Essays in Behavioral and Experimental Economics

Inauguraldissertation

zur Erlangung des Grades eines Doktors
der Wirtschaftswissenschaften

durch

die Rechts- und Staatswissenschaftliche Fakultät der
Rheinischen Friedrich-Wilhelms-Universität Bonn

vorgelegt von

Robin Jakob Musolff

aus Hamburg

2026

Dekan:	Prof. Dr. Martin Böse
Erstreferent:	Prof. Dr. Florian Zimmermann
Zweitreferent:	Prof. Dr. Thomas Dohmen
Tag der mündlichen Prüfung:	4. Mai 2026

Acknowledgements

Writing up a dissertation is, in many ways, a process of looking back, not just at the research itself, but also at the person I was when I worked on each part of it.

When I was accepted into the Bonn-Berkeley Program in the final year of my Bachelor's, I had no real sense of what lay ahead. Spending my first year of the PhD in Berkeley was an invaluable experience, and I am deeply grateful to the entire 2018 Berkeley Economics PhD cohort for sharing the burdens of that first year with me. I also owe a great deal to the exposure to empirical research I gained while overseas, which played a pivotal role in my decision to shift my focus from theoretical work to Behavioral and Experimental Economics.

Returning to Bonn, I was welcomed with open arms by the Bonn PhD cohort. Taking Florian Zimmermann's topics course solidified my decision to specialize in Behavioral Economics, and I am especially grateful for his support during my Master's thesis and my first research project on motivated reasoning in financial domains. Though that project didn't make it into this dissertation, it was the first experiment I designed, ran, and analyzed. Florian's supervision was always insightful but never overbearing, allowing me the space to develop my own research skills. Working with him on two of the dissertation's chapters was an invaluable experience, and I also learned a great deal from supporting him as an RA.

I am also deeply thankful to Thomas Dohmen for his guidance as my second supervisor. His feedback was always thoughtful and generous, not only answering my questions but sparking entirely new ideas and perspectives.

I am grateful to Guy Yanay for our joint work on the perceptions of politically motivated reasoning. His insights and collaboration have been invaluable. I also want to thank Andrea Amelio and Alexander Laubel. Even though our projects didn't reach completion, the process of working on them was incredibly valuable. Each of these collaborations shaped my thinking, challenged me in new ways, and contributed to my growth as a researcher.

The pandemic hit just as I entered the dissertation phase, making an already challenging transition even more difficult. Fortunately, right before lockdowns began, I had moved into a shared flat with Jonah and Tim. I couldn't have asked for a better home during those years. I am beyond grateful for our friendship, for the countless FIFA matches, and for both of your perspectives on life.

To Lina, thank you for putting so many smiles on my face, for the unwavering support, and for the love you've given me. I feel absurdly at ease around you.

Finally, I thank my parents, Inge and Roland, and my brother, Leon. My parents have always supported my academic pursuits without ever pressuring me, striking that rare balance between encouragement and freedom. Perhaps somewhat at odds with my choice to pursue a PhD in Economics, they somehow managed to instill in me a love for music and creativity, which has been fundamental in shaping who I am. My brother has been an incredible role model and a source of inspiration throughout, showing me that an academic path is one worth following.

Robin Jakob Musolff
Bonn, September 2025

Contents

List of Figures

List of Tables

Introduction

Beliefs are central to human decision-making. They inform judgments, guide behavior and shape interpretations in both economic and political contexts. Assuming rational information processing, standard models often treat belief formation as a Bayesian updating process. In reality, however, beliefs can be shaped by cognitive limitations, social identity, and motivated reasoning—especially when information is complex, framed, or politically charged (Benjamin (2019)).

This dissertation explores how informational framing, model uncertainty, and political identity shape beliefs and subsequent behavior. Across four chapters based on experimental studies, it examines fairness judgments under inequality framings, beliefs and their link to actions under model uncertainty, misperceptions of others' political reasoning, and polarization dynamics after the 2024 U.S. election.

