TUM School of Governance TUM School of Management

Munich Papers in Political Economy

Working Paper No. 7/2020

Motivated Beliefs and Anticipation of Uncertainty Resolution

Christoph Drobner

November 2020

Motivated Beliefs and Anticipation of Uncertainty Resolution*

Christoph Drobner[†]

May 28, 2021

Abstract

Manipulating subjects' expectations about the resolution of uncertainty, I show that subjects update beliefs about ego-relevant information optimistically when they expect no resolution of uncertainty but neutrally when they expect immediate uncertainty resolution. This finding highlights an important channel of the supply side of motivated beliefs and informs the discussion about the puzzling evidence on belief updating about ego-relevant information. Moreover, I document that subjects expost rationalize information by manipulating their stated beliefs about the ego-relevance of the underlying event depending on the valence of information. This result suggests an additional channel that subjects use to protect their ego utility.

Keywords: Motivated beliefs, Optimistic belief updating, Ex-post rationalization, Bayes' rule, Expectations about uncertainty resolution
JEL Codes: C91, D83, D84

*I thank Sebastian J. Goerg, Yves Le Yaouanq, Kai Barron, Alexander Coutts and seminar participants at the Bavarian Micro Day, TUM Research Seminar in Economics and Workshop on Behavioral and Experimental Economics (TUMCS) for helpful comments and suggestions. I gratefully acknowledge financial support from the ExperimenTUM.

[†]Technical University Munich, TUMCS Straubing, Am Essigberg 3, D-94315 Straubing, Germany; email: christoph.drobner@tum.de

1 Introduction

Empirical evidence shows that people's beliefs about personal characteristics or future life outcomes are often too optimistic (see Moore and Healy, 2008, for a review). This optimism bias can result in inferior decision-making and explains a variety of behavioral phenomena in the field such as suboptimal investment decisions (Malmendier and Geoffrey, 2005) and polarization in politics (Ortoleva and Snowberg, 2015). An important question in the economics literature is how these optimistic beliefs evolve despite the presence of objective information because standard economic theory assumes that people process information according to Bayes' rule. One behavioral explanation is that people update beliefs about ego-relevant information optimistically, overweighting good news relative to bad news with respect to their preferred state of the world (Eil and Rao, 2011; Möbius et al., 2014; Sharot et al., 2011). However, the experimental evidence on belief updating about ego-relevant information has produced very mixed results (Benjamin, 2019; Barron, 2020). For instance, studies by Ertac (2011) and Coutts (2019) even found empirical support for pessimistic belief updating. Conversely, studies in psychology and neuroscience document consistent evidence for optimistic belief updating (see Sharot and Garrett, 2016, for a review).

Taking this empirical puzzle as a starting point, I examine important differences in the experimental methodology between the disciplines of economics and psychology/neuroscience. My analysis reveals that there are systematic differences in the way that experimenters manage subjects' expectations regarding the short run and long-run resolution of uncertainty. In psychology/neuroscience, subjects form beliefs about future life events, such that the uncertainty will only be resolved in the distant future (if at all) and therefore remains uncertain during and after the experiment. In economics, subjects form beliefs about their relative performance in an experimental task such as relative performance in an IQ test. Subjects may therefore be more likely to expect immediate uncertainty resolution at the end of the experiment. Reviewing the instructions for the experiments in economics, I summarize that subjects either expect the resolution of uncertainty or experimenters have no control over subjects' expectations.

Building on existing theoretical models with belief-based utility, I postulate and experimentally show that these expectations play a key role to activate optimistic belief updating. Theoretical work in behavioral economics explains optimistic belief updating using interactions of preferences and beliefs (see Bénabou and Tirole, 2016, for a review). One implicit assumption of this strand of literature is that subjects derive not only instrumental but also direct utility from beliefs through motives such as anticipatory utility (Brunnermeier and Parker, 2005) or ego utility (Köszegi, 2006). In this paper, I hypothesize that subjects only derive direct utility from inflated beliefs when they expect no immediate uncertainty resolution. Intuitively, subjects may not be able to savor direct belief utility from inflated beliefs if they are immediately exposed to the potentially unpleasant truth. As a result, subjects may update their beliefs about good news and bad news optimistically when they expect no immediate uncertainty resolution and neutrally when the resolution of the true state is imminent.

To test this prediction, I use a controlled laboratory experiment and implement a variant of the belief updating task with relative performance in an IQ test as the underlying ego-relevant event. After the elicitation of prior beliefs about the likelihood to score in different ranks of the performance distribution in a group, I exogenously vary subjects' expectations about the resolution of uncertainty. Before subjects receive piece-wise information about their relative performance, they either get the information that their true rank remains uncertain or will be resolved by the end of the experiment. Subsequently, I provide subjects in both treatments with noisy and piece-wise information about their true rank and elicit posterior beliefs. Comparing subject's belief adjustments with optimal Bayesian belief adjustments confirms my hypothesis and shows that subjects only update their beliefs optimistically when their true rank remains uncertain. Conversely, subjects incorporate new information neutrally as the resolution of uncertainty is imminent.

Moreover, I exploit the noisy signal structure and identify that subjects also manipulate their beliefs about the ego-relevance of the IQ test depending on the valence of information. Using stated beliefs about the importance of the IQ test for study and job performance as a proxy for ego-relevance, I document that subjects perceive the IQ test as being more ego-relevant when they received good news compared to bad news about their relative performance. To this end, I show that subjects ex-post rationalize information depending on its valence, suggesting an additional channel that subjects use to protect their ego utility, which goes beyond biases in information processing. This behavioral mechanism, however, is again only present when subjects expect no resolution of uncertainty, supporting the hypothesis that expectations about the resolution of uncertainty play an important role to activate motivated reasoning.

Overall, this paper makes three contributions. First, I complement recent literature reviews on belief updating about ego-relevant information in economics and psychology/neuroscience. While Sharot and Garrett (2016) summarize the robust evidence for optimistic belief updating in experiments of psychology/neuroscience, Benjamin (2019) and Barron (2020) offer discussions about the mixed evidence in the experimental literature of economics and psychology/neuroscience. I amplify these reviews by highlighting that subjects may form heterogeneous expectations about the resolution of uncertainty.

Second, the results of my controlled laboratory experiment show the importance of these expectations to activate optimistic belief updating and ex-post rationalization of information. These findings provide a direct contribution to our understanding of the supply side of motivated beliefs and complement a literature, which identifies channels that switch optimistic belief updating on and off.[1] Given that most experiments in economics did not control subjects' expectations about the resolution of uncertainty, the findings in this paper cannot resolve the puzzling evidence on belief updating about ego-relevant information in economics. However, the results suggest that future research on motivated belief dynamics may devote particular attention to manage sub-

[1] For instance, Zimmermann (2020) shows that optimistic belief updating is only activated in the long run as subjects find ways to suppress negative feedback.

jects' expectations about the resolution of uncertainty. Moreover, the findings are relevant for theoretical work on motivated beliefs and indicate that deviations from Bayes' rule arise from an active choice to form self-serving beliefs rather than from an automatic updating process.

Third, the findings provide a potential explanation why people's outlooks often become more dire when the resolution of truth is imminent (see Sweeny and Krizan, 2013, for a review). For instance, Sweeny and Krizan (2013) show that voters lower their expectations about the chances of their preferred candidates as the election day approaches. In a related vein, Taylor and Shepperd (1998) document that subjects lower their predictions for a medical condition with severe consequences when they will learn the result of this testing in the near future. The focus of this literature is on the evolution of beliefs in waiting periods before the resolution of uncertainty but it remains agnostic about the process of how subjects incorporate piece-wise information into their beliefs. To this end, I complement this literature by showing that subject's expectations about the resolution of uncertainty result in differential belief updating patterns.

The remainder of this paper proceeds as follows: Section 2 provides an interdisciplinary literature review of experimental work on belief updating about ego-relevant information and derives the main hypothesis of this paper. Section 3 describes the experimental design. Section 4 discusses the results and Section 5 concludes.

2 Literature review and hypothesis

Taken all together, the evidence on preference-biased inference is confusing [...] Sorting out the reasons why different experiments reach different conclusions should be a priority.

Daniel J. Benjamin (2019)

A recent, growing body of experimental research has explored optimistic belief updating (i.e. preference-biased inference) about ego-relevant information and produced a variety of puzzling results. Table 1 summarizes these studies separated by the disciplines of economics (Panel A) and psychology/neuroscience (Panel B).[2]

The standard experimental paradigm in economics to study how people incorporate ego-relevant information into their existing beliefs is as follows: First, subjects solve an ego-relevant performance task such as an IQ test. Second, experimenters elicit subjects' probabilistic beliefs about their relative performance (e.g. the likelihood of scoring in the top half of the performance distribution). Third, subjects receive noisy but objective and informative feedback about their relative performance in form of a binary signal. This environment generates signal valence – information that contains either good news or bad news. Fourth, experimenters elicit subjects' posterior beliefs. This experimental methodology allows a direct comparison of subjects' belief updating process to the normative benchmark of Bayes' rule. Optimistic belief updating is identified when subjects update more strongly to good news than bad news about their relative performance compared to Bayesian updating.

Panel A in Table 1 summarizes the empirical evidence in economics. Studies by Eil and Rao (2011) and Möbius et al. (2014) document empirical support for optimistic belief updating. In contrast, Ertac (2011) and Coutts (2019) find evidence for pessimistic updating — individuals put more weight on bad news relative to good news about their preferred state of the world. Studies by Buser et al. (2018), Grossman and Owens (2012) and Schwardmann and van der Weele (2019) find no asymmetry in information processing. Zimmermann (2020) documents no asymmetry in the short run but optimistic belief updating in the long run, suggesting that this mechanism evolves over time as individuals find ways to suppress negative information. Benjamin (2019) and Barron (2020) offer discussions about potential drivers between the heterogeneity of the results within economics. To this end, Barron (2020) summarizes that experiments in Panel A of Table 1 differ in the size and ambiguity of

[2]The literature review in this section focuses on studies about belief updating with purely ego-relevant information and excludes similar work about valence-dependent belief updating with ego-neutral information such as Barron (2020), Gotthard-Real (2017), Lefebvre et al. (2017), Kuhnen (2015) and Palminteri et al. (2017).

Table 1: Experimental evidence on belief updating in ego-relevant domains

Panel A: Evidence from economics

Study	**Belief updating**	**Subjects expect resolution of uncertainty?**
Buser et al. (2018)	Neutral	Ambiguous
Coutts (2019)	Pessimistic	Yes
Eil and Rao (2011)	Optimistic	Ambiguous
Ertac (2011)	Pessimistic	Ambiguous
Grossman and Owens (2012)	Neutral	Ambiguous
Möbius et al. (2014)	Optimistic	Ambiguous
Schwardmann and v. d. Weele (2019)	Neutral	Ambiguous
Zimmermann (2020)	Neutral (short run) Optimistic (long run)	Ambiguous

Panel B: Evidence from psychology and neuroscience

Study	**Belief updating**	**Subjects expect resolution of uncertainty?**
Garrett and Sharot (2014)	Optimistic	No
Garrett et al. (2014)	Optimistic	No
Garrett and Sharot (2017)	Optimistic	No
Garrett et al. (2018)	Optimistic	No
Kuzmanovic et al. (2015)	Optimistic	No
Kuzmanovic et al. (2016)	Optimistic	No
Kuzmanovic and Rigoux (2017)	Optimistic	No
Korn et al. (2012)	Optimistic	No
Korn et al. (2014)	Optimistic	No
Marks and Baines (2017)	Optimistic	No
Moutsiana et al. (2013)	Optimistic	No
Moutsiana et al. (2015)	Optimistic	No
Shah et al. (2016)	Neutral	No
Sharot et al. (2011)	Optimistic	No
Sharot et al. (2012)	Optimistic	No

priors, information structures, domain of belief updating and stake sizes but none of these differences provides a neat explanation for the puzzling evidence.

Experimental studies on belief updating about ego-relevant information in psychology and neuroscience typically employ the following methodology: First, experimenters elicit subjects' probabilistic beliefs to experience future life events (e.g. the likelihood of getting a specific disease during their life time). Second, subjects receive the base rates of these events for individuals with the same socioeconomic background. This ensures signal valence because subjects either receive better than expected or worse than expected information based on their prior beliefs. Third, experimenters elicit subjects' posterior beliefs. Optimistic belief updating is identified, when subjects react more strongly to better than expected base rates in comparison to worse than expected base rates.

Panel B in Table 1 summarizes the empirical evidence in psychology and neuroscience. Except for Shah et al. (2016), there is consistent evidence for optimistic belief updating. This result is robust for positive and negative life events (Marks and Baines, 2017). Interestingly, Korn et al. (2014) show that optimistic belief updating disappears in a study with clinically depressed individuals and Garrett et al. (2018) observe diminished asymmetry when subjects update beliefs under immediately perceived threats. Furthermore, Moutsiana et al. (2013) document how information processing evolves over a lifespan, suggesting that updating on worse than expected information improves with age, while learning from good news does not change over the lifespan. Studies with fMRI tracking document that optimistic belief updating corresponds with a relative failure to encode negative information in frontal brain regions (Sharot et al., 2011; Sharot and Garrett, 2016).

Taken together, the literature in economics and psychology/neuroscience draws very different conclusions from the evidence about optimistic belief updating. Although there are various differences in the experimental methodologies between the disciplines, I point out one methodological difference, which may be an important driver of the heterogeneity in the results. Participants in studies of psychology/neuroscience form beliefs about future life events such that the true state of the world is unknown by experimenters. Conversely, participants in studies of economics form beliefs about their relative performance in an experimental task (mostly an IQ test) such that the true state of the world is known by experimenters.

