INSTITUTE for REPLICATION

No. 65
I4R DISCUSSION PAPER SERIES

Motivated Beliefs & Anticipation of Uncertainty Resolution: A Note

Alisher Batmanov

Idaliya Grigoryeva

September 2023

I4R DISCUSSION PAPER SERIES

I4R DP No. 65

Motivated Beliefs & Anticipation of Uncertainty Resolution: A Note

Alisher Batmanov[1], Idaliya Grigoryeva[1]

[1]*University of California, San Diego/USA*

SEPTEMBER 2023

E-Mail: joerg.peters@rwi-essen.de
RWI – Leibniz Institute for Economic Research
Hohenzollernstraße 1-3
45128 Essen/Germany
www.i4replication.org

ISSN: 2752-1931

Motivated Beliefs & Anticipation of Uncertainty Resolution: A Note

Alisher Batmanov* Idaliya Grigoryeva†

University of California, San Diego

July 12, 2023

Abstract

Drobner (2022) examines the effect of manipulating experimental subjects' expectations about uncertainty resolution in learning about their performance on their belief updating patterns in an ego-relevant domain. In their preferred empirical specification, the author finds that individuals update their beliefs optimistically as they exhibit a higher belief adjustment in response to good compared to bad news only when they do not expect resolution of underlying uncertainty about their performance in an IQ test and neutrally when they know they will find out their relative performance at the end of the experiment. First, we reproduce the all of the paper's findings without identifying any coding errors. Second, we test the robustness of the results to (1) adding individual covariates and (2) excluding subjects who exhibit a fundamental error in their belief updating from the analysis. We find no substantial changes in the main coefficients of interest with the inclusion of demographic variables in the analysis, consistent with demonstrated balance in covariates between the two experimental groups. Yet, several of the main estimates lose statistical significance and change from conservatism (under-updating) to over-inference (over-updating) in some conditions on the subset of participants excluding those who exhibit fundamental errors in belief updating.

*abatmanov@ucsd.edu

†idagri@ucsd.edu

1 Introduction

A growing strand of literature on motivated reasoning and motivated beliefs postulates that beliefs fulfill key psychological and functional needs of the individual (such as self-confidence, moral self-esteem, hope and anxiety reduction, etc.), thereby individuals extract not only instrumental, but also direct utility from forming such beliefs (Bénabou and Tirole, 2016). Of particular interest is how such motivated beliefs evolve in response to new information. Empirical evidence from experiments in Economics and Psychology in ego-relevant settings has been mixed with some papers pointing to optimistic belief updating characterized by stronger response to "good news" compared to receiving "bad news", while others reporting neutral or even pessimistic belief-updating patterns (Coutts, 2019; Eil and Rao, 2011; Ertac, 2011).

Studying belief adjustments in the response to informative but noisy signals, Drobner (2022) finds in a lab experiment at a large German university that individuals update their beliefs optimistically only when they do not expect resolution of underlying uncertainty about their performance in an ego-relevant task such as an IQ test. The author states that switching off the channel of uncertainty resolution enables individuals to derive direct utility from inflated beliefs about their relative performance, which is not the case in the setting where subjects are immediately exposed to the potentially unpleasant truth. In a controlled lab experiment with 200 university student participants featuring a variant of belief-updating task with a signal on the relative performance in an IQ test as the underlying ego-relevant characteristic, the author exogenously varies subjects' expectations about the resolution of uncertainty by providing them with information that their true rank will either remain uncertain or will be revealed at the end of the experiment. Comparing the data on subjects' belief adjustments against the Bayesian benchmark reveals that subjects update their beliefs optimistically only if they expect their true relative rank to remain uncertain, whereas they exhibit neutral belief updating pattern under imminent resolution of uncertainty despite being incentivized to provide their best guesses about their true relative performance through the BDM payment mechanism under both conditions.

In the present paper, we investigate whether the author's empirical results are reproducible and replicable, and further test their robustness to two specification checks: (1) adding covariates to the main regression specifications and (2) changing how the incorrect belief updaters are defined to test the robustness of results on a different subsample. We are able to fully replicate all of the tables and figures in the original paper and the appendix before turning to our robustness replicability analysis.

In their original analysis, Drobner (2022) relies on data from the lab experiment where the author collected data on three individual characteristics (age, gender, university major), but did not utilize them anywhere in the analysis, nor verified the balance across treatment groups. In our re-analysis, we extend the main specifications including the covariates given the potential sensitivity of the small samples, and constrain the sample excluding the participants who do not update their posterior beliefs "correctly" utilizing the complete information received in the signal (with a different interpretation of incorrect updating complementary to the robustness check in the original paper). We find that the main results are robust to the inclusion of the covariates (consistent with the fact that we found the covariates to be balanced across the treatment groups). Yet, several of the main estimates lose statistical significance on the subset of participants whose actions are more consistent with paying attention and understanding of the experimental tasks and signals, which may be driven by a sizable sample size reduction using this criterion. In addition, we find that excluding from the analysis those subjects who showcase a major probabilistic reasoning error results in a switch from conservative belief updating to over-inference for those subjects who do not expect immediate uncertainty resolution.