The first chapter investigates how informational framing shapes judgments about inequality. In a controlled experiment, participants are asked to distribute additional income between two individuals, one more and one less advantaged based on the outcome of a real-effort task. The key manipulation varies whether the additional incomes are presented in absolute terms or as percentage growth. Although the underlying distributions are informationally equivalent, the framing has a significant effect on participants' allocation choices: those in the absolute framing condition tend to equalize income levels, while those in the relative framing condition tend to equalize percentage gains. Further results suggest that this pattern is not driven by confusion but may stem from framing leading individuals to adopt different underlying notions of inequality, resulting in systematically different fairness standards across conditions.

The second chapter, co-written with Florian Zimmermann, examines belief formation and its connection to actions under model uncertainty. In a financial decision-making experiment, participants estimate company values based on two competing models, knowing that only one of them is correct. They face uncertainty about the correct model, receiving only a signal indicating which model is more likely to hold true. The key manipulation varies the computational complexity of calculating value estimates under each model: in the high-complexity condition participants must compute the estimates under each model themselves, whereas they are provided in the low-complexity condition. While participants' stated beliefs cor-

rectly reflect uncertainty over which model applies, their actions suggest otherwise: particularly under high complexity, participants often neglect model uncertainty, behaving as though one model were certainly correct. Moreover, participants in the high-complexity condition are more confident in the accuracy of their estimates, even though they are less accurate than those in the low-complexity condition. The findings highlight that computational complexity can reduce the extent to which actions correctly reflect model uncertainty while increasing unwarranted confidence in one's judgments.

The third chapter, co-written with Guy Yanay, examines beliefs about others' information processing in political contexts. In a series of incentivized experiments, both lay participants and academic experts predict how others revise their beliefs in response to feedback on politically charged questions. While actual belief updating exhibits no signs of motivated reasoning in the sense of reacting more strongly to pro-party than anti-party news, both groups systematically overestimate the extent to which others engage in such biased updating. These distorted meta-beliefs appear to reflect exaggerated assumptions about partisan extremity and an overapplication of motivated reasoning to others that may reinforce political polarization.

The fourth chapter, co-written with Florian Zimmermann, uses the 2024 U.S. presidential election as a natural experiment to study real-world belief polarization between Democrats and Republicans. Two representative U.S. samples were surveyed shortly before and after the election, allowing for a between-subjects comparison of beliefs over time. The study measures polarization along three dimensions: ideological preferences, factual beliefs, and perceptions of political challenges. Polarization increased following the election, driven particularly by Democrats disappointed by the outcome. Respondents more strongly favored party-aligned categories when indicating their preferred levels of government spending. They also gave more party-aligned responses to politicized factual estimation tasks. Moreover, Democrats were more likely to mention politics as a key societal concern after the election, whereas this became less common among Republicans. The analysis demonstrates how political events can shape both beliefs about the world and the salience of political identity.

Taken together, the four chapters show that beliefs are systematically shaped by how information is presented, by cognitive constraints in processing it, and by political identity—both one's own and that attributed to others. Rather than reflecting purely rational updating, beliefs emerge as context-dependent, influenced by framing, complexity, stereotypes, and identity.

References

Benjamin, Daniel J. 2019. "Errors in probabilistic reasoning and judgment biases." *Handbook of Behavioral Economics: Applications and Foundations 1* 2: 69–186. [1]

Chapter 1

Absolute and Relative Framing of Inequality*

1.1 Introduction

Voters' preferences over redistribution shape policies such as taxes and social transfers in democratic societies. Studying the determinants of these preferences is therefore essential for understanding the feasibility, scope and design of redistributional mechanisms. A growing literature shows that the provision of information about economic inequality can shift fairness perceptions and, in turn, support for redistributive policies (Kuziemko et al. (2015); Hauser and Norton (2017); Cruces, Perez-Truglia, and Tetaz (2013)). Beyond pure informational content, recent work highlights that how inequality is presented matters: differences in framing can shape perceived legitimacy and policy support (Jun et al. (2022), Dietze and Craig (2021), Lorenz, Paetzel, and Tepe (2017)). Building on this idea, this chapter examines a particular kind of framing—whether presenting the same information in absolute versus relative terms affects how individuals evaluate inequality and make distributional choices.

