The Relevance and Implications of Imperfect Self-Knowledge for Search

Armin Falk
IZA, Bonn
CEPR

David Huffman*
IZA, Bonn

Uwe Sunde
IZA, Bonn

April 25, 2006

Abstract

Standard search theory assumes that individuals are perfectly informed about their own abilities. In contrast, we hypothesize that searchers have imperfect self-knowledge, and that this is potentially very important for understanding search behavior. We test our hypotheses in a laboratory experiment. We find that: (1) subjects are uncertain about their relative abilities, women more so than men; (2) subjects update beliefs about themselves based on search outcomes, in the right direction but too conservatively compared to Bayes' rule; (3) this updating affects willingness to continue search; (4) some high ability types to search too little and some low ability types to search too long, due to wrong beliefs; (5) a substantial number do not want to learn their exact rank among their competitors, even though this information is offered for free, and these are overwhelmingly those who are ranked low. These findings strongly suggest that imperfect self-knowledge is relevant for search outside of the lab, given that it is even more difficult to be certain of relative ability in many real-world search settings. In particular, our findings suggest that unsuccessful search is likely to lead to falling self-confidence and reduced willingness to search, and that this is likely to be more pronounced for women. Applying our findings to field evidence on job search helps explain falling exit rates out of unemployment, and explains greater non-participation rates and lower wages among women. Our findings also suggest that some non-participants in the labor market are in fact high ability types who have wrong beliefs, and that low ability types may search too long due to wrong beliefs, leading to congestion effects. Learning negative information about the self also appears to be psychologically painful, suggesting one reason why unemployment is associated with reduced life satisfaction and mental health problems. Our findings also provide a new perspective on other types of search, including search for a mate, or search for a publication in a top journal.

JEL-Classification:

Keywords: Job-search, Self-confidence, Subjective beliefs

*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 their own abilities, and thus about their objective chances of success in search. Until now, however, there has been little effort to test these assumptions, or undertake a systematic empirical investigation of how imperfect self-knowledge affects the search process. In other domains, by contrast, uncertainty about ability has been shown to have an important impact on behavior. For example, overconfidence about own abilities appears to contribute to excess entry, and the high frequency of small business failures (Camerer and Lovallo, 1999). A tendency for women to be under-confident about ability has been implicated as a cause of gender differences in performance under tournament compensation (Gneezy, Niederle and Rustichini, 2003). There has been some recent theoretical work on incorporating imperfect self-knowledge into search models, but little in the way of an empirical foundation.[1]

This paper presents results from a laboratory experiment designed to assess the relevance of imperfect self-knowledge for search. The experiment poses individuals with a prototypical search problem, namely deciding how long to engage in a series of costly trials, where the probability of success is linked to a real measure of relative ability. We elicit an individual's subjective beliefs about relative ability before search begins, and after the outcome of each subsequent round of search is realized.[2] The insights produced by the experiment are relevant for a variety of different types of search, e.g., search for a mate, or search for a journal that will publish a paper, but for concreteness we focus our discussion mainly on the example of job search.

In this single framework we can address a constellation of important questions that arise once one considers the possibility of search with imperfect self-knowledge. First, are individuals in fact substantially uncertain of their own relative abilities, even in a simple

[1] See, e.g., Andolfatto *et al.* (2004), for a two-period search model where searchers are uncertain about own abilities. Until now, the empirical foundation for such a model has been limited to different pieces of evidence from psychology, many of which are hypothetical or are drawn from settings far removed from search.

[2] To our knowledge, this experiment is the first to link success in search to ability, and to directly elicit beliefs during the search process. 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.

experimental setup? If so, this strongly suggests that searchers are at least as uncertain about their abilities and chances of success in the real world. Second, do people use search outcomes as a signal about own ability? This would imply a fundamental change in the nature of search, relative to standard search theory, because confidence about the self, and chances of success, become endogenous to the process. Third, how does updating compare to the Bayesian benchmark in this case? There is little evidence on how updating works when the subject of updating is a potentially psychologically sensitive topic, namely an individual's own ability. Fourth, given that search outcomes are only a noisy signal of ability, do some individuals converge to wrong beliefs, and search too long or too little as a result? Unlike standard search theory, this would imply that individuals could make mistakes in search, which are potentially very costly to the individual and to society. Fifth, is there a psychological cost of unsuccessful search, in the sense that people find it painful to receive negative information about ability? In contrast to standard theory, a cost of this kind implies that unsuccessful search, e.g., unemployment, has a direct negative impact on utility.[3] Sixth, is there a gender difference in confidence? If so, this could explain gender differences in search outcomes, and even lower wages, if lower confidence about job-finding chances affects the female threat point in wage bargaining.

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 ability and chances of success in search. Even more difficult is finding information on an individual's certainty about these beliefs, which is the key issue for the present study. Second, an experiment allows us to study the impact of information about search outcomes on subjective beliefs about the self, ruling out confounding factors that might be present in the field. For example, objective job-finding chances may decline with unemployment duration, due to deterioration of objective skills (human capital depreciation) or employer stigma. Our design rules excludes these confounding factors and thus achieves a clean study of the role of imperfect self-knowledge. 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.

[3] See Köszegi, 2002 and 2005, for theoretical work on preferences for positive beliefs about ability, or"ego utility."

Our experiment intentionally presents searchers with a relatively simple environment, and thus provides a hard test of the hypothesis that people are uncertain of relative abilities. Unlike in the field, searchers in our experiment know the exact number of competing searchers, and there is only a single relevant dimension of ability, which is easily defined. The particular dimension we study is one that student subjects are very experienced with, and have had strong incentives to learn about: ability to do simple mathematical calculations. Thus, if we observe uncertainty about ability in this simple environment, it provides a strong indication real-world searchers are uncertain as well.

In the experiment, we first measure risk preferences in an incentive compatible way. Then individuals take a math test that involves solving as many multiplication problems as possible in five minutes, without pencil or paper or other helping devices. After the test, subjects learn their own score but not the scores of the others in the room. Then they have several opportunities to engage in costly search. If they search, they can win a prize or they can win nothing. Importantly, the probability of success depends on their relative performance on the test. If they scored above the median they are a high type, and have are assigned a high chance of winning the prize. Otherwise they are a low type and have a low probability of winning.

The parameters of the experiment were chosen so that a risk neutral individual who is certain of being the high type should always search, and someone who is certain of being the low type should never search. Before the first period of search, we elicit initial priors, asking subjects how likely they think it is that they solved more problems correctly than half of the other subjects in the room. We then elicit the same beliefs after each subsequent round of search, when subjects have just been informed about the outcome of their most recent search decision. At the very end of the experiment, after being informed whether they are actually a high or low type, subjects are asked whether they want to be informed about their exact rank on the test.

