxx exercises correctly.

Please indicate the percentage that is nearest to your estimate:

Subjects were then given 11 alternative answer categories ranging from 100% to 0%, in steps of 10%. For example, the first answer category read:

“With probability 100% I answered more exercises correctly than half of all participants.”

At the bottom of the screen, subjects were reminded that they would be informed whether they actually solved more exercises correctly than half of all participants at the end of phase 3.

In each of the 8 subsequent periods of search, subjects saw a screen asking whether they wanted to invest their endowment of 80 Talers, then a screen informing them about the outcome of their decision, and then a screen asking them for their subjective assessment whether they solved more multiplication exercises correctly than half of all participants.

After the 8 periods of search, subjects were informed about their earnings and whether they actually solved more multiplication exercises than half of all participants.

At the end of phase 3, subjects saw an additional screen, which asked whether they wanted to know exactly how many participants had solved more multiplication exercises correctly than them.

You have the possibility to learn the exact number of participants that solved more multiplication exercises correctly than you.

In case you would like to obtain this information, an experimenter will communicate this to you verbally when paying out your earnings.

Would you like to obtain this information from an experimenter?

The subjects could then tick a box indicating whether their answer was “Yes” or “No”.

In phase 4 that followed, subjects filled in a questionnaire asking them for socio-economic characteristics.