Economic inequality can be measured in absolute or relative terms. Absolute inequality remains unchanged if everyone's income rises by the same fixed amount, whereas relative inequality remains constant if all incomes rise by the same per-

* This research project was funded by the Bonn Graduate School of Economics, the former briq - Institute on Behavior & Inequality (now part of the IZA Institute of Labor Economics) and the CRC TR 244 through project A01. Funded by the Deutsche Forschungsgemeinschaft (DFG, German Research Foundation) under Germany´s Excellence Strategy – EXC 2126/1– 390838866. I thank Florian Zimmermann, Thomas Dohmen, Holger Gerhardt, Andrea Amelio and Anna Schulze Tilling, the participants of the briq Summer School in Behavioral Economics 2021, as well as the audiences of the Bonn Applied Micro Workshop, the 2022 European ESA Meeting in Bologna and the 2023 Bonn-Frankfurt-Mannheim PhD Conference for helpful comments. Proofreading assistance for this chapter was provided by the LLM ChatGPT; the research design, analyses, and all arguments remain solely the work of the authors.

centage. These perspectives often lead to different conclusions. Globally, absolute inequality has increased over the past five decades, while relative inequality has declined (Niño-Zarazúa, Roope, and Tarp (2017)). Locally, similar patterns emerge across many countries—including Germany, Japan, India, and Italy—where absolute inequality has risen while relative inequality has fallen (Bandyopadhyay (2018)).

Experts may interpret both perspectives correctly, but lay people might be swayed by whichever measure is made salient. A news article emphasizing falling relative inequality could lead readers to conclude that inequality is becoming less of a concern and therefore less in need of policy intervention, while one focusing on rising absolute inequality might make them view growing disparities as a justification for stronger redistribution.

This chapter provides causal evidence that such framing can systematically shift fairness preferences. In an online experiment (n = 200), participants allocate additional income between the winner and the loser of a real-effort competition. In the Absolute treatment, allocation options are expressed as absolute income increases; in the Relative treatment, the same options are expressed as percentage increases relative to the contest prizes. Because the contest awards more to the winner than the loser, equalizing absolute gains requires unequal percentage gains, and vice versa.

Three main findings emerge. First, framing strongly shifts distributional choices: participants in the Absolute treatment choose allocations that lead to significantly lower inequality in absolute incomes, as measured by the Gini coefficient, compared to those in the Relative treatment.

Second, the treatment effect persists even among participants who correctly anticipated the payoff implications of their choices in the baseline scenario, indicating it is not driven by confusion. Instead, framing appears to induce a normative shift in what participants perceive as fair.

Third, a complementary survey study (n = 300) replicates the framing effect in a naturalistic context using a potential scenario of U.S. income growth. The scenario features larger percentage growth for poor income quantiles, but greater absolute changes among the rich. Absolute framing increases both the perceived rise in inequality and support for redistributive policies such as raising the minimum wage or top income taxes.

The results presented in this chapter contribute to several related literatures. First, they add to research on how information provision influences attitudes toward inequality. Previous work shows that attitudes toward redistribution are malleable to information provision (Kuziemko et al. (2015); Hauser and Norton (2017)). Other studies examine how framing moderates this effect. A well-documented finding is that framing inequality as a disadvantage to the poor rather than an advantage to the rich increases support for redistributive policies (Chow and Galak (2012); Dietze and Craig (2021); Bruckmüller, Reese, and Martiny (2017)). Simi-

larly, presenting a redistributive mechanism as a minimum income rather than a tax rate raises the preferred level of redistribution in the experimental setting of Lorenz, Paetzel, and Tepe (2017). More broadly, Starmans, Sheskin, and Bloom (2017) argue that people are not averse to inequality per se but tolerate it when they perceive processes as fair, consistent with the present finding that framing shifts the normative standard applied to distributional choices.

Research in economics and psychology has also examined heuristic equity concerns. Exley and Kessler (2024) find that when multiple payoff components are present (e.g., large and small tokens with different monetary values), many individuals equalize only the dimension they can directly control, even when overall payoff equity is achievable. Messick and Schell (1992) report similar findings in a setting where participants allocate either profits or expenses to firms. Andreoni and Bernheim (2009) highlight the pervasive role of the 50-50 norm in redistribution experiments. Such heuristic equity concerns may help explain the treatment effects observed here: participants in the *Absolute* treatment may be inclined to equalize additional incomes, while those in the *Relative* treatment may prefer allocations that equalize percentage growth.