The first main result is that people have substantial uncertainty about their relative abilities, even in the simple setting of the experiment. 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 half believing they have a 100 percent chance of being the high type. Instead, the distribution of beliefs is more or less uniform. We also find that women are significantly less confident in initial priors than men in the experiment, despite performing better on the test on average. This has important implications for differences in search outcomes by gender, because as 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 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 own abilities. Instead, they are somewhat too conservative, responding too little to new information provided by search outcomes.

Third, updating of beliefs has a systematic impact on future search decisions. The more subjects are convinced 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, the decision to search or not search is endogenous to the process, due to changing beliefs about the self.

Fourth, we find that subjects search too much or too little as compared to a situation in which they know their types. A substantial fraction of low types search, and some high types do not search, as a result of incorrect beliefs. The low types who search are those who believe that they have a high chance of being a high type, and the high types who do not search are those who believe 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, sometimes sub-optimal, search decisions.

The data also suggest that there is a psychological cost to unsuccessful search, above and beyond effort costs. We find that a substantial fraction of individuals do not want to know their exact rank on the test, after being informed wether they were in the top or bottom half. Overwhelmingly, these information averse individuals are those who scored in the bottom half on the test. Apparently, subjects are less willing to accept free information if there is a chance that it is very negative, e.g.., indicating that they are the worst in the room. An important implication of our finding is that unsuccessful search may be a painful process, and that searchers may try to avoid further negative information, for instance by stopping search.

In summary, our findings provide a strong empirical basis for adopting imperfect self-knowledge as a more realistic behavioral assumption for search. Furthermore, we identify systematic ways in which this uncertainty affects the search process. After our empirical analysis, we illustrate how these findings can be incorporated into a partial equilibrium search model, in a parsimonious way. In a companion paper, Falk, Huffman, and Sunde (2006), we extend this model to the general equilibrium case, within a matching framework.

We believe that changing the assumption of perfect self-knowledge has the potential to spark new discussions, and new research, on a variety of important issues surrounding search. For example, our findings suggest that search changes people, in the sense that they come to have different self-perceptions and beliefs about job-finding chances. This offers 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 exit rates arise even in the absence of these forces, due to falling self-confidence. This subjective mechanism for negative duration dependence has not been discussed in the literature, presumably due to the maintained assumption of perfect self-knowledge.

Given that low confidence may play a role in sustaining long-term inactivity, this 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 Heckman 1998). 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, if they are able to use data on the experiences of previous searchers to provide new searchers with reliable predictions regarding job-finding chances. However, it may then be necessary to design institutions that overcome various obstacles to sending credible signals about ability, for example, a perception by searchers that the agency has an agenda focused on maximizing search (Benabou and Tirole on signaling etc).

The importance of confidence in search also points to a new explanation for differences in economic outcomes across 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 one job search (see Bowlus, 1997). Incorporating this lower prior into our general equilibrium model (), which includes wage bargaining, implies lower wages for women due to greater pessimism about their outside option in terms of job finding chances. Until now, there has been little discussion of lower confidence in search as a source of gender differences in search outcomes, and the gender wage gap.

Evidence that unsuccessful search is psychologically painful, because it provides negative information about the self, supports recent efforts to model search with "ego utility", i.e., a preference for positive beliefs about the self (e.g., Andolfatto *et al.* 2004). 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, as suggested by results from our experiment. 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 the self, standard search theory effectively assumes that search outcomes are psychologically neutral, and thus has nothing 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 Self-Knowledge

2.1 Design of the Experiment

The search 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 success chances of search depend on the subject's type.

Stage 1: In this stage we measure subjects' attitudes towards risk. This is important since risk attitudes are likely to affect search decisions as has been pointed out, e.g., by Cox and Oaxaca (1989). To our knowledge our search experiment is the first to control for risk attitudes in an incentive compatible manner. The procedure we use is similar to the one in Holt and Laury (2002). The lottery in our measure incorporates parameters that correspond exactly to the payoffs associated with searching or not searching as will become evident when we discuss the actual search experiment in stage 3. Subjects are presented a table with 11 rows. In each row they have to decide between two lotteries. One lottery is the same in all 11 rows. This lottery is relatively risk free and pays out either 85 or 75 points with equal probability. The other lottery involves 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 is 60 percent, in the second row it is 57 percent, in the third it is 54 percent etc. up to 30 percent in row eleven. If subjects have monotonous preferences, they prefer the second lottery up to a certain level, and then switch to preferring the relatively safe option in all subsequent rows. The switching point informs us about a subject's risk attitude. It is also possible of course that very risk averse subjects always prefer the first lottery and never switch, whereas relatively risk loving subjects always prefer the second one. After a subject has made a decision for each row, it is randomly determined which row becomes relevant for payment. This procedure guarantees that each decision is incentive compatible. Subjects are 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 is to determine subjects' types, which depend on how well they are able to perform simple calculation tasks. All subjects are asked to solve a series of multiplication problems of the form $x*yz$. They have to solve these questions in their head and are not allowed to use any helping device. For each correctly solved problem they receive 25 points. All problems are presented to subjects on computer screens. They can type in their answer into a box and confirm the answer by clicking an "OK"-button with their mouse. Upon having entered the answer in that way, a subject is informed whether or not the answer is correct. If it is correct, a new problem appears automatically on the screen. If the answer is wrong, a subject has to tackle the same problem again until the correct answer is entered. We force subjects to solve a problem before a new question appears on the screen in order to avoid that subjects just guess or search for "easy" problems. Information about the number of correct answers so far is presented on the screen throughout the test. Stage 2 lasts for five minutes. Subjects are also told that the experiment continues after stage 2 but are not given any details. They are, however, told that the amount of correctly answered questions will, amongst other factors, determine "how easy it will be to earn money later in the experiment". After subjects have been told the number of correct answers and the resulting payoff, stage 3 begins.

Stage 3: This is the actual search part of the experiment. In each of the eight periods of stage 3 subjects are endowed with 80 points. They have two options, to search or not to search. If they decide to search they encounter search costs of 80 points and have a chance of receiving 200 points in case of a successful search or to get 0 points in case of an unsuccessful search. Importantly, the chance of being successful when searching depends on a subject's type. A "high" type has a 60 percent chance of being successful, while a "low" type has only a chance of 30 percent. Whether a subject is a high or a low type depends on the correctly solved problems in stage 2. High types are those subjects who solved more multiplication problems in stage 2 than half of all participants in a session. Low types are defined analogously. Thus all n participants in a session are partitioned such that the best $n/2$ are high and the worst $n/2$ are low types. Note that as assumed in our model, under risk neutrality and similar to our model high types should always search while low types should never search.