We are grateful to the author for clarifying our questions related to participant recruitment and data processing, which helped us conduct an informed replication study.

2 Reproducibility

The author has supplied the complete raw data files (10 individual session-specific Excel spreadsheet) that are merged and cleaned in the code the author provided to produce the processed dataset for analysis and all of the tables and figures from the main paper and the appendix. The main code files include a 160-line STATA do file for data cleaning and a 200-line STATA do file creating all tables and figures.

Figure 1: Experimental Study Design Diagram

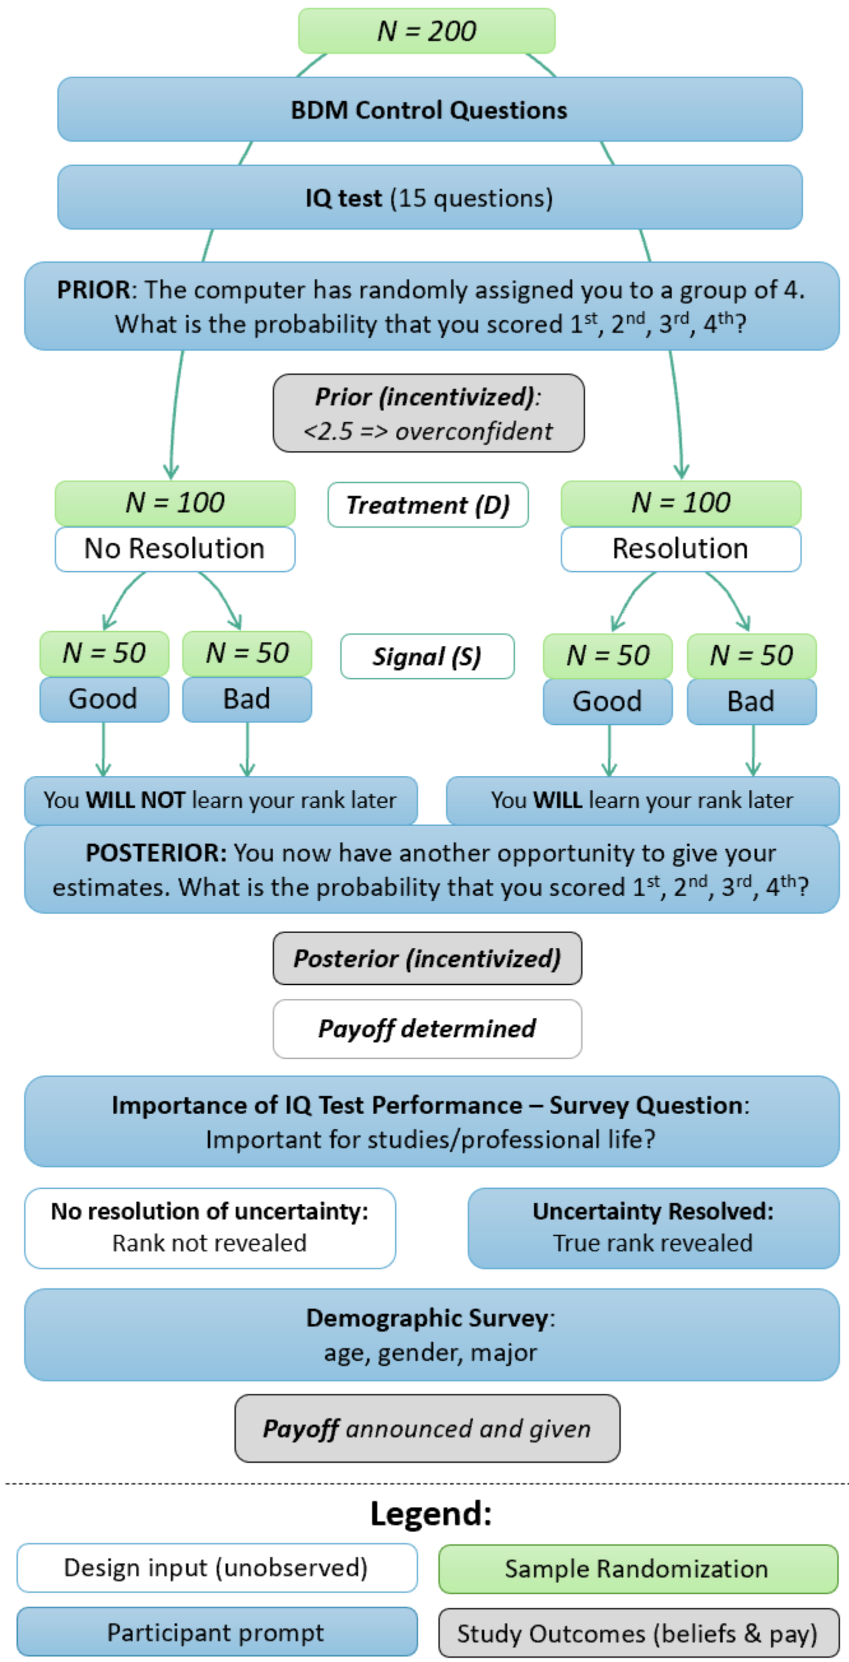

Notes: original diagram created for the replication study.