This methodological difference has important implications for subjects' expectations about the resolution of the true state of the world during the course of the experiment, as summarized in column 3 of Table 1. Among studies in economics, only Coutts (2019) explicitly informed subjects that their true state would be resolved by the end of the experiment and finds pessimistic updating patterns. To the best of my knowledge, participants of the remaining studies of Panel A in Table 1 received no explicit information about the resolution of the true state such that experimenters have no control over subjects' expectations. Moreover, subjects may expect the inference of their true state from several rounds of feedback, accuracy payments for beliefs and performance payments in the IQ test. Conversely, experiments in psychology and neuroscience rule out the resolution of uncertainty by design as outcomes of future life events cannot be resolved during the course of the experiment and the absence of monetary incentives in belief elicitations reveals no additional information.

Building on existing theoretical work with belief-based utility, I hypothesize that subjects' expectations about the resolution of uncertainty play a key role to activate optimistic belief updating patterns. For a standard economic agent, there is no demand for optimistic belief updating because unbiased information processing is instrumentally useful as it leads subjects to make better decisions, i.e. maximize expected earnings in an experiment. Möbius et al. (2014) provide a behavioral model that rationalizes optimistic belief updating, suggesting that subjects do not only derive instrumental utility but also direct utility from beliefs through motives such as anticipatory utility (Brunnermeier and Parker, 2005) or ego utility (Köszegi, 2006). Consequently, optimistic belief updating results in optimistic beliefs with first-order gains through direct belief utility but second-order losses because biased beliefs can lead to sub-optimal decision making.[3]

In this paper, I hypothesize that optimistic belief updating is only activated when subjects expect no resolution of uncertainty in the near future. This hypothesis imposes the assumption that subjects cannot savor direct be-

[3]Note that other more instrumental reasons such as maintaining personal motivation (Bénabou and Tirole, 2002) or persuasive motives (Von Hippel and Trivers, 2011) may also be the driver of optimistic belief updating, which does not weaken the importance of expectations about temporal uncertainty resolution.

lief utility from optimistic beliefs as suggested by Brunnermeier and Parker (2005) when they are immediately exposed to the true state of the world. Consequently, subjects may update their beliefs neutrally upon the receipt of ego-relevant information when they expect immediate uncertainty resolution. Moreover, anticipation of immediate uncertainty resolution may even lead to pessimistic belief updating if utility is reference-dependent on previous beliefs and expectations, as suggested in Köszegi and Rabin (2006). This notion fits in a wider discussion of defensive pessimism, where subjects form pessimistic beliefs to avoid disappointment in the moment of resolution (see e.g. Norem and Cantor, 1986). For instance, some students form pessimistic expectations about grades in order to avoid psychological losses when the resolution occurs.

Hypothesis. *Subjects update beliefs about ego-relevant information optimistically when they expect no resolution of uncertainty and neutrally or even pessimistically when they expect immediate resolution of uncertainty.*

Taken together, this literature review reveals that subjects may form heterogeneous expectations about the resolution of uncertainty, which I propose to be an important channel to activate optimistic belief updating. However, the experiments in economics did not vary this design feature in a controlled environment and the experimental methodology in psychology/neuroscience differs in many aspects such that it is impossible to draw a conclusion based on the evidence in previous studies. The purpose of this paper is to close this gap in the literature and explicitely test this hypothesis in a controlled laboratory experiment.

3 Experimental design

To test this hypothesis in a laboratory experiment, the experimental design requires i) a belief updating task in which the underlying event contains ego-relevant information and ii) exogenous variation in subjects' expectations about the resolution of uncertainty. I capture these features by implementing the following experimental methodology: First, subjects conducted an IQ-related quiz. Second, I elicited subjects' prior beliefs about each possible rank in a randomly assigned group of four subjects. After the prior belief elicitation, I implemented the treatment variation. In a between-subject design, subjects were informed that their true rank would either be resolved (*Resolution* treatment) or remains uncertain (*No-Resolution* treatment) during the course of the experiment. Third, I provided subjects with piece-wise and binary feedback about their true rank and elicited posterior beliefs.

This experimental methodology allows me to compare subjects' belief adjustments to the corresponding Bayesian belief adjustments separately for good news and bad news depending on treatment status. Importantly, the treatment information was assigned after the prior belief elicitation to rule out that other prior related errors such as *base-rate neglect* or *confirmation bias* confound treatment differences in belief updating patterns. Moreover, the treatment variation hinges on the assumption that subjects do not expect to learn their true rank from payments. To this end, I aimed to obfuscate the relationship between payments and the true rank while maintaining the desirable properties of fully incentivized decisions. Subsequently, I provide a detailed outline about the different stages of the experiment.[4]

Quiz stage. To collect the necessary information to rank participants by their cognitive abilities, subjects performed an IQ-related quiz with puzzles of Civelli and Deck (2018) that are behaviorally similar to the commonly used Raven Progressive Matrices. Subjects saw a set of 15 puzzles and had 30 seconds each to choose the correct answer from a set of four possible answers. Subjects received a piece-rate payment that varied between 10 cents and 1 euro for each correct answer in the quiz. The size of the payment was randomly selected for each question to obfuscate the relationship between quiz payments and relative performance.

Prior belief elicitation. Subjects were randomly assigned to a group of

[4]Experimental instructions are provided in Appendix C.

four subjects. Subsequently, I elicited their probabilistic beliefs about the four possible ranks in the performance distribution of the IQ test within their group. To incentivize truthful reporting, I implemented a variation of the Becker-DeGroot-Marschak (BDM) mechanism (Becker et al., 1964) proposed by Grether (1981), Allen (1987) and Karni (2009). Participants were asked to state the likelihood p which makes them indifferent between winning a monetary prize of 2 Euro with likelihood p and winning the same monetary prize if they indeed performed in the specific rank given that this rank has been randomly chosen for payment. This mechanism ensures that truthful reporting maximizes expected utility regardless of subjects' risk preferences. Compared to alternative methods such as the quadratic scoring rule, the BDM mechanism obfuscates the relationship between belief payments and the true rank because it introduces an additional layer of probabilities and subjects never learned the randomly chosen rank for payments.

Feedback stage and posterior belief elicitation. In both treatments, subjects received one binary signal about their true rank. Subjects were randomly matched with another subject of their reference group and received complete information whether their score in the IQ test was higher or lower than the score of their matched group member (Eil and Rao, 2011). Subjects never learned the rank of the matched group member such that the signal contains noise and never reveals the true rank with certainty. Before subjects received the signal, I introduced the treatment variation: In the *No-Resolution* treatment, I informed subjects that their true rank would remain uncertain. In the *Resolution* treatment, I informed subjects that their true rank would be revealed by the end of the experiment. After the feedback stage, I elicited subjects' posterior beliefs about each possible rank in the group using the BDM mechanism. To emphasize the treatment information, I again informed subjects that they would either learn or not learn their true rank before they entered their posterior beliefs.

Questionnaire. In the post-experimental questionnaire, I elicited subjects' beliefs about the importance of the IQ test for their study and job performance on a seven-point Likert scale, which serves as a proxy for the perceived level of ego-relevance of the IQ test. Moreover, I elicited subjects' gender, age and major.

Procedures. The experiments were conducted at the ExperimenTUM of the Technical University Munich. I programmed the computerized experiments with the experimental software *ztree* (Fischbacher, 2007). Recruitment was automated, using the online recruitment software ORSEE by Greiner (2004) and the treatment was randomly assigned at the session level. A total of 200 subjects participated in the experiments: 100 subjects in the *Resolution* treatment and 100 subjects in the *No-Resolution* treatment. I ran 10 sessions of exactly 20 subjects each and the average duration of the experiments was 30 minutes. Subjects were paid in private and received only the combined payments of all parts in the experiment. Again, the purpose for doing so was to obfuscate the relationship between payments and the true rank in the group based on the performance in the IQ test as much as possible.[5]

4 Results

The experimental data include subjects' beliefs about the probability of each possible rank in the group, which consists of four people, before and after receiving the signal. Let $\hat{b}_{jt}$ be the probabilistic belief about each possible rank $j \in \{1, 2, 3, 4\}$ in periods $t \in \{0, 1\}$. Throughout the analysis, I refer to beliefs as the subject's expected rank in the group: $\hat{B}_t = \sum_{j=1}^{4} j \cdot \hat{b}_{jt}$. As a result, subjects' beliefs range in the interval [1,4] with the most confident subjects holding beliefs close to the lower boundary one and the least confident subjects holding beliefs close to the upper boundary four. In period 0, before subjects received the binary signal, I refer to subject's prior beliefs $\hat{B}_0$. In period 1, after subjects received the binary signal, I refer to subject's posterior beliefs

[5]Appendix B depicts the distribution of payments for each rank, showing that subjects could hardly infer their true rank from payments.

$\hat{B}_1$. Based on subject's probabilistic prior beliefs about each possible rank and the corresponding informativeness of the signal, I calculated the Bayesian posterior beliefs B_1, which I use as the normative benchmark in the analysis below. To test the optimistic belief updating hypothesis, I follow Eil and Rao (2011), Grossman and Owens (2012) and Zimmermann (2020) and investigate asymmetry in subjects' belief adjustments after good news and bad news in comparison to the normative benchmark of Bayes' rule. Belief adjustments are defined as the subject's posterior beliefs $\hat{B}_1$ minus prior beliefs $\hat{B}_0$, while Bayesian belief adjustments are defined as the subject's Bayesian posterior beliefs B_1 minus subjective prior beliefs $\hat{B}_0$.

4.1 Aggregate beliefs

Before I delve into the analysis of belief adjustments after good news and bad news, I describe beliefs at the aggregate level. The treatment was assigned after the prior belief elicitation such that I expected no level difference in aggregate prior beliefs between treatments. On average, prior beliefs in the *No-Resolution* treatment are 2.382 compared to 2.418 in the *Resolution* treatment. This small difference in prior beliefs is in fact not statistically different from zero (two-sided t-test, $p = 0.637$). Pooling data from both treatments, aggregate prior beliefs are significantly below 2.5, which is the rational belief at the aggregate level (two-sided t-test, $p = 0.009$). Considering that beliefs are more optimistic if they converge towards rank 1, this result provides evidence for overconfidence in subjects' aggregate prior beliefs.

After subjects received the binary signal, I compare subjects' posterior beliefs to the corresponding Bayesian posterior beliefs. In the *Resolution* treatment, average posteriors are 2.443 and almost equal to the corresponding Bayesian posteriors of 2.452. To this end, the resulting difference is not statistically different from zero (two-sided t-test, $p = 0.837$). In the *No-Resolution* treatment, average posteriors are 2.376 and tend to be more optimistic than the corresponding Bayesian posteriors of 2.423. However, the resulting difference is not statistically different from zero (two-sided t-test, $p = 0.215$).

Result 1 *Subjects' prior beliefs are overconfident. Subjects' posterior beliefs are not significantly different from Bayesian posteriors.*

4.2 Belief adjustments

Looking only at the aggregate data may potentially lead to wrong conclusions about updating behavior because subjects only adjust their beliefs upon receipt of one signal and this may not result in large biases at the aggregate level. To this end, I follow Eil and Rao (2011), Grossman and Owens (2012) and Zimmermann (2020) and analyze how closely subjects' belief adjustments follow Bayesian belief adjustments separately for good news and bad news. Figure 1 plots subjects' belief adjustments with corresponding Bayesian belief adjustments separately for *No-Resolution* and *Resolution* treatments. The green crosses and red dots depict belief adjustments after good news and bad news, respectively. The stacked 45-degree lines represent the normative benchmark with an overlap of subjects' belief adjustments and Bayesian belief adjustments. If subjects update their beliefs in the direction of Bayesian belief adjustments, belief adjustments after good news are negative with posterior beliefs closer to rank 1 and belief adjustments after bad news are positive with posterior beliefs closer to rank 4.

Overall, the patterns in Figure 1 show conservatism in both treatments because the majority of belief adjustments is lower in absolute magnitude compared to Bayesian belief adjustments for both good news and bad news. More importantly, however, the fitted values in the *No-Resolution* treatment show that belief adjustments after good news have a significantly steeper slope and follow the Bayesian predictions more closely than belief adjustments after bad news. This pattern provides visual evidence for optimistic belief updating in the *No-Resolution* treatment. Conversely, no such asymmetry is observed in the *Resolution* treatment where belief adjustments follow the Bayesian predictions similarly for both good news and bad news.[6]

[6]Figure 2 in Appendix A.1 shows similar belief updating patterns as in Figure 1 when I exclude subjects with belief adjustments in the wrong direction.

Figure 1: Belief adjustments on Bayesian belief adjustments

In Table 2, I provide quantitative evidence of the patterns illustrated in Figure 1, regressing subjects' belief adjustments on the corresponding Bayesian belief adjustments separately for good news and bad news including a diff-in-diff specification to identify optimistic belief updating (Eil and Rao, 2011; Grossman and Owens, 2012; Zimmermann, 2020). The estimated coefficients in both treatments for Bayesian belief adjustments β_1 are significantly below one for both good news and bad news, confirming that subjects update their beliefs conservatively. The regression output of the *No-Resolution* treatment (columns 1 and 2) reveals that the slope coefficient for Bayesian belief adjustments is substantially higher for good news ($\beta_1 = 0.665$) than bad news ($\beta_1 = 0.076$), showing that subjects follow the Bayesian prediction much more closely for good news than bad news. The diff-in-diff specification in column 3 confirms this result with a positive and highly significant interaction term

β_3 ($p = 0.004$). This finding provides evidence for optimistic belief updating in the *No-Resolution* treatment and indicates that subjects' inference from bad news does not respect signal strength. In contrast, the slope coefficients for Bayesian belief adjustments in the *Resolution* treatment (columns 4 and 5) show that belief adjustments follow the Bayesian predictions similarly for good news ($\beta_1 = 0.530$) and bad news ($\beta_1 = 0.645$). To this end, the coefficient of the interaction term β_3 in the diff-in-diff specification in column 6 is slightly negative and insignificant ($p = 0.729$), suggesting no asymmetry in subjects' belief updating process.