Another line of research recognizes that inequality can be presented in absolute or relative terms. Amiel and F. A. Cowell (1992) study students' inequality preferences in a framing-free environment and find that most align with absolute inequality, followed closely by relative inequality. The lack of a dominant preference suggests that framing effects could meaningfully shape inequality perceptions. Lembregts and Pandelaere (2014) provide further evidence that qualitative fairness assessments of income changes respond to relative versus absolute framing. Their conditions, however, do not preserve informational equivalence across frames. They also study hypothetical scenarios rather than the incentivized choices used in the present work.

Further evidence on framing effects comes from other economic and behavioral settings. Brekke, Konow, and Nyborg (2017) find that framing contributions in a public goods game in percentage terms reduces contributions among participants with low endowments compared to absolute amounts. Glaser, Iliewa, and Weber (2019) show that framing past asset price changes in absolute versus relative terms shapes expectations about future prices. In marketing, Chen and Rao (2007) and Kruger and Vargas (2008) document that people often struggle with percentages, particularly due to difficulties handling base values and a tendency to overemphasize percentage changes. Yan (2019) further show that the clarity of information determines whether people focus on absolute or relative differences, suggesting that framing may affect which fairness notion becomes cognitively salient when evaluating inequality. Taken together, these studies raise the question of whether framing effects in the present chapter are driven by confusion or by shifts in the fairness notions people use. The results indicate that confusion plays only a minor role, as the treatment effect persists even among participants who fully understand the pay-

off implications, pointing instead to framing-induced changes in perceived fairness standards.

The remainder of the chapter is structured as follows. Section 1.2 outlines the experimental design and hypotheses. Sections 1.3 and 1.4 present the results of the experiment and survey study, respectively, and Section 1.5 concludes.

1.2 Experimental Design and Hypotheses

1.2.1 Experimental Design

Many situational factors may influence people's fairness preferences. For example, distributional attitudes can be shaped by whether incomes are determined by luck or skill, individuals' subjective views on the relative importance of various skills, and the presence of equal opportunities, among other considerations that fall outside the scope of this study. Conducting a controlled experiment allows us to abstract from these factors and directly examine the research question at hand: Does presenting wealth changes as absolute amounts or percentage growth affect people's distributional choices?

An experiment designed to answer this question should include the following key elements. First, participants must engage in an incentivized allocation task, distributing real money between other participants. Second, the framing of the additional money must vary exogenously between absolute income and percentage growth. Finally, the design should establish an initial level of inequality, ensuring that equalizing incomes and percentage growth requires different allocations, thereby making the framing manipulation consequential.

The experiment comprised two main stages that incorporated these design features, along with additional tasks aimed at identifying potential channels through which the treatment effect might operate. Figure 1.1 provides an overview of the experimental design. An English translation of the experimental instructions is available in Appendix 1.A.4.

In the first stage, participants were paired and participated in a real-effort competition designed to generate an initial level of inequality. The task required counting the number of zeros in a binary string composed of zeros and ones. The participant who solved the most questions within two minutes won the competition, with ties resolved by a fair coin toss. The winner received a high prize of 5 euros, while the loser received 1 euro. Participants were not informed of the competition's outcome until after the experiment.

The zero-counting task was chosen as an effort-intensive activity in which performance is not obviously determined by inherent ability, as it might be in a math task or IQ quiz. This ensures that the competition starts from a level playing field, rewarding effort rather than skill. By doing so, the design facilitates and standard-