We did not identify any coding errors in either file, and the two files produce all tables (except for Table 1 described below) and figures from the main paper and the appendix. Table 1 in the published paper constitutes the author's analytical assessment comparing the main findings of a set of papers (1) in Economics and (2) in Psychology and Neuroscience on motivated belief updating, with no additional materials or notes included in the replication for reproducing this table.

In addition to replicating the main results and conducting additional robustness checks, we created a diagram describing the sequential logic of the experimental study to assist the readers in following how the hypotheses relate to the parts of the study and the main results (see Figure 1). It demonstrates the sample splitting by uncertainty resolution and by received signal, as well as the order of belief elicitation and IQ-related questions.

3 Replication

3.1 Motivation & Summary of Previous Studies

The paper is largely motivated by the contradictory results on optimistic vs. neutral / pessimistic belief updating varying between Economics studies, on the one hand, and Psychology and Neuroscience studies, on the other. To showcase these differences, Drobner (2022) provides summary findings on the type of belief updating and uncertainty resolution across a sample of 8 Economics and 15 Psychology & Neuroscience papers (See Table 1 in Drobner (2022)). The paper states that the selection of papers to be included focused on "studies about belief updating with purely ego-relevant information", and there is no additional commentary on the assessment of the presence of uncertainty and main results, whether it is straightforward and unambiguous across all studies or not. To cross-check the Table 1 summary results, we randomly selected 2 out of 8 Economics papers (25% coverage) and 3 out of 15 Psychology papers (20% coverage). The papers were selected with a random number generator resulting in the following list:

- Möbius et al. (2022) (Econ #6): optimistic with ambiguous uncertainty (link)
  - ✓ IQ test with no indication as to whether their actual test performance will be revealed before the end of the experiment *(Drobner's assessment is confirmed)*

- Zimmermann (2020) (Econ #8): short-run neutral, long-run optimistic (link)
  - ✓ IQ test with no indication as to whether their actual test performance will be revealed before the end of the experiment (*Drobner's assessment is confirmed)*
- Garrett et al. (2018) (Neuroscience #4): optimistic without uncertainty (link)
  - ✓ Eliciting beliefs about the probability of an adverse event happening to the participant in the future (e.g., burglary or cancer) (*Drobner's assessment is confirmed)*
- Kuzmanovic et al. (2016) (Neuroscience #6): optimistic without uncertainty (link)
  - ✓ Eliciting beliefs about the probability of an adverse event happening to the participant in the future (e.g., burglary or cancer) (*Drobner's assessment is confirmed)*
- Sharot et al. (2012) (Neuroscience #15): optimistic without uncertainty (link)
  - ✓ Eliciting beliefs about the probability of an adverse event happening to the participant in the future (e.g., burglary or cancer) (*Drobner's assessment is confirmed)*

Out of the five randomly selected papers, we confirmed Drobner's (2022) assessment of their type of belief updating and uncertainty resolution in each paper. The one thing to highlight is that at least in the studies we randomly selected, the subjects are not explicitly told about the conditions or uncertainty resolution – whether they will find out the outcome or not at the end of the experiment, and it is merely assumed by the participants (hence, "ambiguous" uncertainty resolution in Econ papers) or implied by the context of the question (experiencing an event sometime in one's lifetime) that they wouldn't find out the real outcome in the Psych and Neuroscience studies. In Economics studies that we checked, there were no details identified about subject expectations for the resolution of uncertainty, implying that the belief updating results are a composition of people who might expect and those not expecting to find out how well they did in the experimental task. The novelty in Drobner's approach is in being able to compare the belief updating between the subjects who are explicitly told they will know they IQ-test based rank at the end of the experiment to those who are told that they will not.

3.2 Main Results

Prior to delving into the replication of primary empirical results in the paper, we first assess the balance on covariates resulting from randomization. As can be seen from Figure 1, the main experimental variation in the paper is randomly assigning one half of subjects to the *No Resolution* group and the other half to *Resolution* treatment group. The only difference between the two is that the study participants in the former group are explicitly informed that their true rank relative to the other members in their group based on the IQ test will not be revealed at the end of the experiment, while those in the latter group are told, in contrast, that they will learn about their performance-based rank at the end of the study. In addition, in the final stage of the experiment subjects complete a three-question demographic survey in which they report their age, sex and university major. Throughout the paper, Drobner (2022) does not consider the covariates, so in our analysis, we created a binary variable *Econ_Major* based on the subjects' open-text answer to distinguish those subjects whose major is Economics, which we later use as a proxy for higher concept familiarity or higher likelihood of prior exposure to similar experiments.[1]

Table 1 provides a balance test for the three covariates across the *No Resolution* and *Resolution* groups. As can be seen, treatment assignment appears to be balanced across the treatment groups with no mean differences being statistically significant (p values much greater than 0.01 for each covariate in a two-sided t-test).