In case a subject does not search in a given period this subject can either keep the endowment of 80 points or participate in a lottery, where with equal chance the outcome is either 75 or 85 points. We included this lottery option in case of not searching to keep the potential "game value" between searching and not searching relatively similar. Note that the chance of winning in the lottery is independent of the type. In other words learning about one's type is possible only through searching.

Subjects are informed about the outcome and their period incomes in each period. In case of searching they are informed about whether they have been successful or not, receiving either 200 or zero points. In case they do not search they are told that their period income is 80 or, in case they have chosen the lottery option, whether they receive 75 or 85 points. In addition subjects are presented their search history, i.e., in each period they are informed about the total number of search decisions and the number of successful and unsuccessful search outcomes.

Before the first period of search, we asked subjects about their subjective belief regarding the probability of being a high type. More specifically we ask each subject to indicate on an 11 point scale ranging from zero percent to 100 percent in steps of 10 percent, how likely they think it is that they answered more problems correctly than half of the other subjects in the room. For example, a subject who is certain to be a high (low) type can indicate that the probability is 100 (0) percent. A subject who is completely unsure about his or her type would instead indicate a probability of 50 percent etc. After each subsequent round of search, after subjects had been informed about the outcome that period, we asked again for their subjective belief about the likelihood of being the high type.

In summary, subjects decide in each of eight periods whether to search or not to search, are informed about the outcomes, and are then asked to indicate their subjective beliefs of being a high type. Asking beliefs in every period enables us to study potential belief updating as a result of search outcomes. In addition we ask subjects about their prior of being a high type just before the first period. The answers inform us about subjects' initial priors without any information from search as well as about the existence and amount of type uncertainty in a simple and well-defined environment.

Eliciting Psychological Costs of Search: After the last period subjects are informed about the outcome and the associated payoff from the lottery in stage 1, total payoffs from stages 2 and 3 and whether they have solved more problems than half of the subjects, i.e., whether they are low or high type. On a separate screen they are then asked if they want to learn the exact number of subjects in their session that have solved more problems in stage 2 than they have. They are told that in case they want to learn their exact rank, the experimenter will tell them in private when paying them. In case they want to learn their exact rank they can click on a “Yes-Box”; in case they don’t they can click on a ”No-Box”. Receiving this information is materially costless. However, there is a potential psychological cost in the sense that learning to be close or actually the worst may hurt someone’s self-esteem. In fact we expected that subjects who know they are a low type are less likely than high types to ask for the exact rank information.

Procedural Details: The experiment was computerized using the software z-Tree (Fischbacher, 1999). All of the interaction was anonymous. The instructions were presented on the computer using a neutral frame, i.e., we did not speak of search, high or low types etc. We ran 2 sessions with 22 subjects each to end up with 44 independent observations. A session lasted, on average about 40 minutes. Subjects were students from different fields 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, including a show-up fee of 2 Euro, were about 16 Euro.

2.2 Hypotheses and Predictions

The experiment makes it possible to test a collection of hypotheses regarding the relevance and impact of imperfect self-knowledge for search. We summarize these hypotheses below. In the next section we turn to analysis of the data, and organize our discussion by assessing the support for each of these in turn.

The first hypothesis is that subjects have imperfect self-knowledge. Denoting a subject’s prior about the probability of being a high type by p^h, perfect self-knowledge would imply $p^h = 0$ for low types and $p^h = 1$ for high types, whereas limited self-knowledge implies a substantial portion of both types with $0 < p^h < 1$.

The second hypothesis is that subjects will 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. Third, we test whether updating is quantitatively correct according to Bayes’ rule.

The fourth hypothesis is that beliefs will have an impact on subsequent 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 certain of being the high type, i.e., if their belief in period t is above a threshold probability for that period, p_t^{h*}.[4] Depending on an individual's risk preferences this threshold may be higher or lower, 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 high type decreases.

A fifth hypothesis is that search decisions will not necessarily reflect true types. In other words, some high types will not to search, due to low confidence, and some low types will search, due to falsely inflated beliefs about their chances of being the high type.

Sixth, we assess support for the hypothesize that some individuals will not want to know their exact rank after being informed whether they are in the top or bottom half, because they want to avoid information that they are one of the worst performers. Consistent with this hypothesis would be a finding that those who avoid information are those who are low types, especially those low types who performed very poorly on the test, because for these individuals there is a greater likelihood of being informed that they were the worst.

3 Results

3.1 Uncertainty About the Self

The first result concerns our most basic hypothesis. Is there type uncertainty in our simple set-up or are subjects relatively well informed about their relative abilities? Figure 1 gives a definite answer. The upper panel of that figure shows for both types, high and low, their subjective prior of being a high type. Without type uncertainty we would see a bimodal

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

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

where $\tilde{p} = .6 \cdot p^h + .3 \cdot (1 - p^h)$ 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.

distribution, i.e., the subjects who are certain they are high types indicate a probability of 1 while those who are sure being a low type a probability of zero. Quite to the contrary the actual distribution is almost uniform. Only 18 percent of the subjects are certain about their type, i.e., the vast majority of subjects is at least to some extent uncertain about their type. The mean subjective prior is 0.57 with a standard deviation of 0.29.

The fact that the distribution of priors is almost uniform does not imply, however, that individual priors are just random or irrational. As the middle and lower panel of Figure 1 show, low and high types hold quite differen priors. While the average prior of being a high type is 0.39 percent for the low types it is 0.72 percent for the high types. The differences in priors between these two groups are highly significant (Mann Whitney test, $p<.001$). Another indication for the fact that priors are not just random can be inferred from the fact that the degree of certainty of being a high or low type is increasing with having more extreme outcomes with respect to the median outcome in stage 2. The correlation between the absolute difference between the median outcome and an individual's score and the deviation from a prior of 0.5 is highly significant (Spearman's rho = 0.57, $p<.0001$). In other words uncertainty is highest for those whose performance in stage 2 is close to the median and lowest for those who are much better or much worse than the median. Thus we find a considerable degree of uncertainty about one's type even in a relatively simple and well-defined environment. In the next section we investigate whether these priors change with new information as suggested by our model.

3.2 Belief Updating

Initial support for the hypothesis that subjects learn about themselves 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 34 percent of the subjects are certain about their type, which is almost twice as much as in the beginning when only 18 percent indicate being certain about their type.