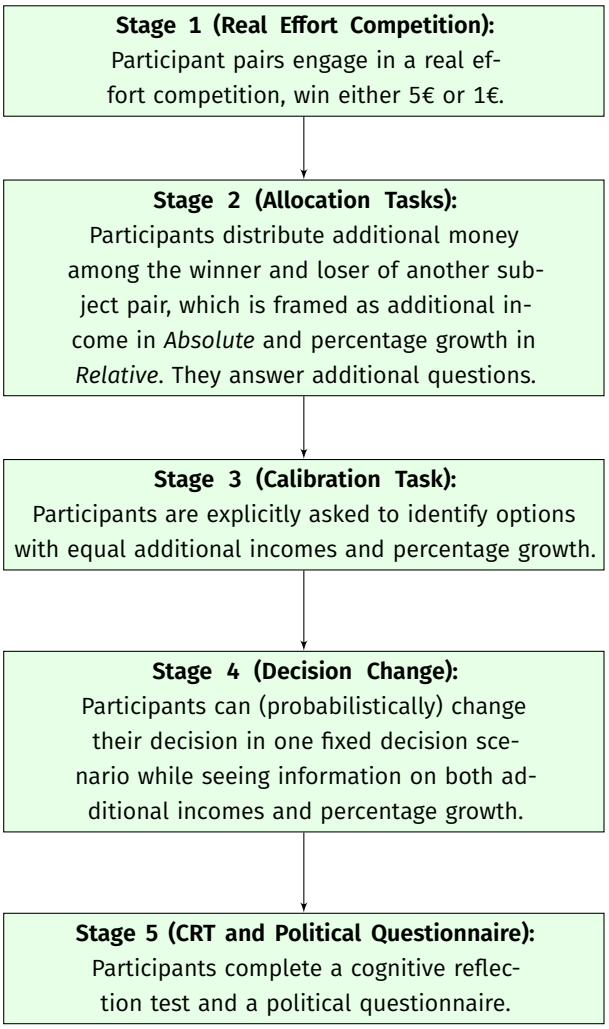

Figure 1.1. Experimental design.

izes acceptance of the allocation of high and low contest prizes, thereby reinforcing the initial level of inequality.

At the beginning of the second stage, participants were informed that all participants could earn additional money on top of their competition prizes. They also learned that they would be responsible for distributing this additional money between the winner and loser of another participant pair. The instructions did not mention or imply any connection between the additional money and the competition.

Before being exposed to any framed information, participants answered two general questions about their views on the additional money. First, they were asked whether they believed the winner should receive a larger share of the additional money (*meritocracy question*). Second, they were asked whether, all else equal, the winner should leave the experiment with more money than the loser (*egalitarian*

question). These questions served to control for participants' fairness attitudes and to identify those inclined toward a more balanced distribution.

Framing effects could be stronger for participants who answered "no" to the *meritocracy question*, as they are likely opposed to allocating all additional money to the competition winner. Similarly, the *egalitarian question* captures the preference for equalizing total payments by awarding most or all of the additional money to the loser. Since these questions were answered before any treatment manipulation, they serve as valid control variables in the analysis.

Afterward, participants were introduced to the details of the allocation stage, in which they distributed additional money between the winner and loser of another pair. Throughout the chapter, I refer to the additional money allocated to each participant as their *income*.

Participants made choices in eight decision scenarios, each presenting seven allocation options. These options were ordered either in ascending or descending order based on the income awarded to the winner, with the order randomized at the subject level. The scenarios varied in both the total amount to be distributed and the specific allocation options available. An overview of all eight decision scenarios is provided in Appendix 1.A.3.

Figure 1.2 illustrates how the options were framed in the two treatment conditions. In the *Absolute* treatment, the wealth change from the additional money was presented as income, whereas in the *Relative* treatment, it was expressed as percentage growth relative to the initial prize. The information on total implied wealth levels was identical across both conditions. Thus, differences in treatment effects cannot be attributed to participants in the *Relative* condition struggling to calculate the total payments implied by various percentage changes.

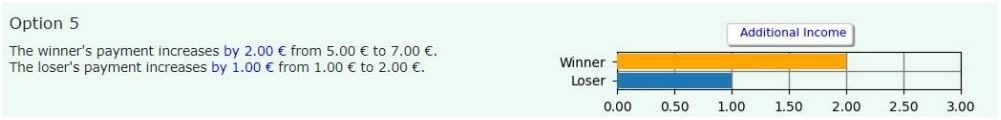

(a) *Absolute* treatment: incomes are highlighted in blue and visualized.

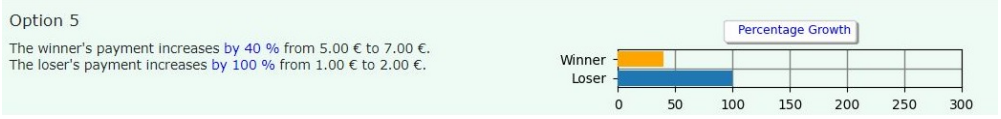

(b) *Relative* treatment: percentage growth is highlighted in blue and visualized.