Next, we assess whether the set of estimation results that jointly tests the main hypothesis of the paper are replicable:

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.*

[1] As subjects' majors were recorded using an open-text entry recorded in German, we construct the *Econ_Major* variable including undergraduate students in Economics ("*Wirtschaft*") and Management ("*Betriebswirtschaftslehre*" (BWL) and "*Technologie- und Managementorientierte Betriebswirtschaftslehre*" (TUM-BWL) and their variations), and Master's in Management students ("TUM-BWL Master" and its variations). For a specific list of manually added majors in each category, you can consult the variables *'major'* and *'econ_major'* in our code files in the additional cleaning code section.

Table 1: Balance Test by Resolution Condition Assignment

		(1)		(2)	(2)-(1)	(2)-(1)
		No Resolution		Resolution	Pairwise t-test	Pairwise t-test
Variable	N	Mean/(SE)	N	Mean/(SE)	Mean difference	P-value
Age	100	24.500	100	23.970	–0.530	0.581
		(0.787)		(0.548)		
Female	100	0.450	100	0.430	–0.020	0.777
		(0.050)		(0.050)		
Econ Major	100	0.390	100	0.410	0.020	0.774
		(0.049)		(0.049)		

Notes: Comparison of randomly assigned treatment groups of No Resolution vs. Resolution.
***, **, * indicate 1, 5, 10% significance.

3.2.1 Aggregate Beliefs

Prior to looking into the potential asymmetry in belief updating in response to good vs. bad news, the paper first starts with describing beliefs at the aggregate level. In the prior belief elicitation stage, the subjects are asked to report their expected probability of being in each possible rank within a randomly assigned group of 4 participants based on their IQ test score. Rank 1 implies the person got the highest score among the four group members and rank 4 implies the lowest score (lowest performance). The average expected rank, calculated as the average rank weighted by reported expected probability of being it in, is 2.38 in the *No Resolution* treatment arm, which is not significantly different from the average of 2.42 in the *Resolution* group. Running a simple two-sided t-test confirms this result with the p-value of 0.637, which corresponds to the one in the paper. When pooled, the average prior belief of 2.4 is statistically different from the rational belief of 2.5 at the aggregate level ($p = 0.009$ and $p = 0.004$ for the two-sided and one-sided tests, respectively).

When comparing aggregate posterior beliefs to the corresponding Bayesian benchmark, for both treatment groups the average reported posterior ranks are not significantly different from the theoretical predictions: 2.38 against 2.42 in the *No Resolution* group ($p = 0.215$ for the two-sided t-test and $p = 0.108$ for the one-sided test, which is close to being marginally significant at the 10% level) and 2.443 against 2.452 in the *Resolution* group ($p = 0.837$ and $p = 0.418$ for the two-sided and one-sided tests, respectively). Hence, the first result of the paper summarized below replicates precisely.

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

3.2.2 Belief Adjustments

To address the main research question in the paper, the analysis proceeds by evaluating belief updating patterns by the valence of the signal (good vs. bad), i.e. by comparing the subjects' belief adjustments (difference between the reported posteriors and Bayesian posteriors) separately for the good and bad news. Table 2 in the paper (Drobner, 2022) summarizes the results by reporting the coefficient estimates from the regression

$$Beliefadjustment_i = \beta_0 + \beta_1 BayesBeliefAdj_i + \beta_2 Goodnews_i + \beta_3 BayesBeliefAdj_i \times Goodnews_i + \epsilon_i$$

where the actual belief adjustment, defined as the difference between the reported posterior belief and the prior belief ($\hat{B}_1 - \hat{B}_0$), is regressed on the Bayesian belief adjustment, defined as the difference between the Bayesian prescribed posterior and the reported prior ($B_1 - \hat{B}_0$), in addition to the dummy indicating receiving a positive signal (vs. the baseline negative) and the interaction term between the two.

When running the regression above for *No Resolution* and *Resolution* groups as in the paper, we were able to exactly replicate the estimates in Table 2 in the original paper (see our Table 2) Columns 3 and 6 in Table 2 confirm the hypothesis of optimistic belief updating activation by showing that subjects respond much stronger to receiving good news when they do not expect uncertainty resolution ($\hat{\beta}_3 = 0.589$, significantly different from zero at 1% level). Subjects expecting their rank to be revealed, on the contrary, exhibit neutral belief updating pattern and do not differentially react to good versus bad news ($\hat{\beta}_3 = -0.115$, not significantly different from zero).