Table 2: Belief adjustments on Bayesian belief adjustments

$$Beliefadjustment_i = \beta_0 + \beta_1 Bayesbeliefadj_i + \beta_2 Good\ news_i + \beta_3 Bayesbeliefadj_i * Good\ news_i + \epsilon_i$$

	No-Resolution			Resolution		
	Good news (1)	Bad news (2)	Diff-in-diff (3)	Good news (4)	Bad news (5)	Diff-in-diff (6)
β_1	0.665	0.076	0.076	0.530	0.645	0.645
	(0.088)	(0.180)	(0.180)	(0.218)	(0.249)	(0.249)
β_2			-0.359			-0.133
			(0.129)			(0.166)
β_3			0.589			-0.115
			(0.200)			(0.331)
Constant	-0.042	0.317	0.317	-0.094	0.039	0.039
	(0.053)	(0.118)	(0.118)	(0.105)	(0.128)	(0.128)
Observations	50	50	100	50	50	100
R^2	0.419	0.003	0.646	0.124	0.137	0.531

Notes:
(i) Subjects' belief adjustments are defined as subjects' posteriors minus priors. Bayesian belief adjustments are defined as Bayesian posteriors minus subjects' priors.
(ii) Analysis uses OLS regressions with robust standard errors in parentheses.

Result 2 *Subjects' belief adjustments follow Bayesian belief adjustments more closely for good news than bad news when they expect no resolution of uncertainty. Conversely, subjects' belief adjustments follow Bayesian belief adjustments similarly for good news and bad news when they expect immediate uncertainty resolution.*

One alternative interpretation of the treatment difference in updating patterns could be that subjects revise their prior beliefs immediately after receiving the treatment information, rather than process the signal differently.

For example, subjects in the *No-Resolution* treatment might become more optimistic after receiving the information that they will not learn their true rank. Before addressing this potential confound, I would like to emphasize that both a revision of prior beliefs and differences in updating represent motivated departures from the Bayesian model, such the results clearly show that expectations about uncertainty resolution influence the formation of motivated beliefs.

Still, these two channels are conceptually different and it is interesting to disentangle the main driver behind the results. To this end, I suggest that the coefficient estimates in Table 2 provide evidence that a shift in priors may not be the leading explanation for the results. Specifically, a signal independent revision of priors would suggest that we see treatment differences in updating patterns after both good news and bad news.[7] However, a treatment comparison of the coefficient estimates for Bayesian belief adjustments β_1 in column 1 and column 4 of Table 2 reveals that subjects follow the Bayesian prediction similarly after the receipt of good news. This is confirmed by a Chow test, showing that the estimated coefficients are not significantly different between treatments ($\beta_1^{Resolution} = \beta_1^{No-Resolution}, p = 0.565$). In contrast, a treatment comparison of the coefficient estimates for Bayesian belief adjustments β_1 after the receipt of bad news in column 2 and column 5 of Table 2 shows that the estimated coefficients are substantially stronger in magnitude in the *Resolution* treatment compared to the *No-Resolution* treatment. This difference in estimated coefficients is marginally significant using a Chow test ($\beta_1^{Resolution} = \beta_1^{No-Resolution}, p = 0.067$). As a result, I conclude that the treatment difference in updating patterns is mostly driven by subjects who received

[7] To illustrate this point, suppose that subjects in the *No-Resolution* treatment become more optimistic, which shifts their prior distribution g over the ranks $\{1, 2, 3, 4\}$ up to f such that f dominates g in the monotone likelihood ratio property (MLRP), i.e. $f(x)g(x+1) > f(x+1)g(x)$ for every rank $x \in \{1, 2, 3, 4\}$. Let p_x be the probability of receiving good news if the rank is x. Writing out Bayes' rule, this means that for every rank $x \in \{1, 2, 3, 4\}$, $\frac{p_x f(x)}{\sum_{x=1}^4 p_x f(x)} \frac{p_{x+1} g(x+1)}{\sum_{x=1}^4 p_x g(x)} > \frac{p_{x+1} f(x+1)}{\sum_{x=1}^4 p_x f(x)} \frac{p_x g(x)}{\sum_{x=1}^4 p_x g(x)} \Leftrightarrow f(x)g(x+1) > f(x+1)g(x)$, which implies that the MLRP is preserved by Bayesian updating regardless of whether subjects receive good news or bad news. As a result, we should observe treatment differences in updating patterns regardless of whether subjects receive good news or bad news if the shift in priors satisfies the MLRP.

bad news such that a signal independent shift in priors may not be the main driver of the results.

To show the robustness of belief updating patterns, Appendix A replicates the regression analysis in Table 2 using different sample selection criteria and controlling for potentially confounding variables. First, Appendix A.1 replicates the regression analysis excluding subjects with belief adjustments in the wrong direction. Second, Appendix A.2 replicates the regression analysis using a restricted sample excluding subjects with wrong and zero belief adjustments. Third, one potential endogeneity issue of the presented analysis and the optimistic belief updating literature in general arises from the fact that subjects who receive good news may systematically differ from subjects who receive bad news (see Barron, 2020, for a discussion). To address this endogeneity concern, Appendices A.3–A.4 provide additional robustness checks by replicating the regression analysis in Table 2, controlling for subjects' ranks in the group and IQ test scores. Moreover, Appendix A.5 replicates the regression analysis using a restricted sample excluding subjects who are ranked first or fourth in their reference group and therefore only received good news or bad news, respectively. Overall, the robustness checks in Appendix A confirm the main results in Table 2 with both qualitatively and quantitatively similar coefficient estimates.

4.3 Ex-post rationalization

To summarize, thus far I have documented that subjects update their beliefs about ego-relevant information optimistically if they expect no immediate resolution of uncertainty but neutrally if they expect immediate uncertainty resolution. In this section, I examine whether subjects ex-post rationalize information by manipulating their beliefs about the ego-relevance of the IQ test depending on the valence of information. To this end, I use subjects' answers in the post-experimental survey about the importance of the IQ test for their study and job performance as a proxy for subjects' beliefs about the ego-relevance of the IQ test and investigate the following question: Do subjects perceive the IQ test as being more ego-relevant if they incidentally received good news instead of bad news about their relative performance? To answer this question, I regress subjects' stated importance of the IQ test for study and job performance measured on a seven-point Likert scale on a dummy for good news. Controlling for IQ test scores and prior beliefs, I exploit the noisy signal structure to exogenously estimate the effect of good news received on subjects' stated ego-relevance.

Table 3 shows the corresponding regression analysis separately for *No-Resolution* and *Resolution* treatments using ordered logistic regressions. The results in columns 1 and 2 show that subjects in the *No-Resolution* treatment in fact state substantially higher beliefs about the importance of the IQ test for study performance ($p = 0.027$) and job performance ($p = 0.010$) when they incidentally received good news instead of bad news about their relative performance. Intuitively, subjects discount the ego-relevance of the underlying event when they received bad feedback about their relative performance. This behavioral mechanism offers an additional channel for subjects to protect their ego utility even though they cannot fully explain away negative feedback through biases in information processing. Looking at the *Resolution* treatment in column 3 and 4 shows that the coefficients for good news are small and not statistically significantly different from zero, suggesting that subjects do not engage in ex-post rationalization of information when they expect immediate uncertainty resolution. To show the robustness of this finding, Appendix A.6 replicates the regression analysis in Table 3 excluding subjects with belief adjustments in the wrong direction with both qualitatively and quantitatively similar coefficient estimates.

Result 3 *Subjects ex-post rationalize information about their relative performance in the IQ test when they expect no resolution of uncertainty.*

Putting this result together with the evidence in Section 4.2, subjects in the *No-Resolution* treatment engage in optimistic belief updating and ex-post rationalization of ego-relevant information, while subjects in the *Resolution* treatment update their beliefs neutrally and evaluate the importance of the

Table 3: Ex-post rationalization of information

	No-Resolution		Resolution	
Dependent variable	Importance study performance (1)	Importance job performance (2)	Importance study performance (3)	Importance job performance (4)
Good news	0.992	1.168	0.122	0.317
	(0.450)	(0.455)	(0.434)	(0.423)
IQ test score	-0.009	-0.077	0.036	-0.059
	(0.080)	(0.082)	(0.078)	(0.078)
Prior belief	-0.729	-1.190	-0.916	-0.931
	(0.416)	(0.419)	(0.345)	(0.341)
Observations	100	100	100	100
Pseudo R^2	0.043	0.060	0.029	0.023

Notes:
(i) Subjects' stated importance of the IQ test for study and job performance is measured on a seven-point Likert scale.
(ii) Analysis uses Ordered Logistic Regressions with standard errors in parentheses.

IQ test independent of signal valence. One behavioral explanation for these results is that subjects cannot derive direct utility from motivated beliefs when they expect the immediate resolution of uncertainty. To this end, I conclude that motivated reasoning is not activated when subjects expect the immediate resolution of the true state of the world.

5 Concluding remarks

Optimistic belief updating about ego-relevant information can explain why people often end up with too optimistic beliefs about their abilities and future prospects. The results of my controlled laboratory experiment show that this behavioral mechanism is bounded by subjects' expectations about the resolution of uncertainty. While subjects update beliefs about ego-relevant information optimistically when the resolution of uncertainty is absent, they process information neutrally when they expect the resolution of the true state of the world. This result highlights an important dimension of the supply side of motivated beliefs, which is relevant for future empirical and theoretical work.

Moreover, I document that expectations about the resolution of uncertainty also play a key role to activate ex-post rationalization of information

– a potentially important behavioral mechanism for subjects to protect their ego utility, which has been previously undetected in the literature about motivated beliefs. One caveat of this result is that I focused on stated beliefs such that belief distortions about the ego-relevance of the underlying event are not costly in terms of expected earnings in the experiment. To this end, future research may departure from this finding and investigate the impact of this behavioral mechanism in environments with economic consequences.

Overall, one limitation of my analysis is that I focused on the two boundary cases of immediate resolution and no-resolution of uncertainty. As a result, future research is needed to understand whether optimistic belief updating and ex-post rationalization of information is activated when the uncertainty is resolved but the moment of resolution lies in the future. This environment represents a variety of dynamic economic decisions such as human capital formation in which beliefs about returns on human capital investments are formed in the present and the resolution of returns arrives at some point in the future.

References

Allen, F. (1987): "Discovering personal probabilities when utility functions are unknown," *Management Science*, 33, 542–544.

Barron, K. (2020): "Belief Updating: Does the 'Good-news, Bad-news' Asymmetry Extend to Purely Financial Domains?" *Experimental Economics*, 1–28.

Becker, G. M., M. H. DeGroot, and J. Marschak (1964): "Measuring utility by a single-response sequential method," *Behavioral Science*, 9, 226–232.

Bénabou, R. and J. Tirole (2002): "Self-Confidence and Personal Motivation," *The Quarterly Journal of Economics*, 117, 871–915.

——— (2016): "Mindful Economics: The Production, Consumption, and Value of Beliefs," *Journal of Economic Perspectives*, 30, 141–164.

Benjamin, D. J. (2019): "Errors in probabilistic reasoning and judgment biases," in *Handbook of Behavioral Economics - Foundations and Applications 2*, Elsevier, vol. 2, 69–186.

Brunnermeier, M. K. and J. A. Parker (2005): "Optimal expectations," *American Economic Review*, 95, 1092–1118.

Buser, T., L. Gerhards, and J. van der Weele (2018): "Responsiveness to feedback as a personal trait," *Journal of Risk and Uncertainty*, 56, 165–192.

Civelli, A. and C. Deck (2018): "A Flexible and Customizable Method for Assessing Cognitive Abilities," *Review of Behavioral Economics*, 5, 123–147.

Coutts, A. (2019): "Good News and Bad News are Still News: Experimental Evidence on Belief Updating," *Experimental Economics*, 109, 21.

Eil, D. and J. M. Rao (2011): "The Good News-Bad News Effect: Asymmetric Processing of Objective Information about Yourself," *American Economic Journal: Microeconomics*, 3, 114–138.

Ertac, S. (2011): "Does self-relevance affect information processing? Experimental evidence on the response to performance and non-performance feedback," *Journal of Economic Behavior & Organization*, 80, 532–545.

Fischbacher, U. (2007): “z-Tree: Zurich toolbox for ready-made economic experiments,” *Experimental Economics*, 10, 171–178.

Garrett, N., A. M. González-Garzón, L. Foulkes, L. Levita, and T. Sharot (2018): “Updating Beliefs under Perceived Threat,” *The Journal of Neuroscience*, 38, 7901–7911.

Garrett, N. and T. Sharot (2014): “How robust is the optimistic update bias for estimating self-risk and population base rates?” *PLoS ONE*, 9, 1–8.

——— (2017): “Optimistic update bias holds firm: Three tests of robustness following Shah et al,” *Consciousness and Cognition*, 50, 12–22.

Garrett, N., T. Sharot, P. Faulkner, C. W. Korn, J. P. Roiser, and R. J. Dolan (2014): “Losing the rose tinted glasses: neural substrates of unbiased belief updating in depression,” *Frontiers in Human Neuroscience*, 8, 639.

Gotthard-Real, A. (2017): “Desirability and information processing: An experimental study,” *Economics Letters*, 152, 96–99.

Greiner, B. (2004): “An Online Recruitment System for Economic Experiments.” *University of Cologne Working Paper Series in Economics*, 10.

Grether, D. M. (1981): “Financial incentive effects and individual decision-making,” *Working paper*.

Grossman, Z. and D. Owens (2012): “An unlucky feeling: Overconfidence and noisy feedback,” *Journal of Economic Behavior & Organization*, 84, 510–524.

Karni, E. (2009): “A Mechanism for Eliciting Probabilities,” *Econometrica*, 77, 603–606.

Korn, C. W., K. Prehn, S. Q. Park, H. Walter, and H. R. Heekeren (2012): “Positively biased processing of self-relevant social feedback,” *The Journal of Neuroscience*, 32, 16832–16844.

Korn, C. W., T. Sharot, H. Walter, H. R. Heekeren, and R. J. Dolan (2014): “Depression is related to an absence of optimistically biased belief updating about future life events,” *Psychological Medicine*, 44, 579–592.

Köszegi, B. (2006): “Ego Utility, Overconfidence, and Task Choice,” *Journal* *of the European Economic Association*, 4, 673–707.