Figure 1.2. Example for the treatment manipulation.

Immediately after the allocation decisions, participants answered a series of additional questions. One question required them to estimate the wealth consequences of their choices in the *Baseline* decision scenario. Accuracy in this estimation was incentivized and later used to classify participants as *confused* or *non-confused* based on their awareness of the consequences of their actions. The remaining additional

questions asked in the second stage will be introduced in detail as they are analyzed in the next section.

In the third stage of the experiment, participants were asked to identify the options that led to equal additional incomes and equal percentage growth in the *Baseline* decision scenario. This assessment tested whether participants could correctly recognize the key options of interest in both treatment conditions. A bonus of 1 euro was awarded for correct answers.

In the fourth stage, participants had the opportunity to probabilistically revise their decisions from the *Inbetween* decision scenario in the first stage. In this stage, both treatments displayed information on additional incomes and percentage growth. Participants were reminded of their initial decision to minimize noise and ensure that any changes reflected deliberate choices rather than random preference fluctuations. Revisions were implemented with a 20% probability. This stage aimed to provide within-subject evidence of the framing effect.

In the fifth and final stage, participants completed a six-item cognitive reflection/IQ test and a short political questionnaire on redistribution in Germany. These data were collected as control variables and to assess whether treatment effects were mediated by cognitive reflection. The test was a variant of the cognitive reflection test, designed to include answers that could not be easily found online. A time limit was imposed to further discourage searching for answers. In the political questionnaire, participants indicated their level of agreement with four statements about redistributive policies in Germany. All questions are provided in Appendix 1.A.4.

1.2.2 Hypotheses

The first two hypotheses pertain to participants' allocation decisions in the second stage of the experiment, focusing on the between-subjects comparison of distributional behavior. The expectation is that participants in the *Absolute* condition will tend to equalize additional incomes, whereas those in the *Relative* condition will favor options that equalize percentage growth.

Hypothesis 1.1. Participants in the *Absolute* condition make allocation decisions that result in more equal additional incomes than those in the *Relative* condition. Consequently, income inequality between the winner and loser is lower in the *Absolute* condition than in the *Relative* condition.

Hypothesis 1.2. Participants in the *Relative* condition make allocation decisions that result in more equal percentage growth than those in the *Absolute* condition.

The third hypothesis concerns behavior in the fourth stage. In this stage, participants have the opportunity to revise a decision from the first stage while being presented with information on both additional incomes and percentage growth. I hypothesize that this opportunity leads to convergence in behavior across both treatments.

Hypothesis 1.3. The opportunity to revise decisions while viewing information on both additional incomes and percentage growth leads to convergence in allocation decisions across the two treatments.

1.3 Results

1.3.1 Procedures

The experiment was conducted online in March 2022 using the subject pool of the BonnEconLab. Participants were primarily University of Bonn students from various fields of study. Participants were recruited via hroot (Bock, Baetge, and Nicklisch (2014)), and the experiment was implemented in oTree (Chen, Schonger, and Wickens (2016)). A total of 200 participants—100 per treatment—completed the study.[1] On average, participants earned 11.67 euros and took 43 minutes to complete the experiment. All sessions were supervised via Zoom, and participants were required to keep their cameras on throughout the study.

1.3.2 Statistical measures

To examine how absolute and relative framing affect the inequality implied by participants' allocation decisions, I use the Gini coefficient, the most widely used statistical measure of inequality in economics. Let y_1 denote the winner's income and y_2 the loser's income, excluding the high and low competition prizes. That is, y_1 and y_2 are determined solely by the allocation choice of the decision-maker assigned to a participant pair. For a society consisting of two people, namely the winner and loser of the competition, the Gini coefficient of additional incomes is defined as follows:

$$\text{Additional Income Gini} = \frac{|y_1 - y_2|}{y_1 + y_2} \tag{1.1}$$

Note that the Gini coefficient simply measures how close the winner's and loser's incomes are. To quantify the proximity of chosen percentage growth values, we can define an analogous measure. Let p_i denote the competition prizes. Then, let $g_i = \frac{y_i}{p_i}$ represent the percentage growth values.

$$\text{Additional Percentage Gini} = \frac{|g_1 - g_2|}{g_1 + g_2} \tag{1.2}$$

The Gini coefficient of incomes equals zero when incomes are fully equalized and 1 when either the winner or the loser receives all of the additional money. The Gini coefficient of percentage growth has analogous properties, equaling zero when percentage growth is equalized. Using alternative measures, such as the standard deviations of incomes and percentage growth, yields qualitatively identical results.