An overview about all changes in subjective beliefs is provided by Table 1. Given that we have 44 subjects who are first asked about their priors and then in each of the eight periods about their subjective beliefs, there are 352 cases of potential belief changes. Column (1) shows that in 224 cases beliefs were not changed (64 percent). In 64 cases subjects thought it had become more likely that they are a high type. Equally often subjects thought it had become more likely that they are a low type. Given that the outcome of search is informative in the sense that conditional on an unsuccessful search the likelihood of being a high type is decreasing while it is increasing for a successful search outcome we expected that search success should 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 are a high type. Likewise, 81 percent of all subjects who changed their belief after an unsuccessful search outcome became less convinced of being a hight type. Subjects who did not invest did not change beliefs in 79 percent of the cases and if they changed beliefs, these changes were relatively symmetric (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. All significance results are based on robust standard errors adjusted for clustering on the 44 subjects. In column (1) we regress the dependent variable on the dummy variable "Unsuccessful search in t", which takes the value one if the search outcome was unsuccessful. 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 have an unsuccessful search outcome compared to those who were successful. In the second column we add as explanatory variables 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 investments enters positively and significantly. This finding simply reflects the fact that subjects who invest frequently have 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 in period t on the previous period's belief in addition to the unsuccessful search dummy. Controlling for the previous belief the coefficient on unsuccessful search is negative and significant, indicating that subjects become more pessimistic about being a high type after an unsuccessful search episode.

Note that if subjects use previous outcomes for their belief updating, the coefficient on previous beliefs should also capture the information from previous unsuccessful search outcomes and the number of search decisions. We test this claim in column (4) where we regress current beliefs jointly on previous beliefs and previous outcomes. 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 comprise information from previous outcomes, i.e., that subjects update according to outcomes. As before, the unsuccessful search success dummy in column (4) is negative and significant.

Thus far we have shown that subjects update their beliefs in the direction predicted by our hypothesis of rational updating. We now check whether belief formation is quantitatively consistent with rational updating, i.e., whether subjects act as though they calculate probabilities using Bayes' rule. In light of previous studies which have found that individual's typically update too little or too much compared to Bayes' rule, this is not a very likely result (see, *e.g.*, Grether, 1980, 1992; El-Gamal and Grether, 1995; Kahneman and Tversky, 1972; Tversky and Kahneman, 1971 and 1973). However, previous studies have not examined updating when the subject of uncertainty is an individual's own ability. If people have preferences for positive beliefs about the self, this could potentially have an impact on the updating process. Previous research has also not examined updating in a search framework. Thus, for the purposes of informing search theory, it is interesting to establish if and how individuals deviate from Bayesian updating in the current setting.

In order to test whether individuals adhere 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:[5]

$$\frac{q_t\left(h|loss\right)}{1-q_t\left(h|loss\right)} = e^{\alpha} \cdot \left(\frac{q_{t-1}\left(h|loss\right)}{1-q_{t-1}\left(h|loss\right)}\right)^{\beta_1} \cdot \left(\frac{1-p^h}{1-p^l}\right)^{\beta_2} \cdot e^{u_{it}} \, . \tag{2}$$

The advantage of the odds-ratio form is that taking logs decomposes the right hand side into separate components that are meaningful:

$$ln\left(\frac{q_t^h}{1-q_t^h}\right) = \alpha + \beta_1 \cdot ln\left(\frac{q_{t-1}^h}{1-q_{t-1}^h}\right) + \beta_2 \cdot ln\left(\frac{1-q^h}{1-q^l}\right) + u_{it} \, . \tag{3}$$

The left hand side is the log of the individual's subjective posterior odds of being the high type, conditional on unsuccessful search in period t. The first logged term on the right hand side is the individual's prior subjective odds of being the high type, held at

[5] The odds of an event are the probability of that event divided by the probability of its negation.

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.[6] 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.

In Table 4 we report estimates for the parameters of the updating model.[7] Column 1 shows that, consistent with Bayesian updating, α is not significantly different from zero and β_1 is not significantly different from 1. The coefficient reflecting the impact of new information, β_2, is significantly less than 1 at the one percent level, however, implying that subjects in the experiment are somewhat too conservative in updating their beliefs in response to new information. This finding is consistent with El-Gamal and Grether (1995), who find that although some people are too sensitive to information, a substantial portion of individuals are too conservative in updating their priors. In order to investigate whether people who did worse on the math test also deviate more from calculating Bayesian probabilities, Columns 2 and 3 report separate regressions for low and high types. Both types exhibit qualitatively similar updating, however, with β_2 significantly lower than 1, and differences in the coefficients between types are not statistically significant, lending little support to this hypothesis.[8] We return to Columns 4 and 5 of the table in the discussion of the psychological costs of search below.

One reason why we might underestimate β_2, and exaggerate the conservatism of subjects, is that 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, and thus 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

[6] In the case of successful search the likelihood ratio is the reciprocal of the one shown above.

[7] We estimate the coefficients using OLS, with robust standard errors adjusted for clustering on individual. Because beliefs are elicited in intervals, prior and posterior odds are constructed using midpoints of belief intervals. *E.g.*, 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.

[8] The significance tests are based on a regression for all investors together, both low and high types, but including interaction terms between high type and both log prior odds and the log likelihood ratio. Neither interaction term is statistically different from zero.

correct but censored beliefs as the dependent variable, we estimate the same regression as column (1) of Table 4. The resulting coefficients show how much this bias affects our estimates. We find that β_2 is approximately equal to 0.75, indicating that censoring does bias the estimate of responsiveness to information downward. However, Column 1 shows that subjects' actual beliefs implied a β_2 of about 0.35. Thus, subjects are still too conservative, even after taking interval elicitation into account.

To sum up, belief updating is qualitatively consistent with our hypotheses in the sense that beliefs are systematically affected by the success of the search outcome and previous search decisions. Belief updating is not perfectly rational, however, because subjects place too little weight on information provided by the search process.

3.3 Is Search Affected by Subjective Beliefs?

Given that beliefs are updated in the process of search, what needs to be shown next is whether these updated beliefs affect the search decision. According to our hypothesis, search activity should be positively correlated with the subjective belief of being a high type. Figure 2 confirms this prediction. It shows the fraction of subjects who search 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 is the search activity. Probit regressions 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).[9] This result remains strong and significant if we control for risk preferences[10], the number of times invested 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.

[9] Significance results are based on robust standard errors adjusted for clustering on the 44 subjects. Regression results are available on request.

[10] 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.