Kőszegi, B. and M. Rabin (2006): "A Model of Reference-Dependent Preferences," *The Quarterly Journal of Economics*, 121, 1133–1165.

Kuhnen, C. M. (2015): "Asymmetric Learning from Financial Information," *The Journal of Finance*, 70, 2029–2062.

Kuzmanovic, B., A. Jefferson, and K. Vogeley (2015): "Self-specific Optimism Bias in Belief Updating Is Associated with High Trait Optimism," *Journal of Behavioral Decision Making*, 28, 281–293.

——— (2016): "The role of the neural reward circuitry in self-referential optimistic belief updates," *NeuroImage*, 133, 151–162.

Kuzmanovic, B. and L. Rigoux (2017): "Valence-Dependent Belief Updating: Computational Validation," *Frontiers in Psychology*, 8, 1–11.

Lefebvre, G., M. Lebreton, F. Meyniel, S. Bourgeois-Gironde, and S. Palminteri (2017): "Behavioural and neural characterization of optimistic reinforcement learning," *Nature Human Behaviour*, 1, 1–9.

Malmendier, U. and T. Geoffrey (2005): "CEO overconfidence and corporate investment," *The journal of financ*, 60, 2661–2700.

Marks, J. and S. Baines (2017): "Optimistic belief updating despite inclusion of positive events," *Learning and Motivation*, 58, 88–101.

Möbius, M. M., M. Niederle, P. Niehaus, and T. S. Rosenblat (2014): "Managing Self-Confidence," *Working paper*.

Moore, D. A. and P. J. Healy (2008): "The trouble with overconfidence," *Psychological Review*, 115, 502–517.

Moutsiana, C., C. J. Charpentier, N. Garrett, M. X. Cohen, and T. Sharot (2015): "Human Frontal-Subcortical Circuit and Asymmetric Belief Updating," *The Journal of Neuroscience*, 35, 14077–14085.

Moutsiana, C., N. Garrett, R. C. Clarke, R. B. Lotto, S.-J. Blakemore, and T. Sharot (2013): "Human development of the ability to learn from bad news," *Proceedings of the National Academy of Sciences of the United States of America*, 110, 16396–16401.

Norem, J. K. and N. Cantor (1986): "Anticipatory and post hoc cushioning strategies: Optimism and defensive pessimism in risky situations," *Cognitive Therapy and Research*, 10, 347–362.

Ortoleva, P. and E. Snowberg (2015): "Overconfidence in political behavior," *American Economic Review*, 105, 504–35.

Palminteri, S., G. Lefebvre, E. J. Kilford, and S.-J. Blakemore (2017): "Confirmation bias in human reinforcement learning: Evidence from counterfactual feedback processing," *PLoS computational biology*, 13.

Schwardmann, P. and J. van der Weele (2019): "Deception and Self-Deception," *Nature Human Behaviour*, 1–7.

Shah, P., A. J. L. Harris, G. Bird, C. Catmur, and U. Hahn (2016): "A pessimistic view of optimistic belief updating," *Cognitive Psychology*, 90, 71–127.

Sharot, T. and N. Garrett (2016): "Forming Beliefs: Why Valence Matters," *Trends in Cognitive Sciences*, 20, 25–33.

Sharot, T., M. Guitart-Masip, C. W. Korn, R. Chowdhury, and R. J. Dolan (2012): "How dopamine enhances an optimism bias in humans," *Current Biology*, 22, 1477–1481.

Sharot, T., C. W. Korn, and R. J. Dolan (2011): "How unrealistic optimism is maintained in the face of reality," *Nature Neuroscience*, 14, 1475–1479.

Sweeny, K. and Z. Krizan (2013): "Sobering up: a quantitative review of temporal declines in expectations," *Psychological bulletin*, 139, 702–724.

Taylor, K. M. and J. A. Shepperd (1998): "Bracing for the worst: Severity, testing, and feedback timing as moderators of the optimistic bias," *Personality and Social Psychology Bulletin*, 24, 915–926.

Von Hippel, W. and R. Trivers (2011): "The evolution and psychology of self-deception." *Behavioral and brain sciences*, 34, 1–56.

Zimmermann, F. (2020): "The dynamics of motivated beliefs," *American Economic Review*, 110, 337–361.

Appendices

A Robustness checks

Subsequently, I provide robustness checks for the results in Section 4.2 and Section 4.3.

A.1 Belief adjustments - excluding wrong belief adjustments

Figure 2 plots subjects' belief adjustments on Bayesian belief adjustments, excluding belief adjustments in the wrong direction.

Figure 2: Belief adjustments on Bayesian belief adjustments

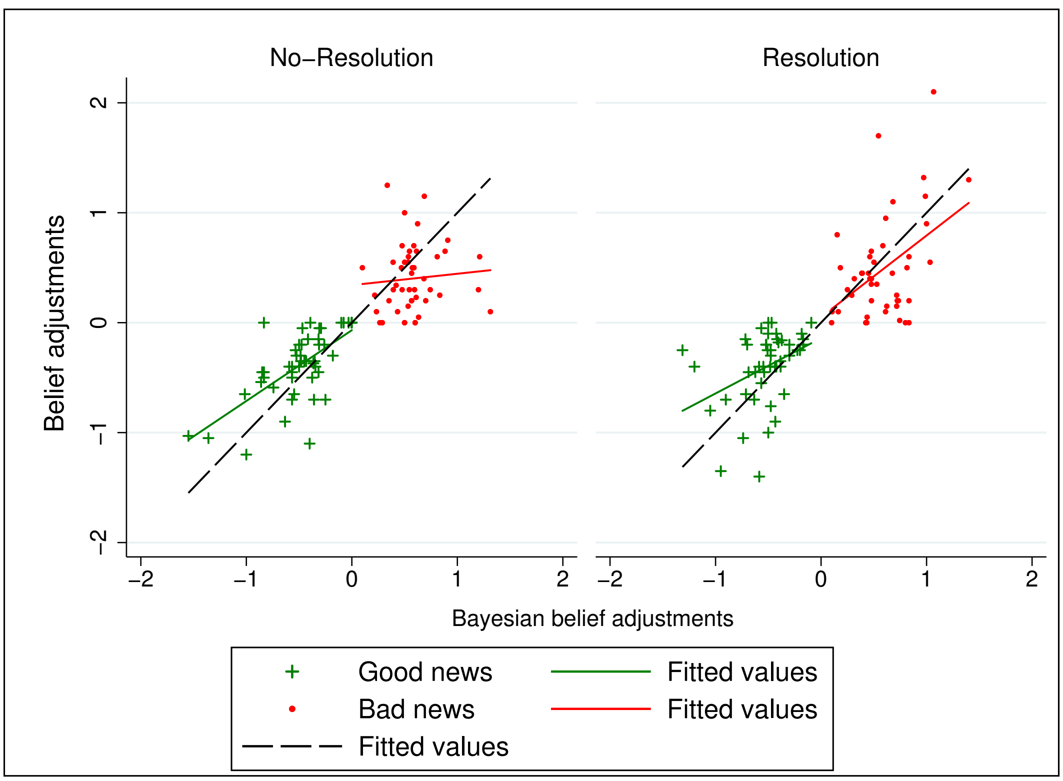

In Table 4, I replicate the regression analysis in Table 2 of Section 4.2, excluding belief adjustments in the wrong direction.

Table 4: Belief adjustments - excluding wrong belief adjustments

$$Beliefadjustment_i = \beta_0 + \beta_1 Bayesbeliefadj_i + \beta_2 Good\ news_i + \beta_3 Bayesbeliefadj_i * Good\ news_i + \epsilon_i$$

	No-Resolution			Resolution		
	Good news (1)	Bad news (2)	Diff-in-diff (3)	Good news (4)	Bad news (5)	Diff-in-diff (6)
β_1	0.644	0.104	0.104	0.504	0.742	0.742
	(0.089)	(0.179)	(0.178)	(0.219)	(0.239)	(0.239)
β_2			-0.410			-0.191
			(0.129)			(0.161)
β_3			0.540			-0.239
			(0.200)			(0.324)
Constant	-0.068	0.341	0.341	-0.140	0.051	0.051
	(0.053)	(0.118)	(0.118)	(0.102)	(0.125)	(0.125)
Observations	48	47	95	48	45	93
R^2	0.426	0.007	0.711	0.146	0.200	0.637

Notes:
(i) Subjects' belief adjustments are defined as subjects' posteriors minus priors. Bayesian belief adjustments are defined as Bayesian posteriors minus subjects' priors.
(ii) Analysis uses OLS regressions with robust standard errors in parentheses.

A.2 Belief adjustments - excluding wrong and zero belief adjustments

In Table 5, I replicate the regression analysis in Table 2 of Section 4.2, excluding belief adjustments in the wrong direction and zero belief adjustments.

Table 5: Belief adjustments - excluding wrong and zero belief adjustments

$$Beliefadjustment_i = \beta_0 + \beta_1 Bayesbeliefadj_i + \beta_2 Good\ news_i + \beta_3 Bayesbeliefadj_i * Good\ news_i + \epsilon_i$$

	No-Resolution			Resolution		
	Good news (1)	Bad news (2)	Diff-in-diff (3)	Good news (4)	Bad news (5)	Diff-in-diff (6)
β_1	0.625	-0.023	-0.023	0.459	0.775	0.775
	(0.095)	(0.176)	(0.176)	(0.226)	(0.240)	(0.240)
β_2			-0.584			-0.271
			(0.137)			(0.168)
β_3			0.648			-0.316
			(0.200)			(0.330)
Constant	-0.110	0.474	0.474	-0.185	0.085	0.085
	(0.067)	(0.120)	(0.120)	(0.107)	(0.129)	(0.129)
Observations	42	41	83	45	40	85
R^2	0.411	0.000	0.780	0.126	0.225	0.686

Notes:
(i) Subjects' belief adjustments are defined as subjects' posteriors minus priors. Bayesian belief adjustments are defined as Bayesian posteriors minus subjects' priors.
(ii) Analysis uses OLS regressions with robust standard errors in parentheses.

A.3 Belief adjustments - controlling for ranks

In Table 6, I replicate the regression analysis in Table 2 of Section 4.2, controlling for subjects' ranks in the group.

Table 6: Belief adjustments - controlling for ranks

$$Beliefadjustment_i = \beta_0 + \beta_1 Bayesbeliefadj_i + \beta_2 Good\ news_i + \beta_3 Bayesbeliefadj_i * Good\ news_i + \epsilon_i$$

	No-Resolution			Resolution		
	Good news (1)	Bad news (2)	Diff-in-diff (3)	Good news (4)	Bad news (5)	Diff-in-diff (6)
β_1	0.657	0.093	0.082	0.529	0.666	0.654
	(0.098)	(0.182)	(0.181)	(0.224)	(0.242)	(0.245)
β_2			-0.322			-0.038
			(0.171)			(0.190)
β_3			0.593			-0.124
			(0.199)			(0.324)
Rank	✓	✓	✓	✓	✓	✓
Constant	-0.022	0.152	0.259	-0.068	-0.387	-0.147
	(0.065)	(0.267)	(0.195)	(0.176)	(0.316)	(0.234)
Observations	50	50	100	50	50	100
R^2	0.420	0.012	0.647	0.125	0.175	0.536

Notes:
(i) Subjects' belief adjustments are defined as subjects' posteriors minus priors. Bayesian belief adjustments are defined as Bayesian posteriors minus subjects' priors.
(ii) Analysis uses OLS regressions with robust standard errors in parentheses.

A.4 Belief adjustments - controlling for IQ test scores

In Table 7, I replicate the regression analysis in Table 2 of Section 4.2, controlling for subjects' IQ test scores.

Table 7: Belief adjustments - controlling for IQ test scores

$$Beliefadjustment_i = \beta_0 + \beta_1 Bayesbeliefadj_i + \beta_2 Good\ news_i + \beta_3 Bayesbeliefadj_i * Good\ news_i + \epsilon_i$$

	No-Resolution			Resolution		
	Good news (1)	Bad news (2)	Diff-in-diff (3)	Good news (4)	Bad news (5)	Diff-in-diff (6)
β_1	0.574	0.031	0.040	0.525	0.647	0.638
	(0.123)	(0.191)	(0.185)	(0.225)	(0.269)	(0.264)
β_2			-0.568			-0.153
			(0.171)			(0.197)
β_3			0.470			-0.110
			(0.197)			(0.344)
Quiz score	✓	✓	✓	✓	✓	✓
Constant	-0.320	-0.010	0.055	-0.306	0.052	-0.002
	(0.261)	(0.172)	(0.143)	(0.288)	(0.218)	(0.173)
Observations	50	50	100	50	50	100
R^2	0.435	0.109	0.672	0.134	0.138	0.531

Notes:
(i) Subjects' belief adjustments are defined as subjects' posteriors minus priors. Bayesian belief adjustments are defined as Bayesian posteriors minus subjects' priors.
(ii) Analysis uses OLS regressions with robust standard errors in parentheses.

A.5 Belief adjustments - excluding rank 1 and rank 4

In Table 8, I replicate the regression analysis in Table 2 of Section 4.2, excluding subjects who are ranked first or fourth in their reference group.

Table 8: Belief adjustments - excluding rank 1 and rank 4

$$Beliefadjustment_i = \beta_0 + \beta_1 Bayesbeliefadj_i + \beta_2 Good\ news_i + \beta_3 Bayesbeliefadj_i * Good\ news_i + \epsilon_i$$

	No-Resolution			Resolution		
	Good news (1)	Bad news (2)	Diff-in-diff (3)	Good news (4)	Bad news (5)	Diff-in-diff (6)
β_1	0.579	0.046	0.046	0.453	0.526	0.526
	(0.145)	(0.198)	(0.198)	(0.431)	(0.244)	(0.244)
β_2			-0.460			-0.166
			(0.186)			(0.246)
β_3			0.533			-0.073
			(0.246)			(0.495)
Constant	-0.125	0.335	0.335	-0.177	-0.011	-0.011
	(0.109)	(0.151)	(0.151)	(0.208)	(0.131)	(0.131)
Observations	25	25	50	25	25	50
R^2	0.316	0.001	0.641	0.070	0.134	0.535

Notes:
(i) Subjects' belief adjustments are defined as subjects' posteriors minus priors. Bayesian belief adjustments are defined as Bayesian posteriors minus subjects' priors.
(ii) Analysis uses OLS regressions with robust standard errors in parentheses.