This result is also presented in Figure 1 in the original paper, where the subjects' actual belief adjustment is plotted against the Bayesian belief adjustment, separately for two treatment arms and for good and bad news. Using the replication package provided by the author, we obtained the same figure which portrays the main finding that is summarized in the paper as follows (the identical figure is omitted in this replication report):

Result 2. *Subjects' belief adjustments follow Bayesian belief adjustments more closely for good*

Table 2: Replicating Drobner's Table 2

	(1) NoRes-Good	(2) NoRes-Bad	(3) NoRes-DiD	(4) Resol-Good	(5) Resol-Bad	(6) Resol-DiD
Bayes. Belief	0.665***	0.076	0.076	0.530**	0.645**	0.645**
Adjustment	(0.088)	(0.180)	(0.180)	(0.218)	(0.249)	(0.249)
Good News			-0.359***			-0.133
			(0.129)			(0.166)
Bayes Adj ×			0.589***			-0.115
Good News			(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
$R^2 adj.$	0.41	-0.02	0.64	0.11	0.12	0.52
F stat	57.05	0.18	66.06	5.89	6.69	31.90

Notes: This table is the exact replica of Table 2 in Drobner (2022) obtaining the same results. Analysis uses OLS regressions with robust s.e. Subjects' belief adjustments are defined as subjects' posterior beliefs minus prior beliefs. Bayesian belief adjustments (Bayes. Belief Adj. / Bayes Adj.) are defined as Bayesian posterior beliefs minus subjects' prior beliefs (higher prior belief value corresponds to a 'higher' expected rank, implying lower performance – e.g., rank 4 is the lowest rank).

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.

3.2.3 Ex-post Rationalization

The final section of the main analysis in the paper looks at whether subject ex-post rationalize information by manipulating their beliefs about ego-relevance of the IQ test based on whether they received good or bad news at the feedback stage. Table 3 in the original paper (Drobner, 2022) uses the self-reported importance of the IQ test for subjects' study and job performance as the outcome variable recorded using a 7-point Likert scale. Controlling for the actual IQ test score and reported prior beliefs, the stated categorical importance is regressed on *Good news* separately for the two treatment groups. We were able to re-create the estimates from the original Table 3 (see our Table 3), which show that when subjects do not expect uncertainty resolution, they report significantly higher importance of the IQ test after receiving the good news relative to receiving the bad news. Columns 1 and 2 in Table 3 below illustrate that the coefficients for *Good news* are significant for both study and job performance importance: upon receiving a positive signal subjects' reported importance is 0.992 and 1.168 points higher, respectively. In contrast, columns 3 and 4 reveal that in the *Resolution* group there is no significant gap in reported importance in response to good or bad news (0.122 and 0.317 points, for study and job performance, respectively). This result suggesting a potential underlying mechanism for the observed difference in belief updating patterns is summarized as follows in the paper.

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

3.3 Robustness

In this section, we extend the original analysis in two ways. First, we make use of the demographic information collected as part of the experiment but not present in the analysis to assess whether the main results are robust to inclusion of covariates (we expect the results to be robust since the covariates were balanced across the *Resolution* and *No Resolution* treatments, but there might be differences in the small 50-person subsamples varying by resolution treatment and signal valence. Second, we run the main regressions (from Table 2 and Table 3) on the sub-samples of the original experiment subject pool by excluding those who violated the basic belief updating rule, namely those who put a positive weight on the highest (lowest) rank upon receiving bad (good) news stating that they performed worse (better) than at least one other group member.

3.3.1 Controlling for covariates

At the end of the experiment, the author elicits subjects' age, gender and major. Throughout the paper (including the Appendix), however, this information is not utilized in any part of the analysis. To enrich the set of results and assess their robustness to the inclusion of covariates, we re-run the regression equation that produces Table 2 in the original paper but additionally add

Table 3: Replicating Drobner's Table 3 – Ex-Post Rationalization of IQ Importance

	(1) NoRes-Study	(2) NoRes-Work	(3) Resol-Study	(4) Resol-Work
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	0.04	0.06	0.03	0.02

Notes: This table is the exact replica of Table 3 in Drobner (2022) obtaining the same results. Outcome in each column is self-reported importance of IQ for studying (cols. 1, 3) or professional activities (referred to as 'Work' above) (cols. 2, 4) on a 7-point Lickert scale. Analysis uses OLS regressions with robust s.e. Prior beliefs about rank are elicited before receiving signals and before learning about resolution of uncertainty.

three demographic variables as regressors.

$$Beliefadjustment_i = \beta_0 + \beta_1 BayesBeliefAdj_i + \beta_2 Goodnews_i + \beta_3 BayesBeliefAdj_i \times Goodnews_i + \gamma_1 Age_i + \gamma_2 Female_i + \gamma_3 Econmajor_i + \epsilon_i$$