3.4 Inefficient Search Decisions Due to Type Uncertainty

In the experiment, risk neutral subjects who are certain of being the high type should always search. Similarly, subjects who know they are low types should never search. This follows straightforwardly from calculating the respective expected payoffs from searching (for high types: 0.6* 200>0.5 75+0.5*85 and 0.3*200<0.5*75+0.5*85 for low types). Of course if subjects are extremely risk averse, they may not want to search even if they are certain of being a high type. Similarly, very risk loving subjects may prefer searching over receiving the safe option even if they know they are a low type. With the help of the elicited risk attitudes, however, we can essentially rule out both possibilities. First note that risk attitudes are not a significant predictor of search decisions in the experiment, as was reported in the previous section. Second, 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. Also, 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 rule out 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 3 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 search in 50 percent of all possible cases. Put differently there is substantial excess search activity among low types, which would not exist if low types were perfectly informed about their true relative abilities. But there is also inefficiency on the side of the high types. In 18 percent of all possible cases they do not search even though they have a high chance of success.[11] These high types correspond to workers who have become overly discouraged by the search process. They do not search because they think their chances of being successful are not sufficiently large, despite the fact that their true chances are quite high.

Table 3 also displays the subjective beliefs of those who search and those who do not search. Low types who search have a substantially higher median belief of being a high type than those who don't (0.5 vs. 0.2). Among high types the differences are quite

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

substantial as well. High types who do not search in particular periods have a median belief that places the probability of being a high type at only 0.4. In contrast, the ones that search have a median subjective belief of 0.8. Moreover, low types who search have a higher subjective belief than high types who don't (0.5 vs. 0.4). In summary, these results confirm the that 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 os 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 unattractive. This experience probably does more than just affect subjective beliefs about the probability of success: it hurts directly. The psychological cost of receiving negative information about the self has been discussed in terms of "ego utility" by Köszegi(2002 and 2005), 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.

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 know that they are above the median face no risk of hurting their ego utility but those below do. Information neglect should therefore be more pronounced among the low types compared to the high types. Table 5 shows that this is in fact the case. Of the 44 subjects 15 don't want to learn their exact ranking. 84 percent of these information averse subjects are low types (column (2). In contrast among those who want to learn only 31 percent are low types. Columns (3) to (5) further show that subjects who are information averse hold lower final beliefs, have a lower expected rank and a lower actual rank than those who are 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 are 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 are 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 A Simple Model of Search with Imperfect Self-Knowledge

Consider a frictional labor market in discrete time. An unemployed worker has to decide whether to search actively or not in each given period t. The decision at the beginning of a given period about active search or inactivity is denoted $d_t = 1$, and $d_t = 0$, respectively. The unemployed worker makes his decision in order to maximize his expected discounted stream of available income, $\max_d E \sum_{t=0}^{\infty}(1-r)^t y_t$, where r is the discount rate.[12] To save on notation, let U denote the expected discounted income of an unemployed worker in a given period. This income stream depends on the unemployed worker's search decision at the beginning of the period. While unemployed, the worker receives a time-independent unemployment benefit $b_t = b$, sufficient to cover his basic needs low enough that he accepts any job offer he might get. If the unemployed worker continues to search actively, he has to incur a cost of $c_t = c$ for searching during a period, but with expected (subjective) probability $\tilde{p}$ he will find a job that offers an expected discounted lifetime income of W. If, on the other hand, the worker does not search, his future income stream includes

[12] As will become clear below, considering the case with infinite time horizon is without loss of generality. Under a finite time horizon, all results hold qualitatively unchanged, but the effects of uncertainty about one's type are even stronger.

the benefits he receives plus the discounted income from a potential job that he finds accidentally, with a baseline probability $\underline{p}$. This probability reflects the possibility that even without searching actively, an unemployed might encounter a job offer by a firm. The wage w_t is the same for every worker regardless of his type, and is determined exogenously.

To generate uncertainty, we assume that there are two types of workers. The low type workers have a strictly lower probability to encounter a job offer in a given period of search p^l than the high types with $0 \leq p^l < p^h < 1$. Both probabilities are common knowledge, but a worker does not know whether he is high or low type. Relevant for a worker's search decision is therefore solely his subjectively perceived job finding probability $\tilde{p}$. All workers have the same productivity on the job, however, the heterogeneity only affects the job finding probability, e.g. through attractiveness of these workers.[13] The value of an unemployed's expected future income then satisfies the Bellman equation

$$rU = \max\{b + \tilde{p}\,(W - U) - c, b + \underline{p}\,(W - U)\}\,. \tag{4}$$

Clearly, an unemployed worker continues to search actively if and only if the expected gains from search outweigh the expected gains of abandoning search. The corresponding value function for holding a job reads

$$rW = w + \delta(U - W)\,, \tag{5}$$

where w is the wage to be earned for every period on a job, and δ denotes the probability that the employment relation dissolves and the worker becomes unemployed. Assume for the time being that no dynamic changes in parameters are expected to affect the income stream of the worker on the job, W, or the income stream from unemployment, U, and assume that the income stream from benefits is exogenously given and time invariant. Using both value functions, one can solve for the expected income stream from unemployment as

$$rU = \max\left\{\frac{(b-c)(r+\delta) + \tilde{p}w}{r+\delta+\tilde{p}} \qquad , \frac{(b)(r+\delta) + \underline{p}w}{r+\delta+\underline{p}}\right\} : . \tag{6}$$

One can now express the search strategy of an unemployed worker by solving for the subjective job finding probability $\tilde{p}^*$ for which he is indifferent between active search

[13] The framework laid out here does not require assumptions about whether firms know the type of an unemployed worker. Even if firms perfectly observe all relevant components of the workers' profiles, they possibly discriminate between ex ante identical applicants, inducing uncertainty in the workers' perception about their marketability and job finding prospects, see Mailath, Samuelson, and Shaked (2000). In the present model, types could also be related to productivity, but we abstain from considering this case for illustrative purposes.

and inactivity. Using the above condition (6), this probability is given by

$$\tilde{p}^* = \underline{p} + \frac{(r + \delta + \underline{p})c}{w - b} \,. \tag{7}$$

If the individually perceived job finding probability exceeds this threshold, active search is optimal, if it is less than or equal to this threshold, the individual prefers not to search actively. Note that $\tilde{p}^*$ is time invariant as long as the right hand side of (7) is time invariant. Moreover, note that if search is costless, unemployed only find it optimal to search if active search leads to a higher job offer arrival rate than sitting at home and waiting for offers. In order for some unemployed to become discouraged from further active job search, search must be optimal for individuals with high job finding probability, while termination of search is optimal for individuals with low job finding probability, i.e. $1 > p^h > \tilde{p}^* > p^l \geq \underline{p} \geq 0$. Clearly, if $\tilde{p}^*$ were higher or lower than both job finding probabilities p^l and p^h, an unemployed would never even start to search, or would never terminate search, respectively.