A.6 Ex-post rationalization - excluding wrong belief adjustments

In Table 9, I replicate the regression analysis in Table 3 of Section 4.3, excluding subjects with belief adjustments in the wrong direction.

Table 9: Ex-post rationalization of information - excluding wrong belief adjustments

	No-Resolution		Resolution	
Dependent variable	Importance study performance (1)	Importance job performance (2)	Importance study performance (3)	Importance job performance (4)
Good news	0.919	1.129	0.152	0.483
	(0.453)	(0.458)	(0.453)	(0.439)
IQ test score	-0.037	-0.114	0.036	-0.070
	(0.085)	(0.087)	(0.085)	(0.087)
Prior belief	-0.712	-1.207	-1.008	-0.939
	(0.424)	(0.430)	(0.363)	(0.355)
Observations	95	95	93	93
Pseudo R^2	0.034	0.054	0.032	0.026

Notes:
(i) Subjects' stated importance of the IQ test for study and job performance is measured on a seven-point Likert scale.
(ii) Analysis uses Ordered Logistic Regressions with standard errors in parentheses.

B Payments by ranks

Figure 3 shows the distribution of payments for each rank in the group.

Figure 3: Payments by ranks

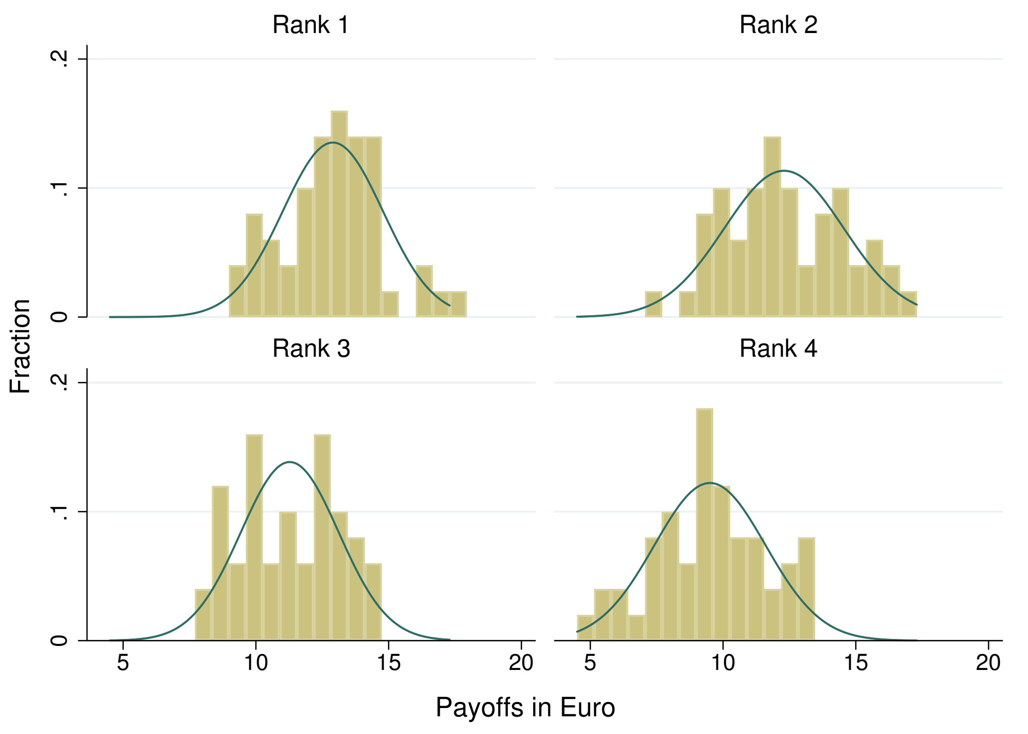

C Experimental instructions

GENERAL INSTRUCTIONS (on paper)

Welcome to this experiment! Please read the instructions carefully.

At the end of the experiment, you will be paid in cash according to your decisions and the decisions of other participants. In addition, you will receive a fixed payment of 4 Euro for your punctual appearance. Please make sure that your mobile phone is switched off. During the experiment, it is not allowed to communicate with other participants, use mobile phones, or start other programs on the computer. If you violate this rule, we regrettably must exclude you from the experiment and all payments.

If you have questions, please raise your hand. A lab manager will then come to your place and answer your question quietly.

Belief elicitation instructions

During the experiment, you will give your estimates for the likelihood of four different scenarios of an event. The likelihood that you will report will influence your earnings. For each estimate, you can receive an additional payoff of 2 euros. The payoff mechanism is designed such that you have the highest chance of receiving an additional payoff of 2 euros when you report your best estimate.

In the following, we will explain the payoff mechanism in detail. We will use the event "average temperature in Germany in 2018" as an example. This example is for **illustrative purposes** only and will be replaced by another event in the experiment.

Assume in the following that there are four possible scenarios for the "average temperature in Germany in 2018", and that exactly one of the scenarios has occurred.

- Scenario A: The average temperature in Germany in 2018 was below 9 degrees Celsius.
- Scenario B: The average temperature in Germany in 2018 was at least 9 degrees Celsius and below 10 degrees Celsius.
- Scenario C: The average temperature in Germany in 2018 was at least 10 degrees Celsius and below 11 degrees Celsius.
- Scenario D: The average temperature in Germany in 2018 was over 11 degrees Celsius.

In the experiment, it would now be the task to give your assessment for the likelihood of the occurrence of each respective scenario. Since only one of these scenarios has occurred, the sum of the probabilities adds up to 100%.

After you have made your assessment for the different scenarios, the computer will randomly select exactly **one scenario** as payoff relevant. This selection is random and does not mean that this scenario occurred.

The computer then randomly selects **a number X between 0 and 100**. The probability to be selected is equal for each number.

Payoff:

- If your specified likelihood for the selected scenario is at least as high as the number X, then you will receive 2 Euros if the scenario has occurred.
- If, on the other hand, your specified likelihood is lower than the number X, then you receive 2 euros with a probability of X%.

According to these rules, it is always beneficial for you to report the likelihood that you truly believe.

For example, assume that your true estimate for the probability of scenario A is 50% and you specify a probability of 30%. Then it can happen that the computer selects scenario A for the payout and the number 40 is taken for X. In this case, your probability of winning 2 Euros is 40%. If you had entered 50%, you would, according to your true estimate, win the 2 euros with a probability of 50% - exactly when scenario A occurred.

Control questions:

In order to increase your understanding of the payoff mechanism, we now ask you to answer some control questions on screen. Therefore, we will use the example above, "Average temperature in Germany in 2018". Your answers to these questions will not affect your payouts in the experiment. However, we will not proceed to the next phase of the experiment until all participants have answered the questions correctly. You may keep this leaflet during the experiment.

INSTRUCTIONS (on screen)

Control questions

The following control questions relate to the exemplary event "average temperature in Germany in 2018" with the following four scenarios:

- Scenario A: the average temperature in Germany in 2018 was below 9 degrees Celsius.
- Scenario B: the average temperature in Germany in 2018 was at least 9 degrees Celsius and below 10 degrees Celsius.
- Scenario C: the average temperature in Germany in 2018 was at least 10 degrees Celsius and below 11 degrees Celsius.
- Scenario D: the average temperature in Germany in 2018 was over 11 degrees Celsius.

Assume that your best estimate for the probability of scenario A is 50%, scenario B is 30% and scenario C is 15%.

1. Which of the following answers maximizes your chance of a payoff of 2 euros?
   - A=25%, B=25%, C=25%, D=25%
   - A=50%, B=10%, C=15%, D=25%
   - A=50%, B=30%, C=15%, D=5%
   - A=100%, B=0%, C=0%, D=0%

Suppose that you reported your true beliefs and the computer has randomly selected scenario C and the number X equal to 25.

2. What is your chance to win 2 Euros?
   - 25%
   - 20%
   - 15%

3. Would you have had a higher chance of winning the 2 Euros by reporting 40% instead of 15%?
   - Yes
   - No

OK

QUIZ STAGE

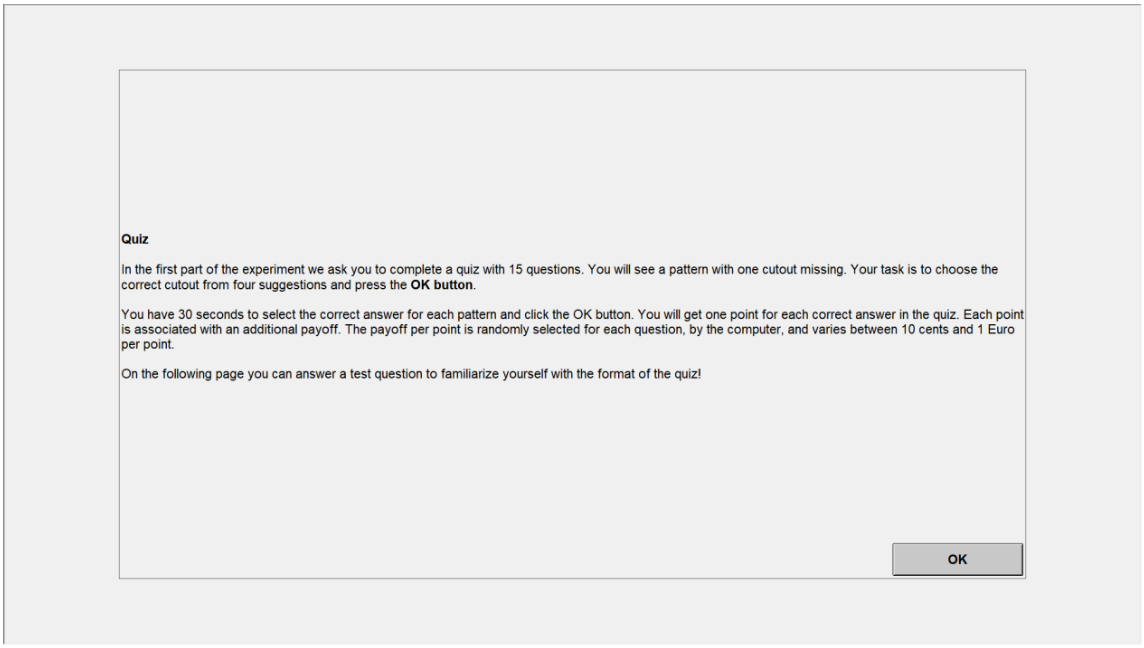
Quiz
In the first part of the experiment we ask you to complete a quiz with 15 questions. You will see a pattern with one cutout missing. Your task is to choose the correct cutout from four suggestions and press the OK button.
You have 30 seconds to select the correct answer for each pattern and click the OK button. You will get one point for each correct answer in the quiz. Each point is associated with an additional payoff. The payoff per point is randomly selected for each question, by the computer, and varies between 10 cents and 1 Euro per point.
On the following page you can answer a test question to familiarize yourself with the format of the quiz!
OK

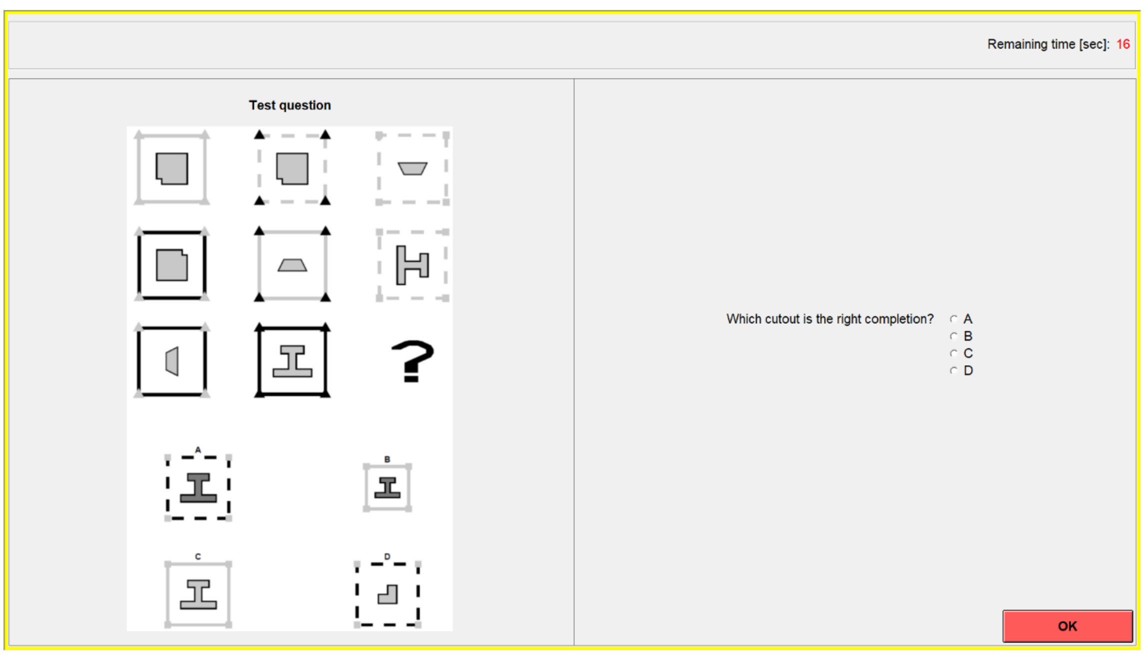
Remaining time [sec]: 16
Test question
A
B
C
D
Which cutout is the right completion?
A
B
C
D
OK