The results are reported in Table 4. It can be observed that all the main results generally remain qualitatively intact with only minor differences in the magnitude of most estimates. In particular, when controlling for individual covariates, subjects' belief adjustments in *No Resolution* treatment arm do not follow the Bayesian belief adjustment benchmark ($\hat{\beta}_1 \approx 0$ in column 2), which slightly increases the coefficient for $Bayesbeliefadj_j \times Goodnews_i$ in column 3 of Table 4 (0.627 compared to 0.589 without covariates). It remains highly significant as before, which once again confirms the hypothesis about optimistic belief updating activation in absence of uncer-

Table 4: Extending Drobner's Table 2 with Covariates

	(1) NoRes-Good	(2) NoRes-Bad	(3) NoRes-DiD	(4) Resol-Good	(5) Resol-Bad	(6) Resol-DiD
Bayes. Belief	0.688***	-0.001	0.032	0.564**	0.599**	0.638**
Adjustment	(0.094)	(0.181)	(0.173)	(0.213)	(0.248)	(0.243)
Good News			-0.393***			-0.140
			(0.141)			(0.171)
Bayes Adj ×			0.627***			-0.075
Good News			(0.185)			(0.315)
Age	0.007	-0.004	0.001	-0.010	-0.015**	-0.014***
	(0.006)	(0.005)	(0.003)	(0.012)	(0.006)	(0.005)
Female	-0.046	-0.131	-0.071	0.121	-0.090	0.006
	(0.083)	(0.114)	(0.069)	(0.117)	(0.141)	(0.091)
Econ Major	-0.119	0.097	-0.032	0.035	0.249	0.151*
	(0.080)	(0.107)	(0.067)	(0.114)	(0.150)	(0.089)
Observations	50	50	100	50	50	100
$R^2 adj.$	0.43	-0.03	0.63	0.08	0.21	0.54
F stat	18.98	0.72	35.24	2.05	5.14	18.69

Notes: This table is replicating Table 2 in Drobner (2022) with added covariates. Analysis uses OLS regressions with robust s.e. Subjects' belief adjustments are defined as subjects' posterior beliefs minus prior beliefs. Bayesian belief adjustments (Bayes. Belief Adj. / Bayes Adj.) are defined as Bayesian posterior beliefs minus subjects' prior beliefs (higher prior belief value corresponds to a 'higher' expected rank, implying lower performance – e.g., rank 4 is the lowest rank).

tainty resolution. This matches our expectation that the covariates balanced by treatment do not affect the main experimental result.

Similar to results for Table 2, the results on ex-post rationalization depicted in Table 3 in the original paper remain largely unchanged once we control for participants' age, gender and majoring in Economics. As before, for both studies and professional life, subjects in *No Resolution* group report significantly higher importance after they are exposed to good news about their performance on the IQ test as can be seen in columns 1 and 2 in Table 5. This rationalization effect is not present in *Resolution* treatment arm as evidenced by the coefficients for *Good news* in columns 3 and 4, which are not statistically different from zero.

3.3.2 Excluding impossible rank belief reporters

Perhaps the most interesting part of the replication report is our exploration of the robustness of results to excluding participants with another type of incorrect belief updating in response to the signal. In the original paper, the author conducts several robustness checks that are included in the Appendix, among them is a table comparing the effect sizes after excluding subjects who adjust their beliefs in the opposite direction to the prescribed Bayesian belief adjustment and/or do not adjust them at all.

In this section, we complement these tests by following a slightly different approach to defining "wrong" belief updating. In the original paper, the author focuses on zero and "wrong" belief adjustments defined by subjects' expected posterior minus prior rank beliefs as in Drobner (2022): $\hat{B}_1 - \hat{B}_0 = \sum_{j=1}^{4} j \cdot \hat{b}_{j1} - \sum_{j=1}^{4} j \cdot \hat{b}_{j0}$, where $\hat{b}_{jt}$ is the probabilistic belief about each possible rank $j \in \{1, 2, 3, 4\}$ in periods $t \in \{0, 1\}$) (See the original paper Appendix – sections A.1 & A.2).

Here, we propose to exclude subjects who exhibit a different fundamental belief updating mistake. Receiving "good" or "bad" news in the experiment is defined as being matched with another subject of their reference group of 4 and learning whether their score in the IQ test was higher or lower than the score of one other randomly matched group member (the study design ensures that pair-matches result in one half of the subjects receiving a "good" signal and the other half receiving a "bad" signal of receiving a lower score than one other group member). This implies that upon receiving *good* news, subjects must infer that they cannot be ranked 4th in their group (there is at least one other group member who necessarily got a lower score than the subject) and, thus, must report the probability of rank 4 as 0 ($\hat{b}_{41} = 0$) to be consistent with the signal. Similarly, receiving *bad* news can be interpreted with certainty as zero chance of being

Table 5: Extending Drobner's Table 3 with Covariates– Ex-Post Rationalization of IQ Importance