At the beginning of each period, the unemployed worker observes whether his search effort in the previous period had been successful in terms of getting a job offer. Consequently, workers can use the information of the success of previous application rounds to update their belief about their true type, and therefore their job finding probability. For notational convenience, denote the observation of having received a job offer in the previous round as $\omega = 1$, and having received no offer as $\omega = 0$, respectively. In case the worker received an offer, he leaves unemployment and starts working on his new job. In case he received no job offer previously, the unemployed has to decide whether to continue searching for a job, or whether to terminate search. The worker's subjective job finding probability, which determines this decision, is a linear combination of the job finding probabilities of the two types of workers, p^l and p^h, weighted by the subjective assessment of the probability of belonging to one group or the other. Denoting the subjective probability of belonging to the group of high types, $\tilde{p}^h$, the subjective job finding probability is

$$\tilde{p} = p^l(1 - \tilde{p}^h) + p^h \tilde{p}^h \tag{8}$$

All job offers are immediately accepted. Therefore, learning about one's true type only pertains to those unemployed who are not successful in obtaining a job offer. Consider a worker, who has been unemployed for one period, and despite searching could not manage to get a job offer. Using Bayes' rule, his subjective probability of belonging to the group of high types is

$$\tilde{p}_1^h = Pr(h|\omega_0 = 0) = \frac{Pr(\omega_0 = 0|h)Pr(h)}{Pr(\omega_0 = 0|l)Pr(l) + Pr(\omega_0 = 0|h)Pr(h)} . \tag{9}$$

In order to calculate the posterior $\tilde{p}_1^l$, however, one needs to know the expressions for the conditional probabilities and priors. The probability of receiving only rejections during the first period of search conditional on being a low type is $Pr(\omega_1 = 0|h) = 1 - p^h$. Likewise, the probability of being rejected given that the worker belongs to the low types is $Pr(\omega_1 = 0|l) = 1 - p^l$. Furthermore, to compute $\tilde{p}_1^l$, a prior about the unconditional probabilities $Pr(h)$ and $Pr(l)$ is required. Let a fraction q^h of the population be low types, while a fraction $q^l = 1 - q^h$ are high types with job finding probability p^h. We assume in the following, that unemployed know these population shares and have a corresponding prior $Pr(h) = q^h$. Of course, any alternative prior is conceivable *a priori* and would affect the search decision. More generally, if a newly unemployed worker has a prior about their subjective job finding probability that induces him to search during the first period of unemployment, his posterior can be computed as in (9). We study the case of a flat prior for illustration only. Workers who lost their job and have too low a prior might directly go into inactivity (long-term unemployment) without ever searching.

In the following, we assume for simplicity that workers update their beliefs about themselves only downwards, i.e. when they receive negative information.[14] The conditional probabilities for unemployed who had searched unsuccessfully for n rounds are thus given by $Pr(\omega_0 = \omega_1 = \omega_2 = ... = \omega_{n-1} = 0|l) = (1 - p^l)^n$ and $Pr(\omega_0 = \omega_1 = \omega_2 = ... = \omega_{n-1} = 0|h) = (1 - p^h)^n$. Hence, the subjective probability of being a high type after an unemployment history of n periods, i.e. after n rounds of unsuccessful search (including period 0), is

$$\tilde{p}_n^h = \frac{(1 - p^h)^{n-1} q^h}{(1 - p^l)^{n-1} q^l + (1 - p^h)^{n-1} q^h} , \tag{10}$$

with $\lim_{n\to\infty} \tilde{p}_n^h = 0$, since $p^h > p^l$.[15] This also implies that $\lim_{n\to\infty} \tilde{p}_n^l = 1$. Using this result, and the fact that $\tilde{p}_n^h = 1 - \tilde{p}_n^l$, the subjective job finding probability after n rounds

[14] Of course, nothing precludes individuals to update their beliefs upwards if they are successful in finding a job. We neglect this possibility here for computational simplicity and to avoid complications arising from updating of beliefs by inactive searchers who accidentally find a job.

[15] Clearly, positive updating upon a successful job accession could be taken into account. This would affect the results only quantitatively and would complicate the analysis.

of failed search can be expressed as

$$\tilde{p}_n = p^l + (p^h - p^l)\tilde{p}_n^h \,. \tag{11}$$

This expression can be substituted for the critical subjective job finding probability that characterizes the optimal search strategy of a worker, who has been unemployed (and therefore been searching unsuccessfully) for n periods by employing condition (7). Solving for $\tilde{p}_n^h$ and substituting using the expression given in equation (10), the optimal search decision is eventually characterized by

$$d_n = \left\{ \begin{array}{c} 1 \\ 0 \end{array} \right\} \qquad \Leftrightarrow \qquad \frac{\tilde{p}^* - p^l}{p^h - p^l} \left\{ \begin{array}{c} < \\ \geq \end{array} \right\} \underbrace{\frac{(1-p^h)^{n-1}q^h}{(1-p^h)^{n-1}q^h + (1-p^l)^{n-1}q^l}}_{=(1-\tilde{p}_n^l(n))} \,. \tag{12}$$

From condition (7), the left hand side of the inequality determining the decision is a function of the parameters, while the right hand side additionally depends on the number periods of previous unsuccessful search, n.

This simple framework allows to generalize some of our experimental findings. For every individual, there exists exactly one threshold duration of unemployment, i.e. of unsuccessful search, n^*, from which on he stops searching and becomes discouraged and inactive. To see this result, take a closer look at the right hand side of the condition in (12), namely $(1 - \tilde{p}_n^l(n))$. Note that the derivative of $\tilde{p}_n^l$ with respect to n is positive, so that higher n shift the decision toward the termination region.[16] Moreover, it is clear from the time invariance of the threshold job finding probability $\tilde{p}^*$, that once an unemployed stops searching, he does so for good and becomes long-term unemployed. Thus, it is more likely that individuals with a longer search record and unemployment duration terminate search for good. The ultimate reason for 'giving up' is the individually perceived low probability of possessing a highly marketable profile, that is, essentially, lack of self-esteem reflected by a low $\tilde{p}^h$. It should be emphasized that this result is true *irrespective* of what type an unemployed person actually is, because only the individually experienced search history is relevant for the decision to continue or terminate search. The effects of parameters for the optimal decision to quit search implied by conditions (7) and (12) are rather intuitive.

[16] Note that

$$\frac{\partial(1-\tilde{p}_n^l)}{\partial n} = \frac{q^h q^l (1-p^h)^{n-1}(1-p^l)^{n-1}\left[\ln(1-p^h) - \ln(1-p^l)\right]}{\left(q^h(1-p^h)^{n-1} + q^l(1-p^l)^{n-1}\right)^2} < 0 \,,$$

since by assumption $0 \leq p^l < p^h < 1$.