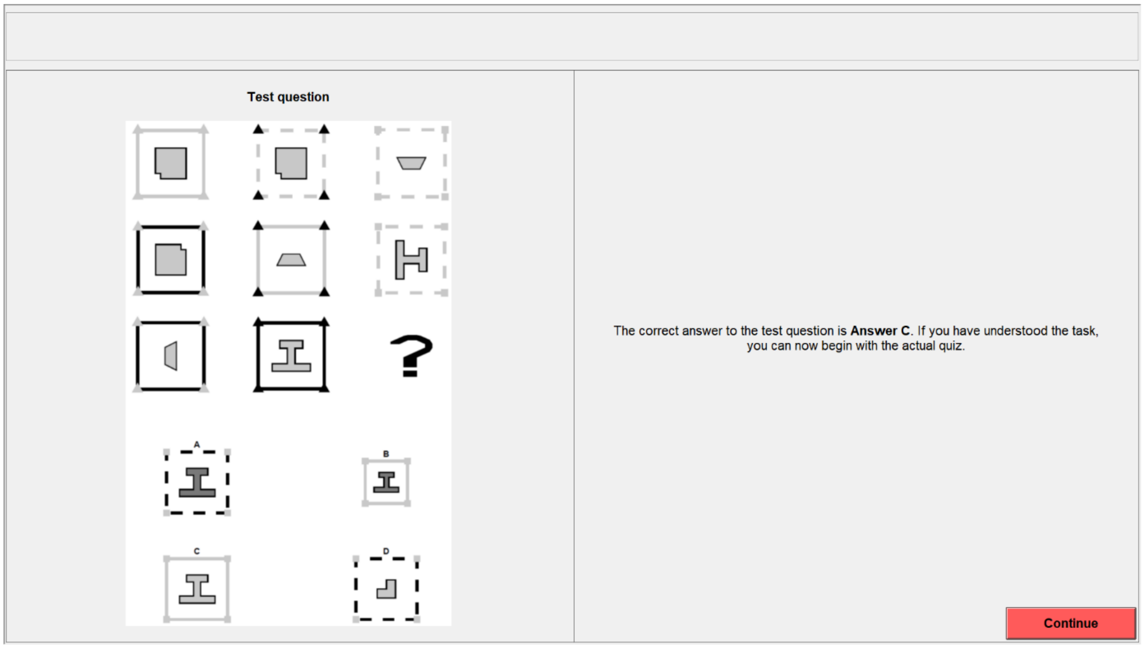

Test question
A
B
C
D
The correct answer to the test question is **Answer C**. If you have understood the task, you can now begin with the actual quiz.
Continue

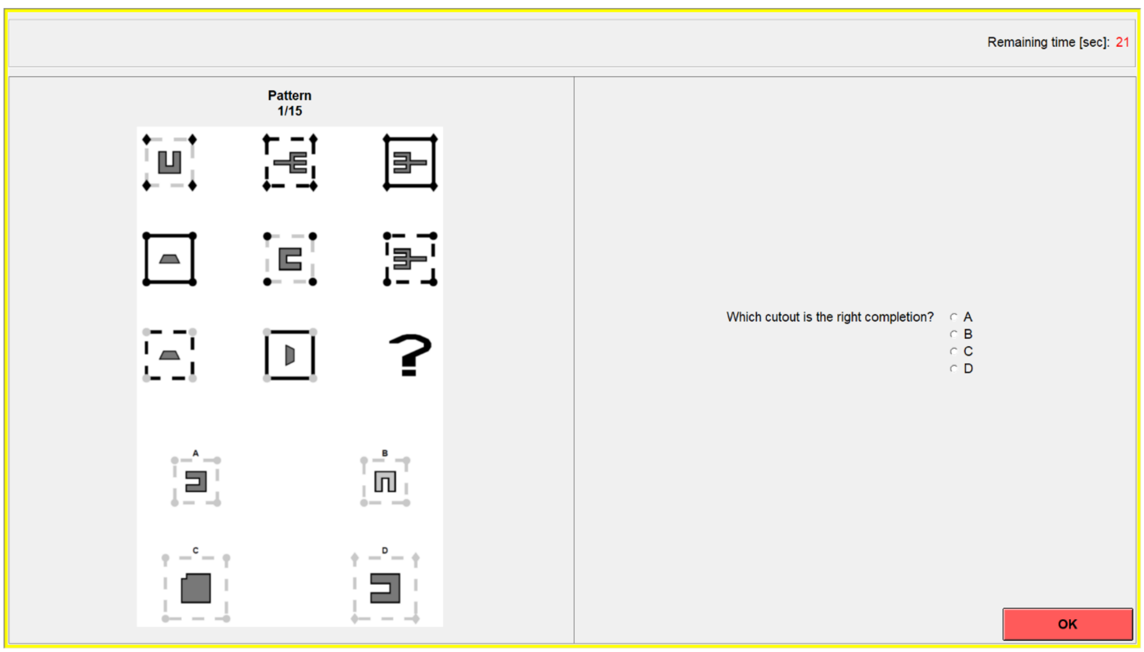

Remaining time [sec]: 21
Pattern
1/15
A
B
C
D
Which cutout is the right completion?
A
B
C
D
OK

PRIOR BELIEF ELICITATION

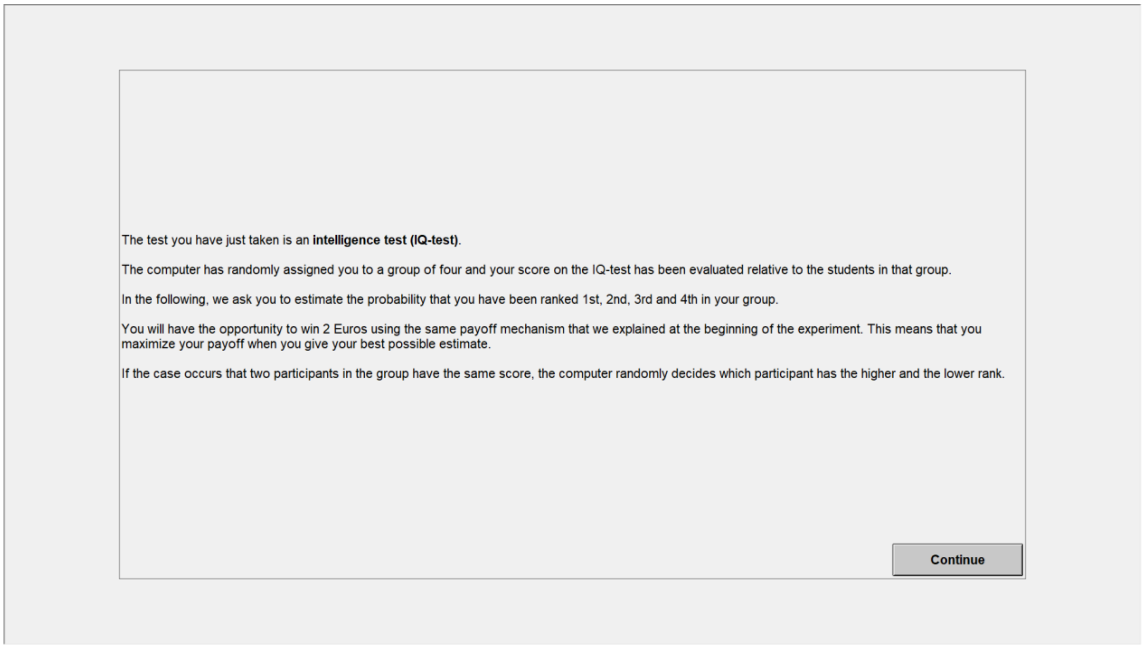

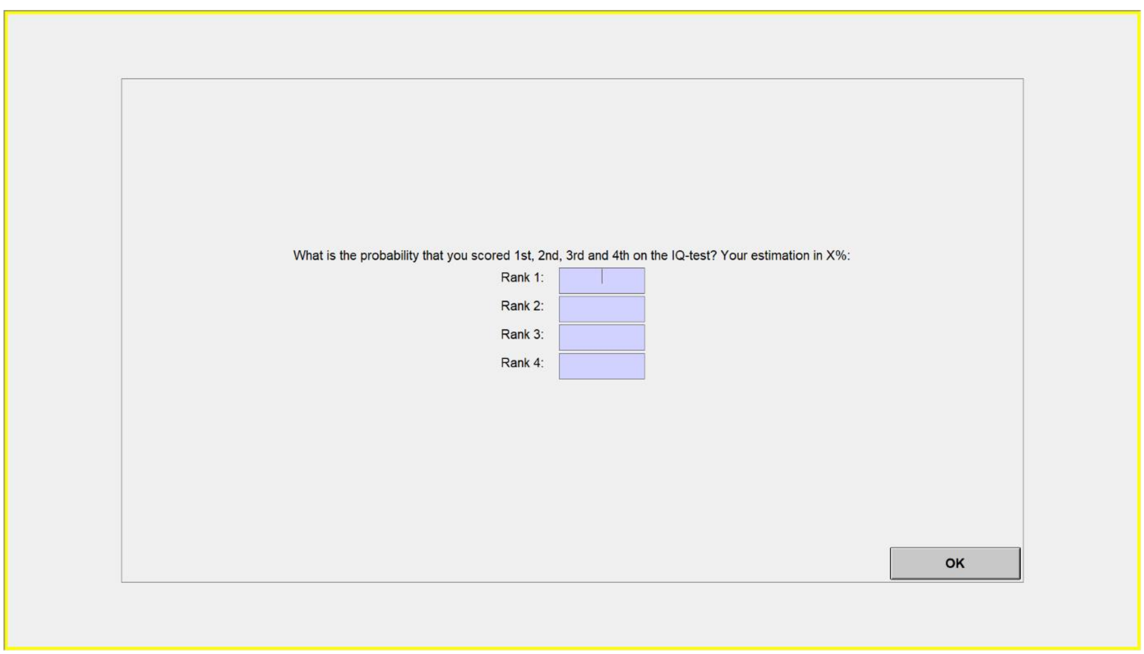

FEEDBACK STAGE AND POSTERIOR BELIEF ELICITATION (RESOLUTION-TREATMENT)