1. Six participants with partially missing demographic data were excluded from the analysis.

1.3.3 Between-subject evidence

1.3.3.1 Main Result

In this section, I compare the distributional behavior of participants in the *Relative* and *Absolute* treatment conditions across the eight second-stage allocation tasks. The results support Hypotheses 1.1 and 1.2.

Figure 1.3 summarizes the main result. The chosen incomes for the winner and loser are closer in the *Absolute* condition than in the *Relative* condition, as measured by the Gini coefficient of additional incomes ($p < 0.01$).

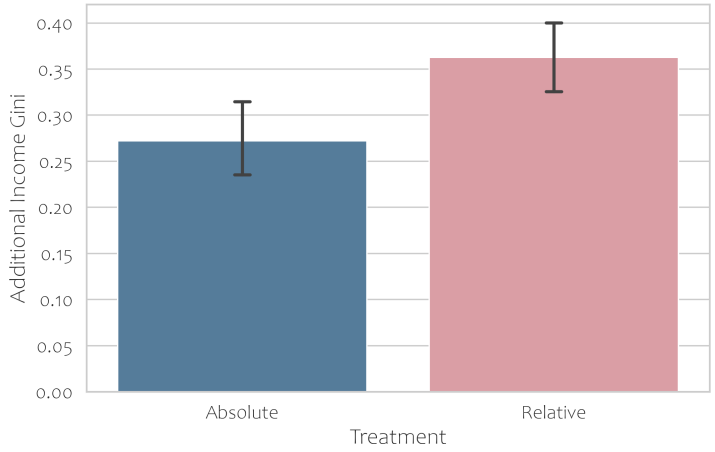

Figure 1.3. Mean additional income Gini in both treatments. *This figure plots the average additional income Gini as defined in Equation 1.1 for both treatment conditions. The figure is based on the main experiment.*

Table 1.1 presents results from an OLS regression, confirming the main finding. In column (1), the Additional Income Gini is regressed on a constant and a dummy variable for the *Absolute* treatment. Column (2) adds controls for subject characteristics, while column (3) further includes controls for participants' responses to the meritocracy and egalitarian questions, as well as their answers to the political questionnaire on redistribution in Germany. Across all specifications, the treatment effect is negative and highly significant.

Examining each decision scenario separately, I find a significant treatment difference in the Additional Income Gini at the 5% level in seven of the eight comparisons.

Figure 1.A.1 and Table 1.A.1 present the corresponding analysis for the Gini coefficient of percentage growth. Consistently, percentage growth values are more equalized in the *Relative* condition than in the *Absolute* condition ($p < 0.01$).

The treatment difference in income inequality translates directionally into a difference in the total earnings of the contest loser. Table 1.2 shows that losers take

Table 1.1. Income Gini Coefficient

	Additional Income Gini		
	(1)	(2)	(3)
Constant	0.363***	0.255***	0.100
	(0.018)	(0.083)	(0.084)
Absolute	−0.091***	−0.091***	−0.106***
	(0.026)	(0.025)	(0.023)
Participant characteristics	✗	✓	✓
Participant answers	✗	✗	✓
R^2	0.029	0.041	0.112
Observations	1552	1552	1552

This table presents OLS regression results for the outcome variable *Additional Income Gini* as defined in Equation 1.1. The table is based on the main experiment. *Absolute* is a treatment dummy for the Absolute treatment. Stars highlight significant differences from 0 with * for $p < 0.10$, ** for $p < 0.05$, *** for $p < 0.01$. Participant-clustered standard errors in parentheses. Participant characteristics controls include age, male dummy, education level and a dummy for an above median CRT score. Participant answer controls include answers to the questions on meritocracy and egalitarianism as well as answers to the political questionnaire.

home more money in the *Absolute* condition than in the *Relative* condition. While this difference is not significant without controls, it becomes highly significant when accounting for the decision-maker's attitudes toward the additional money, as measured by the meritocracy and egalitarian questions.