The duration until discouragement n^* decreases in benefit payments b, search costs c, the discount factor r and the separation rate δ, and increases in the wage w, and the job finding probabilities p^h and p^l.[17]

Due to their beliefs about their job finding probability, unemployed workers give up searching too early. This finding is best illustrated with a little numerical example. Consider a world characterized by the following parameter values: $p^h = 0.25$, $p^l = 0.05$ (the good types have a five times larger instantaneous job finding probability), $w = 1.5$, $b = 1$, $c = 0.75$, and $r = 0.95$. Also assume that half of the population has a highly marketable profile, while the other half has low job finding probability, i.e. $q^h = q^l = 1/2$. In this environment, an unsuccessful unemployed terminates search after five rounds of search according to the continuation decision displayed in equation (12), and accepts living of unemployment benefits for the rest of his life. But after five rounds, the perceived probability of being a person with a highly marketable profile is still about ten percent! Statistically speaking, the unemployed stops searching before he can reject the null of being a good type on any level of significance.

5 Conclusion

Standard search theory assumes that individuals are perfectly informed about their own abilities. In this paper we have explored a simple idea: what happens if searchers have imperfect self-knowledge, regarding their own abilities and regarding their abilities relative to other competing searchers. We used a laboratory experiment to investigate, and found the following results: (1) subjects are uncertain about their relative abilities, even in the simple environment of the experiment, indicating that real-world searchers are uncertain as well; (2) women have lower priors about own ability than men, conditional on the same objective ability; (3) subjects update beliefs about themselves based on search outcomes, in the right direction but too conservatively compared to Bayes' rule; (4) this updating affects willingness to continue search; (5) some high ability types to search too little and some low ability types to search too long, due to wrong beliefs; (6) a substantial proportion

[17] In a companion paper, Falk, Huffman, and Sunde (2006), we extend the model to a general equilibrium framework, where firm behavior is explicitly accounted for. That model derives an endogenous wage schedule that is decreasing in unemployment duration. There we also show that the comparative statics results of the present model continue to hold in general equilibrium.

of individuals are averse to receiving information about ability if it is likely to be negative.

These findings provide a strong case for adopting imperfect self-knowledge as a more realistic behavioral assumption for search. They also identify systematic ways in which this uncertainty affects the search process. As discussed in the introduction, we believe there are many important implications for understanding job search behavior: an alternative explanation for negative duration dependence; a different interpretation of discouraged workers; new questions and challenges for active labor market policy; an additional factor explaining gender differences in search behavior and a gender wage gap; a channel through which unsuccessful search has a direct impact on utility, with potential consequences for willingness to search, and on the psychological well-being of the unemployed.

We illustrated how our findings can be incorporated into a partial equilibrium search model, in a parsimonious way. In a companion paper we also develop a general equilibrium model of the labor market in which searchers have imperfect self-knowledge (Falk, Huffman, and Sunde, 2006). We motivate the alternative behavioral assumptions of the model using the evidence from our experiment. We show that the model predicts the stylized facts discussed above, but also find that the model makes novel predictions regarding the dynamics of the labor market, for example greater volatility of unemployment in response to productivity shocks.

The mechanisms highlighted in our experiment are relevant for search more generally, e.g., search for a mate, or search for a journal that will publish a paper. In most types of search, individuals are probably uncertain of their own abilities, and uncertain of their abilities relative to others. As a result, the process of updating and loss of self-confidence described in our model is likely to be relevant for these search behaviors as well. For example, a young researcher trying unsuccessfully to publish papers in top tier journals may eventually lose confidence in his or her own abilities, and begin to aim lower or even switch to another career.

We tested our hypotheses using a laboratory experiment, because existing surveys typically do not contain the necessary information, just as survey do not contain information necessary to study human capital depreciation or stigma. In principle, however, it should be possible to design surveys that to elicit individuals' beliefs about their relative abilities and their job-finding chances, and their certainty about these beliefs. This would allow an investigation of how duration of unsuccessful search affects confidence and search decisions in the field. We believe that this is a fruitful direction for future research on search behavior.

References

ANDOLFATTO, D., S. MONGRAIN, AND G. MYERS (2005): "Self-Esteem and Lbour 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.

EL-GAMAL, M. A., AND D. M. GRETHER (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 KLAAUW (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.

GNEEZY, U., M. NIEDERLE, AND A. RUSTICHINI (2003): "Performance in Competitive Environments: Gender Differences," *Quarterly Journal of Economics*, 118(3), 1049–1074.

GRETHER, D. M. (1980): "Bayes Rule as a Descriptive Model: The Representativeness Heuristic," *The Quarterly Journal of Economics*, 95(3), 537–557.

——— (1992): "Testing Bayes' Rule and the Representative Heuristic: Some Experimental Evidence," *Journal of Economic Behavior and Organization*, 17, 31–57.

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.

KAHNEMAN, D., AND A. TVERSKY (1972): "Subjective Probability: A Judgement of Representativeness," *Cognitive Psychology*, 3, 430–454.

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.

Mailath, J., L. Samuelson, and A. Shaked (2000): "Endogenous Inequality in Integrated Labor Markets with Two-Sided Search," *American Economic Review*, 90(1), 46–72.

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 invest (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 the Self

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

	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 4: Is Updating Bayesian?

Dependent Variable: Log Posterior Odds of Being the High Type, End of Period t					
	All	Low types	High types	Information averse	Information seeking
	(1)	(2)	(3)	(4)	(5)
Ln(Posterior odds $t-1$) (β_1)	0.944***	0.886***	0.964***	0.950***	0.944***
	[0.039]	[0.077]	[0.052]	[0.036]	[0.056]
Ln(Likelihood ratio) (β_2)	0.347***	0.422*	0.278***	0.224*	0.387***
	[0.104]	[0.223]	[0.089]	[0.113]	[0.135]
Constant (α)	-0.006	-0.074	0.028	0.022	-0.019
	[0.082]	[0.113]	[0.117]	[0.056]	[0.122]
Test $\beta_1 = 1$	$p < 0.16$	$p < 0.15$	$p < 0.50$	$p < 0.19$	$p < 0.32$
Test $\beta_2 = 1$	$p < 0.001$	$p < 0.05$	$p < 0.001$	$p < 0.001$	$p < 0.001$
R^2	0.79	0.71	0.80	0.95	0.74
Observations	232	88	144	54	178

Notes: OLS regression coefficient estimates. The dependent variable is an individual's subjective posterior odds of being the high type, at the end of period t. In the odds ratio form of Bayes' Rule, the prior odds ratio reflects the impact of previous period beliefs, and the likelihood ratio reflects the impact of the search outcome in period t. Under the hypothesis that individuals use Bayes' Rule to update beliefs, β_1 and β_2 must be equal 1. Because beliefs are elicited in intervals, prior and posterior odds are constructed using midpoints of belief intervals. *E.g.*, 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. Robust standard errors are in brackets, adjusted for clustering at the individual level; ***, **, * indicate significance at 1-, 5-, and 10-percent level, respectively.