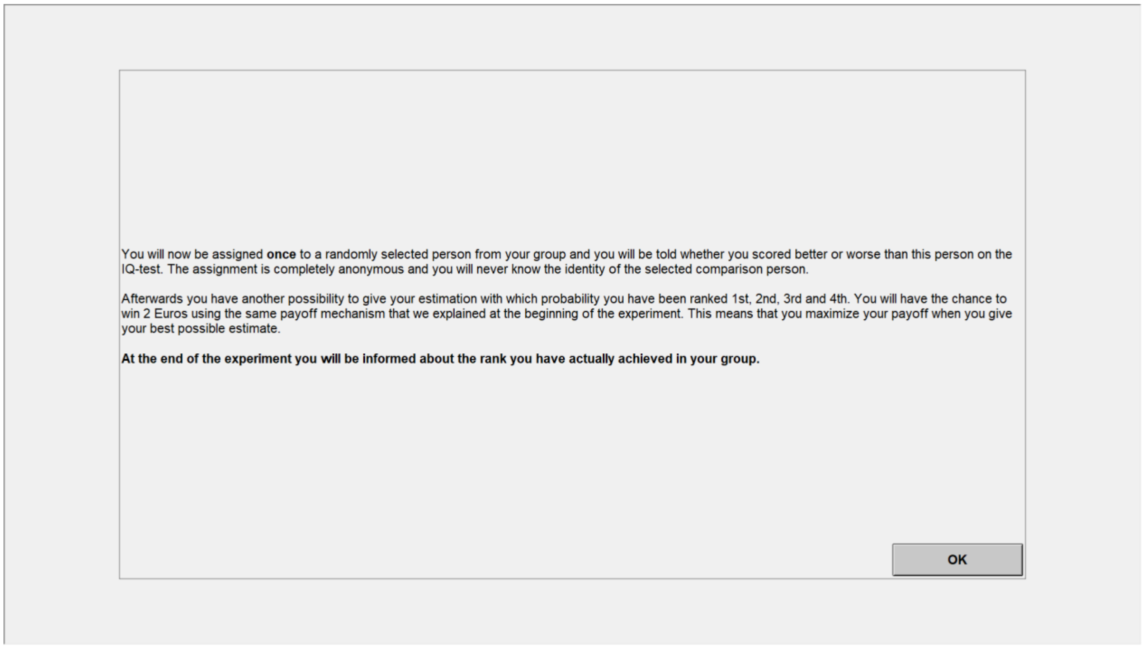

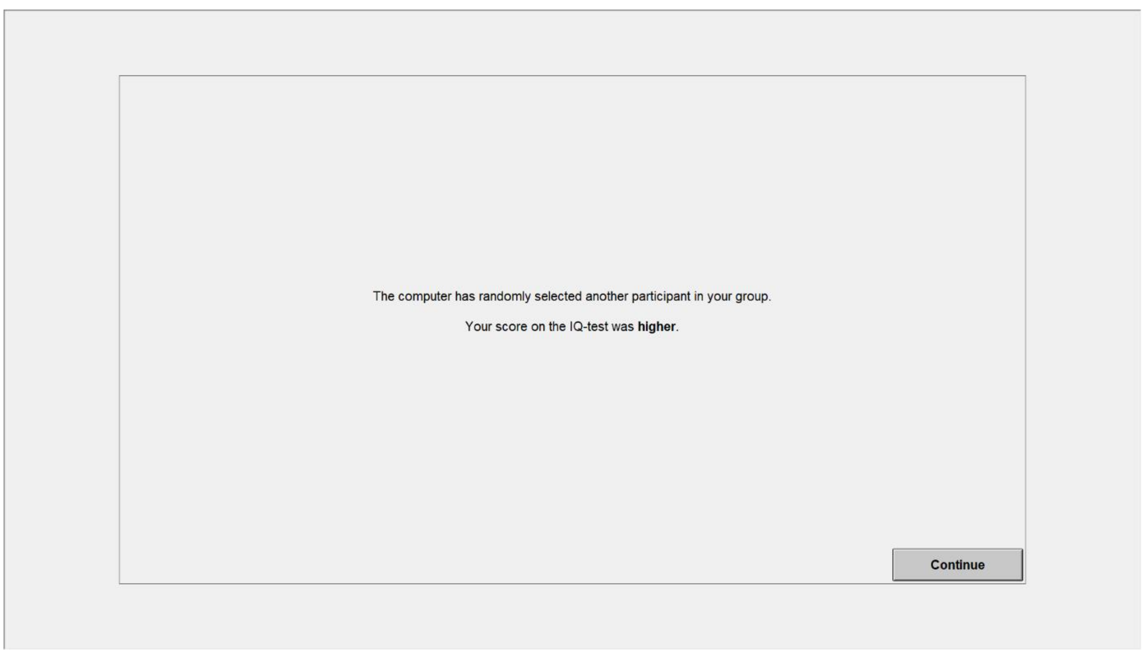

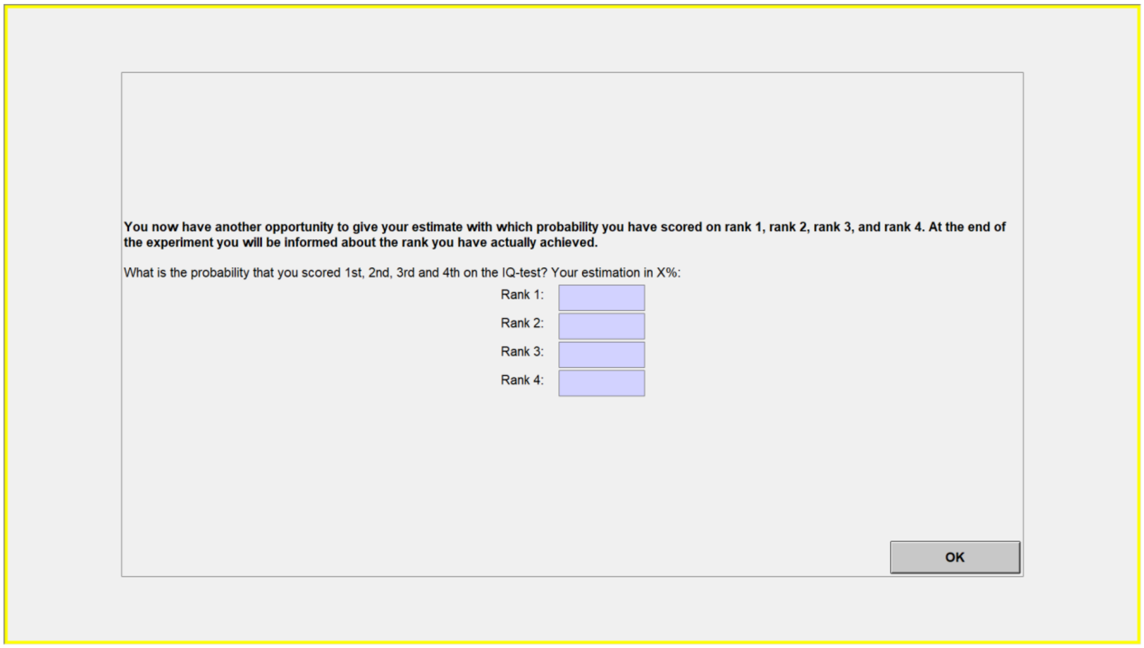

FEEDBACK STAGE AND POSTERIOR BELIEF ELICITATION (NO-RESOLUTION-TREATMENT)

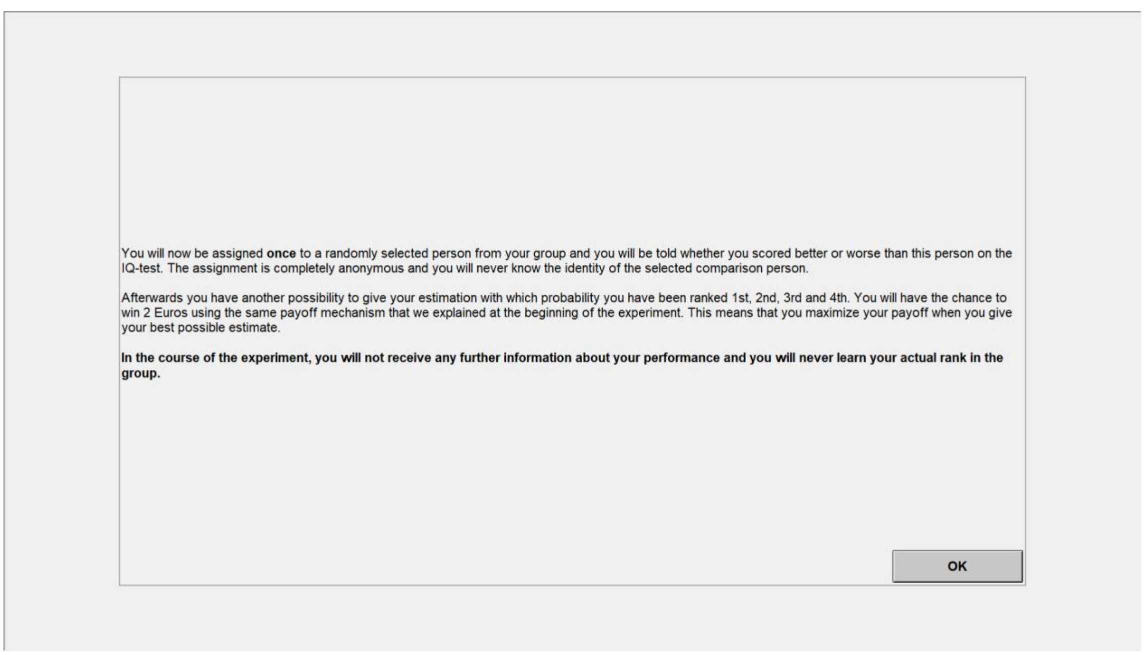

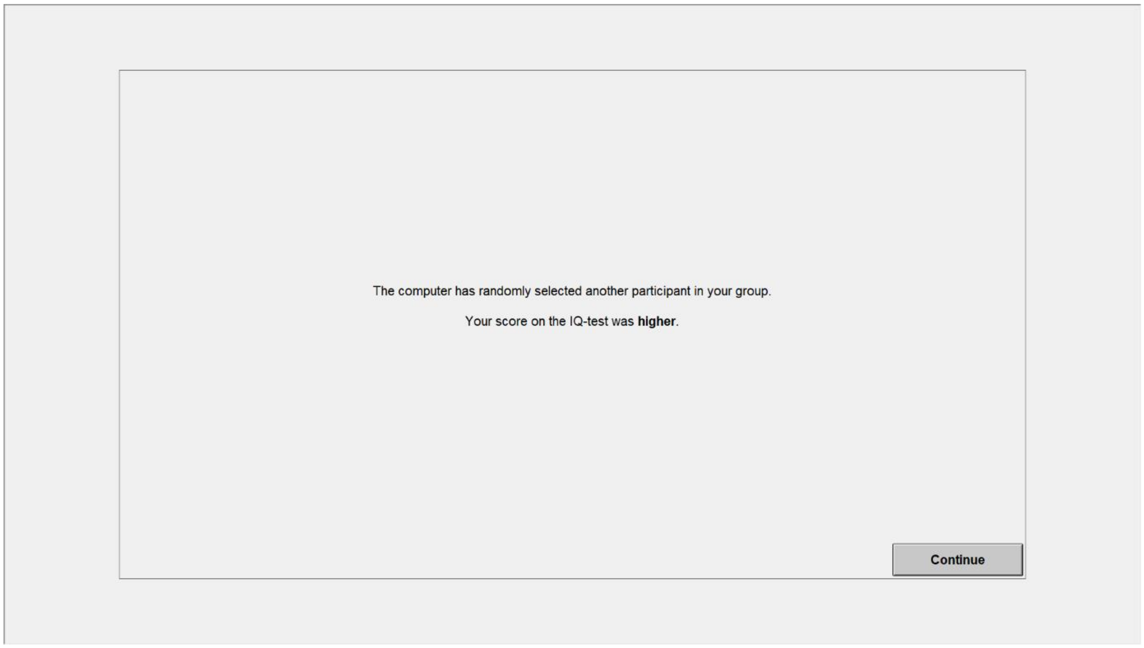
The computer has randomly selected another participant in your group.
Your score on the IQ-test was **higher**.
Continue

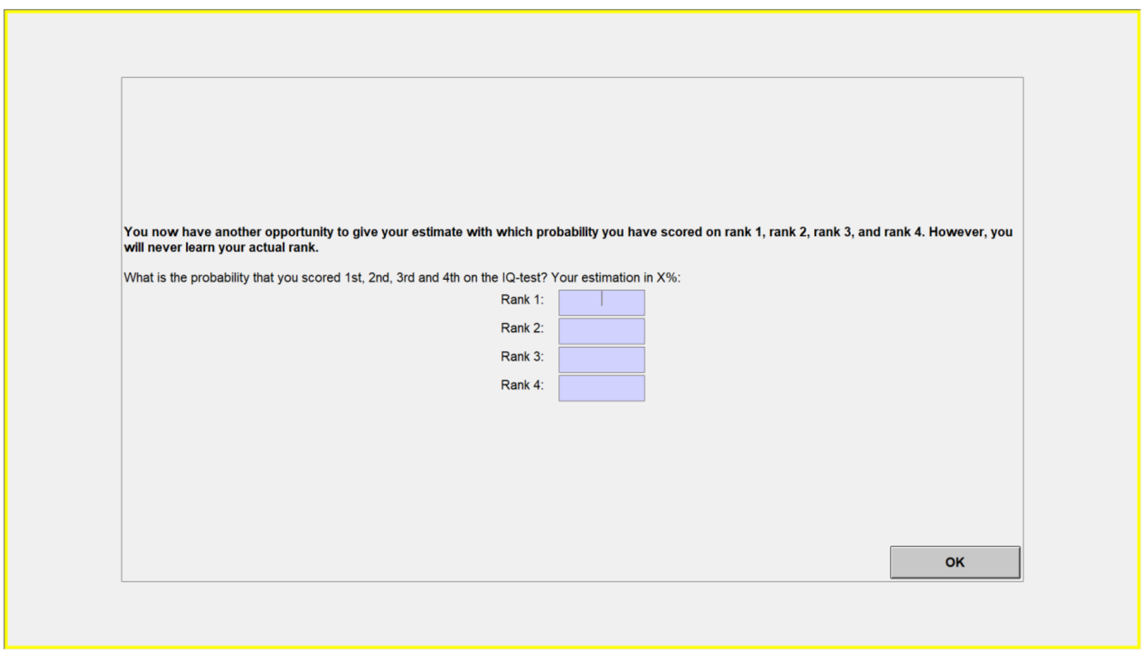
You now have another opportunity to give your estimate with which probability you have scored on rank 1, rank 2, rank 3, and rank 4. However, you will never learn your actual rank.
What is the probability that you scored 1st, 2nd, 3rd and 4th on the IQ-test? Your estimation in X%:
Rank 1:
Rank 2:
Rank 3:
Rank 4:
OK