	(1) NoRes-Study	(2) NoRes-Work	(3) Resol-Study	(4) Resol-Work
Good News	1.008**	1.144**	0.247	0.374
	(0.454)	(0.458)	(0.441)	(0.425)
IQ test score	0.013	-0.077	0.081	-0.030
	(0.082)	(0.083)	(0.082)	(0.080)
Prior Belief	-0.680	-1.189***	-0.949***	-0.986***
	(0.437)	(0.433)	(0.357)	(0.349)
Age	0.035	-0.003	0.080**	0.048
	(0.024)	(0.023)	(0.040)	(0.041)
Female	-0.030	0.010	0.775*	0.690*
	(0.392)	(0.383)	(0.396)	(0.386)
Econ Major	0.165	0.618*	0.977**	0.555
	(0.375)	(0.375)	(0.386)	(0.381)
Observations	100	100	100	100
Pseudo R^2	0.05	0.07	0.06	0.04

Notes: This table is the exact replica of the Table 2 in Drobner (2022) with added covariates. Outcome in each column is self-reported importance of IQ for studying (cols. 1, 3) or professional activities (referred to as 'Work' above) (cols. 2, 4) on a 7-point Lickert scale. Analysis uses OLS regressions with robust s.e. Prior beliefs about rank are elicited before receiving signals and before learning about resolution of uncertainty.

ranked 1st and subject should thereby report the probability of rank 1 as zero ($\hat{b}_{11} = 0$) as their posterior belief. Violating this simple rule is an indicator of a fundamental probabilistic reasoning error, which we will hereinafter refer to as ***impossible rank belief***.

Excluding subjects who exhibit this type of updating error results in 113 out of 200 subjects remaining for the analysis (59 from the no-resolution group and 54 from the resolution group), which we should say is a strinkingly high number for us. Table 6 reproduces the analysis of belief adjustments for this sub-sample as in Table 2 in the original paper. Notably, while the estimates for the *Resolution* group remain highly significant for both good and bad news, we now see that the coefficients on $BayesBeliefAdj_i$ are greater than one in the magnitude, thus indicating over-updating and not conservatism. More strikingly, for *No Resolution* treatment the optimistic belief updating pattern characterized by differential response to good versus bad news is no longer present. In the original Table 2 in the paper the coefficient for $BayesBeliefAdj_i$ under *Bad news* is 0.076, which is much smaller than 0.442 in column 2 of Table 6 once we exclude the subjects with impossible rank beliefs. This difference drives down the gap in belief adjustments between positive and negative signals in *No Resolution* group, which translates into not statistically significant coefficient estimate of 0.299 in front of $BayesBeliefAdj_i \times Goodnews_i$ in column 3 of Table 6.

Likewise, by excluding subjects with this error from the ex-post rationalization analysis (based on the original Table 3), we see a similar discrepancy with the original results in the paper. Columns 1 and 2 in Table 7 reveal that the coefficients for *Good news* in *No Resolution* group are no longer significant, which comes in contrast to the original evidence in favor of ex-post rationalization.

Overall, this highlights the importance of conducting various tests of robustness to "mistakes" in participant understanding and interpretation of the experimental questions. In this case, we found a large number of participants acting inconsistently with basic probability concepts, which resulted in substantial changes in the significance of the main results. Still, this is also driven by the fact that the sample size drops substantially once we exclude these "impossible-rank" believers, implying that we do not invalidate the main results, but rather highlight the sensitivity of the estimates to potential lack of attention or misinterpretation of the experimental information by the participants.

Table 6: Extending Drobner's Table 2 – Excluding Subjects With Impossible Rank Beliefs

	(1) NoRes-Good	(2) NoRes-Bad	(3) NoRes-DiD	(4) Resol-Good	(5) Resol-Bad	(6) Resol-DiD
Bayes. Belief	0.741***	0.442	0.442	1.112***	1.051***	1.051***
Adjustment	(0.110)	(0.307)	(0.305)	(0.220)	(0.222)	(0.222)
Good News			-0.401**			-0.079
			(0.194)			(0.138)
Bayes Adj ×			0.299			0.061
Good News			(0.325)			(0.312)
Constant	-0.079	0.323*	0.323*	0.038	0.116	0.116
	(0.060)	(0.185)	(0.184)	(0.079)	(0.113)	(0.113)
Observations	34	25	59	27	27	54
$R^2 adj.$	0.49	0.04	0.79	0.38	0.49	0.81
F stat	45.54	2.07	77.14	25.51	22.51	61.89

Notes: This table is replicating Table 2 in Drobner (2022) with added covariates on different subsamples. Subsample is a subset of the corresponding column in Table 2 excluding individuals with impossible rank beliefs. Impossible rank beliefs refer to reporting a positive probability of the highest (lowest) rank when given a bad (good) signal. Analysis uses OLS regressions with robust s.e. Subjects' belief adjustments are defined as subjects' posterior beliefs minus prior beliefs. Bayesian belief adjustments (Bayes. Belief Adj. / Bayes Adj.) are defined as Bayesian posterior beliefs minus subjects' prior beliefs (higher prior belief value corresponds to a 'higher' expected rank, implying lower performance – e.g., rank 4 is the lowest rank).