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.

Table 7: Impact of Beliefs About the Self and Risk Attitudes on Search Behavior

Dependent Variable: Search in Period *t*	(1)	(2)	(3)
Belief in *t-1*	1.139***	1.124***	1.010***
	[0.157]	[0.142]	[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 investments 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. Robust standard errors are in brackets, adjusted for clustering at the individual level; ***, **, * indicate significance at 1-, 5-, and 10-percent level, respectively.

Figures

Figure 1: Uncertainty About the Self

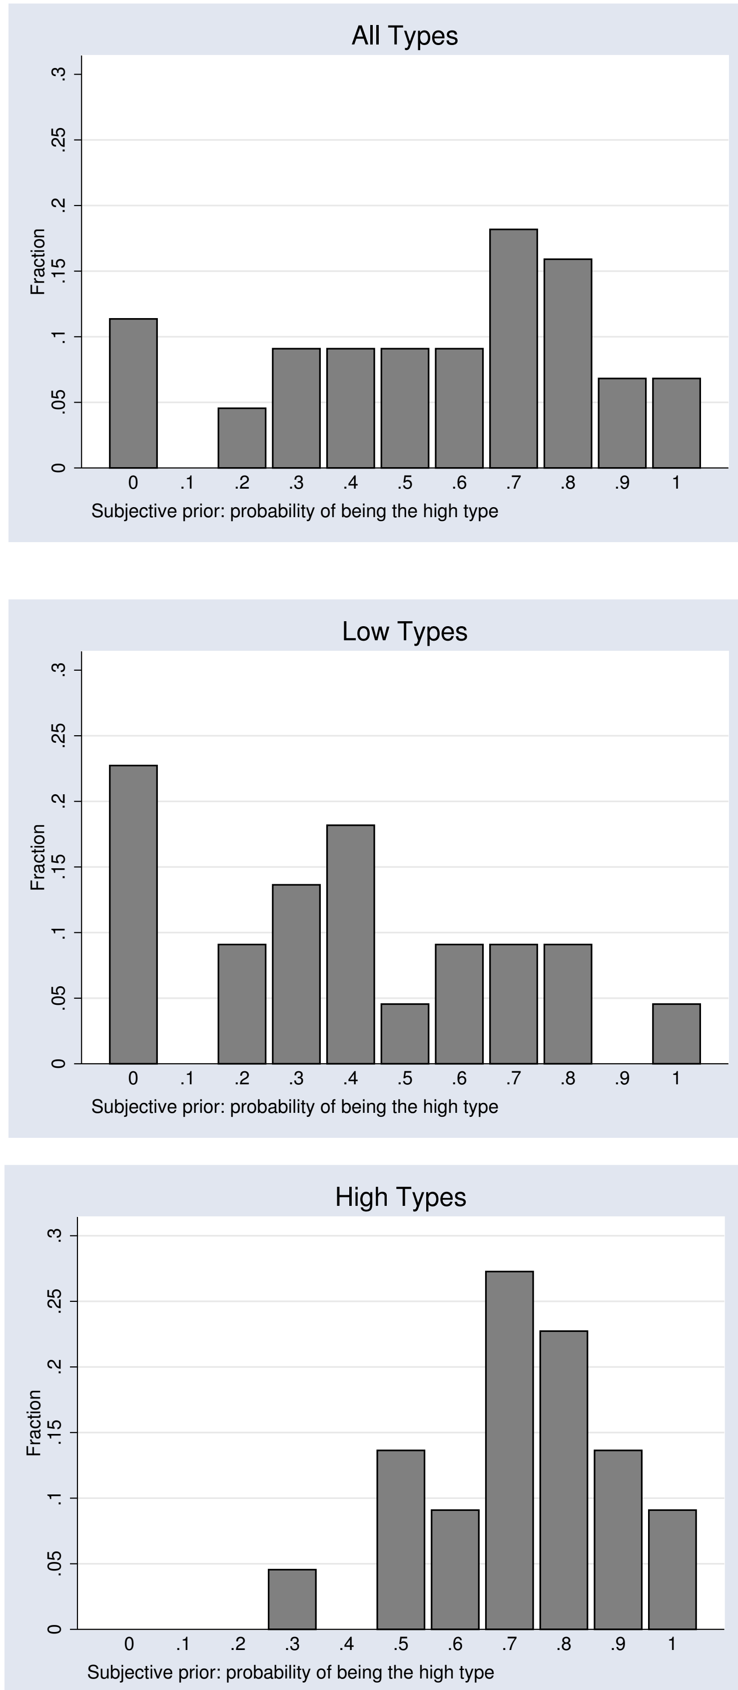

Notes: Subjects were assigned a high job finding probability in the search experiment (high type) if they scored higher than the median on an initial math test. After being informed of their own test score, but not the scores of others, subjects were asked: how likely do you think it is, in percentage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

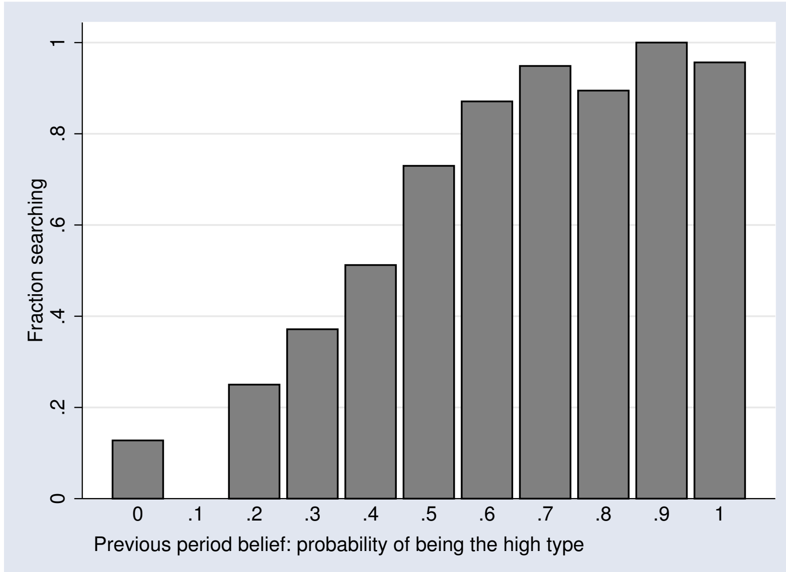

Notes: The figure shows the fraction of individuals searching, as a function of previous period belief about the probability of being a high type (having a high probability of success in search).