QUESTIONNAIRE

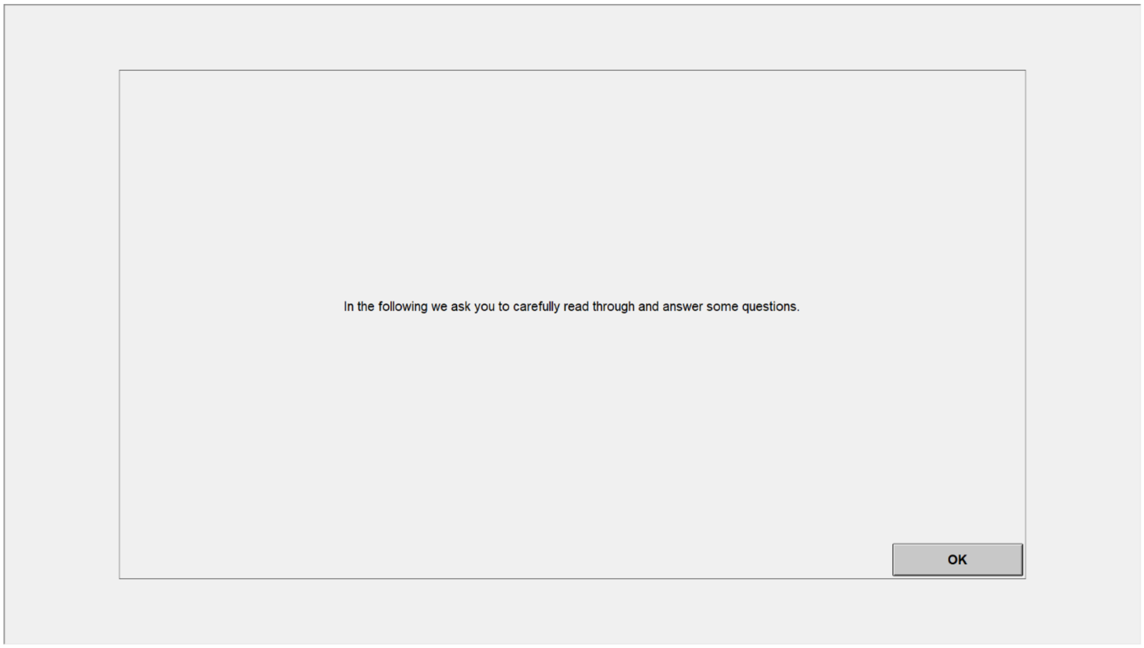

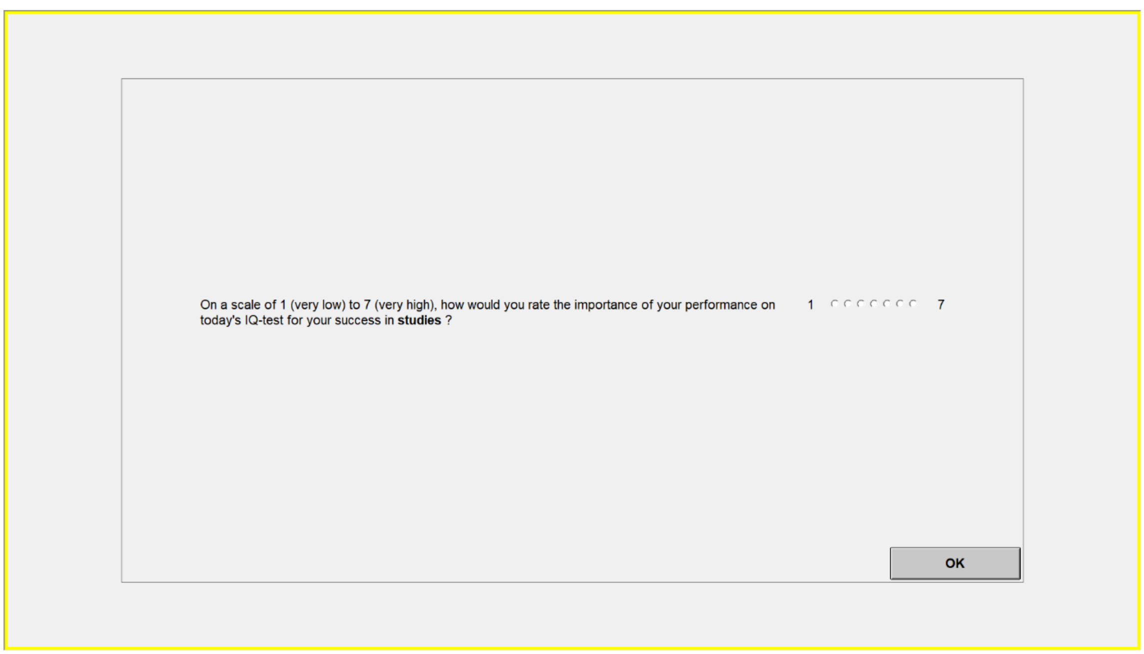

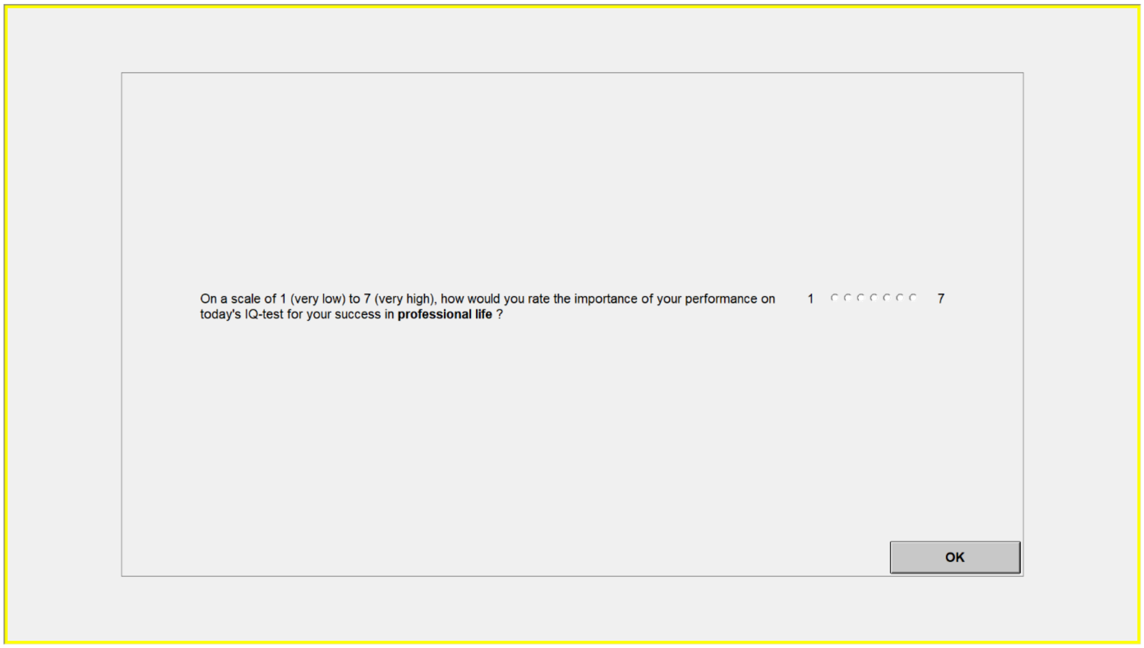

RESOLUTION OF TRUE RANK (RESOLUTION-TREATMENT)

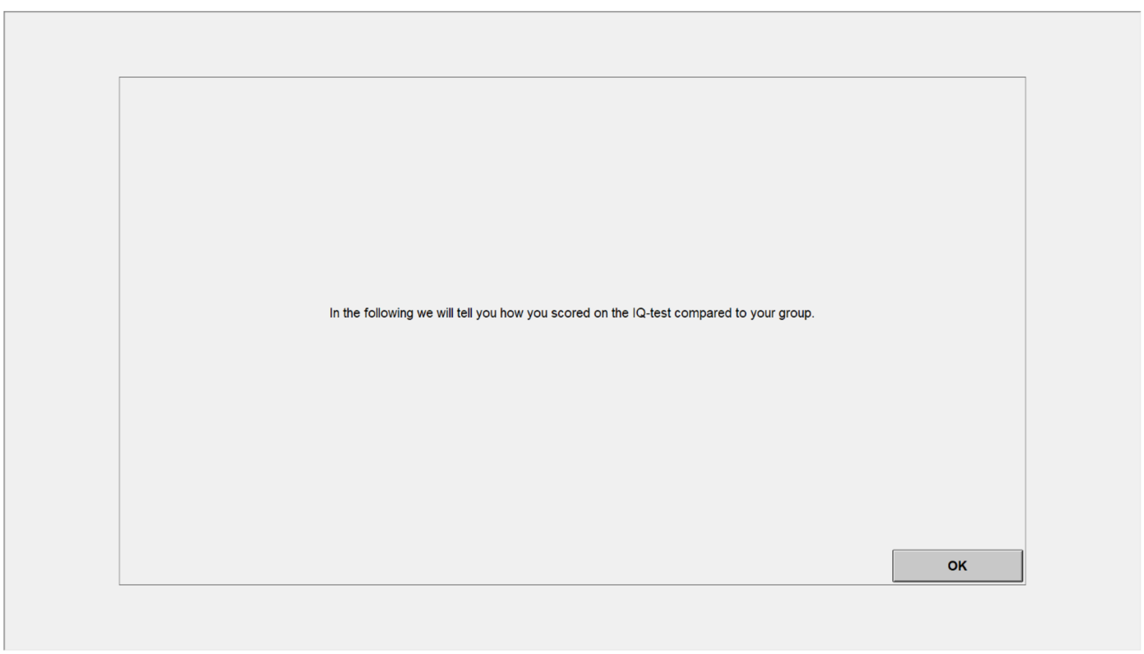

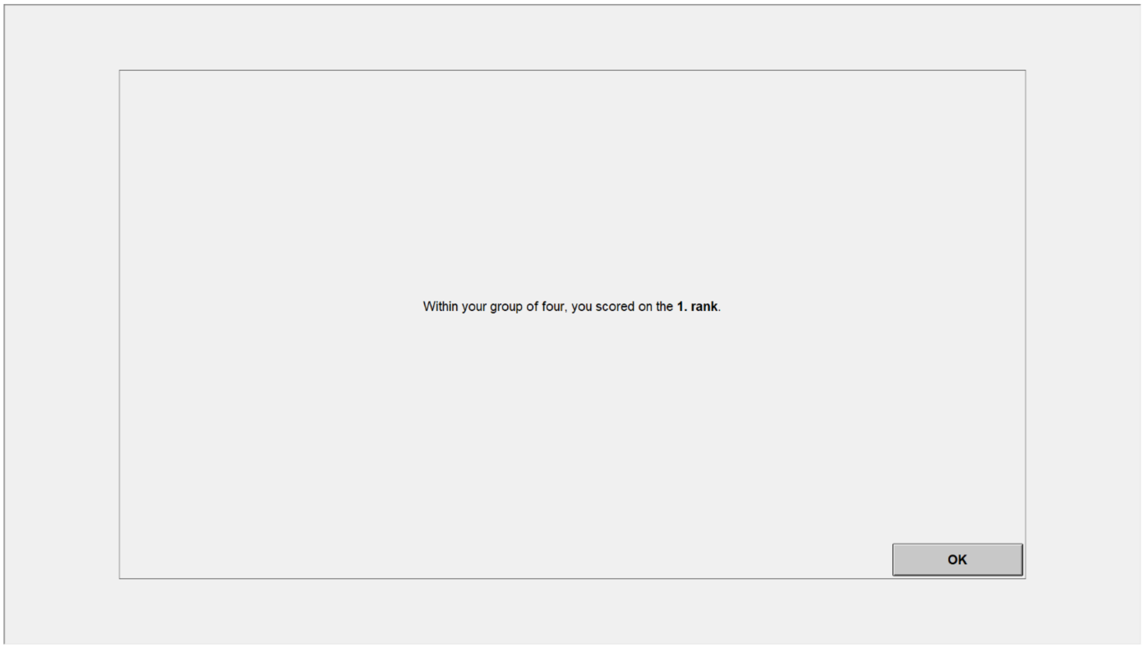

DEMOGRAPHICS AND PAYOFF

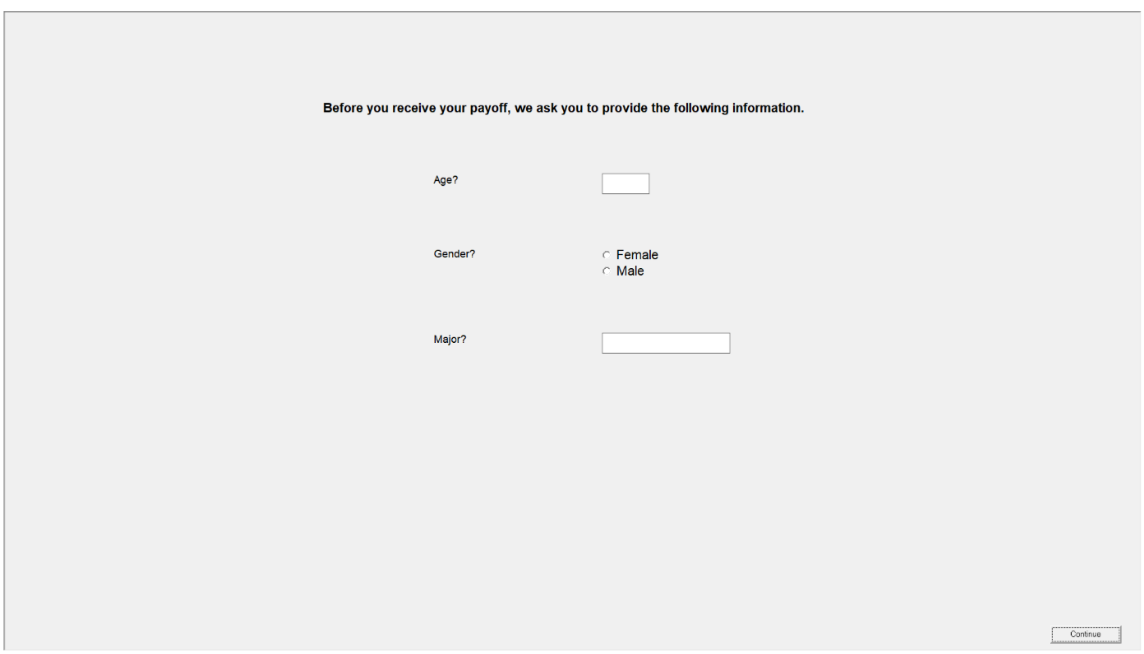

Thank you very much for your participation in the experiment!

Your payoff in the experiment is 11.60 Euro.

Please fill out the receipt. As soon as your PC number is called up by the laboratory manager, you can collect your payoff by handing in the receipt, the instructions and the PC number.

Previous Munich Papers in Political Economy:

Betz, Timm and Amy Pond. “Political Ownership”, MPPE No. 1/2020, Munich. (ISSN)2701-3456

Chatziathanasiou, Konstantin, Hippel, Svenja and Michael Kurschilgen. “Property, Redistribution, and the Status Quo. A laboratory study”, MPPE No. 2/2020, Munich. (ISSN)2701-3456

Becker, Annette, Hottenrott, Hanna and Anwesha Mukherjee. “Division of Labor in R&D? Firm Size and Specialization in Corporate Research”, MPPE No. 3/2020, Munich. (ISSN)2701-3456

Steinert, Janina Isabel, Satish, Rucha Vasumati, Stips, Felix and Sebastian Vollmer. “Commitment or Concealment? Impacts and Use of a Portable Saving Device: Evidence from a Field Experiment in Urban India, MPPE No. 4/2020, Munich. (ISSN)2701-3456

Janz, Nicole and Luca Messerschmidt. ”Unravelling the `race to the bottom' argument: How does FDI affect different types of labour rights?”, MPPE No. 5/2020, Munich. (ISSN)2701-3456.

Chowdhury, Subhasish M., Esteve-González, Patricia and Anwesha Mukherjee. ”Heterogeneity, Leveling the Playing Field, and Affirmative Action in Contests”, MPPE No. 6/2020, Munich. (ISSN)2701-3456

Drobner, Christoph. ”Motivated Beliefs and Anticipation of Uncertainty Resolution”, MPPE No. 7/2020, Munich. (ISSN)2701-3456

Impressum:

ISSN: **2701-3456**
Editors: Tim Büthe, Hanna Hottenrott
Associate Editors: Timm Betz, Sebastian Goerg, Eugénia da Conceição Heldt, Michael Kurschilgen, Amy Pond, Sebastian Schwenen, Janina Steinert, Matthias Uhl
Managing Editor: Luca Messerschmidt

Contact: Technicial University of Munich, Arcisstraße 21, 80333 München
mppe@gov.tum.de, mppe@wi.tum.de
https://www.wi.tum.de/mppe/
Twitter: @MunichPapers