4 Conclusion

Overall, we were able to replicate all of the main results in the paper including the analysis of aggregate beliefs, belief adjustments by the valence of the signal and ex-post rationalization of information using the replication package provided by the author and we did not identify any apparent coding errors. We further extended the analysis making use of the covariates elicited in the experiment to assess the robustness of estimation results to the inclusion of individual

Table 7: Extending Drobner's Table 3 – Excluding Subjects With Impossible Rank Beliefs

	(1) NoRes-Study	(2) NoRes-Work	(3) Resol-Study	(4) Resol-Work
Good News	0.633	0.507	-0.187	0.676
	(0.719)	(0.613)	(0.630)	(0.587)
IQ test score	-0.052	-0.150	0.100	-0.115
	(0.139)	(0.157)	(0.111)	(0.131)
Prior Belief	-0.574	-1.937***	-0.870*	-1.067*
	(0.820)	(0.689)	(0.513)	(0.623)
Age	0.057***	0.027	0.079	0.056
	(0.020)	(0.021)	(0.068)	(0.059)
Female	-0.510	0.002	-0.285	0.461
	(0.609)	(0.552)	(0.665)	(0.675)
Econ Major	0.303	1.261**	0.773	0.031
	(0.563)	(0.619)	(0.582)	(0.562)
Observations	59	59	54	54
Pseudo R^2	0.05	0.11	0.06	0.04

Notes: This table is replicating Table 3 in Drobner (2022) with added covariates on different subsamples Outcome in each column is self-reported importance of IQ for studying (cols. 1, 3) or professional activities (referred to as 'Work' above) (cols. 2, 4) on a 7-point Lickert scale. Subsample is a subset of the corresponding column in Table 2 excluding individuals with impossible rank beliefs. Impossible rank beliefs refer to reporting a positive probability of the highest (lowest) rank when given a bad (good) signal. Analysis uses OLS regressions with robust s.e.

controls. With minor differences in specific estimates, we find all of the key conclusions in the paper holding, as expected from the balance of covariates across the treatment groups, which we also verified.

In addition, we complemented the analysis by testing the robustness of results to a differently defined set of subjects engaged in "incorrect" belief updating, which could be a signal of lack of attention or experimental task misinterpretation among some subjects despite the financial incentives. We exclude subjects who commit a fundamental belief updating mistake by reporting a positive posterior belief about being ranked the highest (lowest) in their group upon receiving the bad (good) news that they did worse (better) than one other group member. The author partially addresses the concern and conducts a robustness check by excluding subjects who update their beliefs in the wrong direction or do not update them at all (which does not change the qualitative results of the paper as per Appendix A.1 and A.2). Instead, we constrain our sample to those who are able to correctly utilize the signal and at the very least attach zero probability weight to the certainly *impossible* (lowest/highest) rank upon receiving good/bad news - this restriction results in neutral belief updating pattern for both groups and weaker evidence of ex-post rationalization.

In this case, we lose over 40% of participants across the treatment groups and obtain the results where we no longer find some of the key estimates from the original paper statistically significant. While not necessarily questioning the main results in the paper, this draws attention to the concern of sample size and "incorrect" belief updating among subjects, whose behavior is not necessarily informative for our desired exploration of optimistic vs. pessimistic belief updating. This highlights the importance of conducting more checks of the consistency of behavioral subject actions in experimental tasks and responses to the provided signals.

References

Bénabou, R. and Tirole, J. (2016). Mindful economics: The production, consumption, and value of beliefs. *Journal of Economic Perspectives*, 30(3):141–64.

Coutts, A. (2019). Good news and bad news are still news: Experimental evidence on belief updating. *Experimental Economics*, 22(2):369–395.

Drobner, C. (2022). Motivated beliefs and anticipation of uncertainty resolution. *American Economic Review: Insights*, 4(1):89–105.

Eil, D. and Rao, J. M. (2011). The good news-bad news effect: asymmetric processing of objective information about yourself. *American Economic Journal: Microeconomics*, 3(2):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(3):532–545.

Garrett, N., González-Garzón, A. M., Foulkes, L., Levita, L., and Sharot, T. (2018). Updating Beliefs under Perceived Threat. *Journal of Neuroscience*, 38(36):7901–7911.

Kuzmanovic, B., Jefferson, A., and Vogeley, K. (2016). The role of the neural reward circuitry in self-referential optimistic belief updates. *NeuroImage*, 133:151–162.

Möbius, M. M., Niederle, M., Niehaus, P., and Rosenblat, T. S. (2022). Managing Self-Confidence: Theory and Experimental Evidence. *Management Science*.

Sharot, T., Guitart-Masip, M., Korn, C. W., Chowdhury, R., and Dolan, R. J. (2012). How dopamine enhances an optimism bias in humans. *Current Biology*, 22(16):1477–1481.

Zimmermann, F. (2020). The Dynamics of Motivated Beliefs. *American Economic Review*, 110(2):337–363.