Overall, the findings so far can be summarized as follows.

Result 1.1. Additional incomes are more equal in **Absolute** than in **Relative**, while percentage growth is more equal in **Relative** than in **Absolute**. When controlling for the planner's general fairness attitudes, contest losers earn more if the planner is in the **Absolute** treatment.

1.3.3.2 Dummy Analysis

The treatment variation may influence the salience of different allocation options. For example, in the *Absolute* condition, the option that equalizes additional incomes may be particularly salient, whereas in the *Relative* condition, participants may instead focus on the option that equalizes percentage growth. Figure 1.4 examines how frequently these options are selected in their respective treatments.

Panel (a) presents the analysis for the options that equalize percentage growth. These options are selected more frequently in the *Relative* condition ($p < 0.01$) and appear salient in this condition, as they are chosen at a rate marginally exceeding

Table 1.2. Implied Total Wealth of the Loser

	Total Wealth of the Loser		
	(1)	(2)	(3)
Constant	2.516***	2.742***	3.334***
	(0.053)	(0.217)	(0.191)
Absolute	0.103	0.101	0.158***
	(0.072)	(0.071)	(0.057)
Participant characteristics	✗	✓	✓
Participant answers	✗	✗	✓
R^2	0.005	0.028	0.198
Observations	1552	1552	1552

This table presents OLS regression results for the outcome variable *Total Wealth of the Loser*, which is defined as the sum of the contest prize of 1 euro and the additional income awarded to the loser of the contest. The table is based on the main experiment. *Absolute* is a treatment dummy for the Absolute treatment. Stars highlight significant differences from 0 with * for $p < 0.10$, ** for $p < 0.05$, *** for $p < 0.01$. Participant-clustered standard errors in parentheses. Participant characteristics controls include age, male dummy, education level and a dummy for an above median CRT score. Participant answer controls include answers to the questions on meritocracy and egalitarianism as well as answers to the political questionnaire.

the random choice probability of $\frac{1}{7}$ ($p < 0.1$). In the *Absolute* condition, this option is selected significantly less often than predicted by random choice.

Panel (b) reports how frequently participants in both treatments select the option that equalizes additional incomes. The treatment comparison indicates that this option is chosen somewhat more often in the *Absolute* condition than in the *Relative* condition ($p < 0.1$). However, it remains salient in both treatments. Since participants in both conditions received information on the total payments implied for the winner and loser, those in the *Relative* condition could still easily identify the option that equalizes additional incomes. This may explain why the option is salient across both treatments.

The results from the graphical analysis are confirmed by the regressions in Table 1.3. In the regressions without controls (columns (1) and (3)), participants in the *Absolute* condition are 7 percentage points more likely to choose the option that equalizes additional incomes. Meanwhile, participants in the *Relative* condition are 13 percentage points more likely to select the option that equalizes percentage growth.

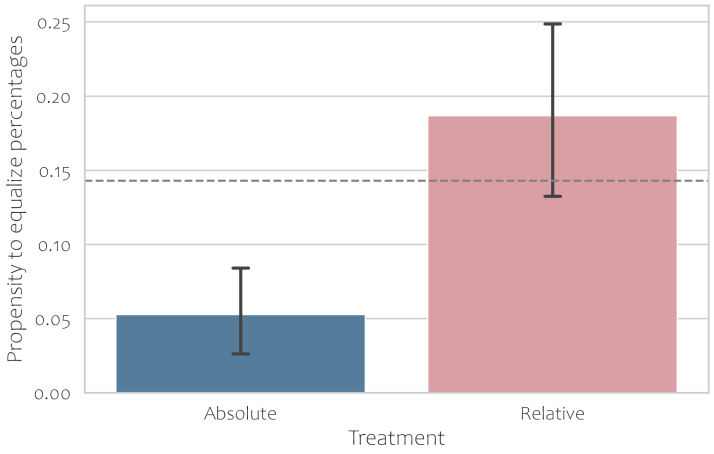

(a) Equal percentage growth dummy.

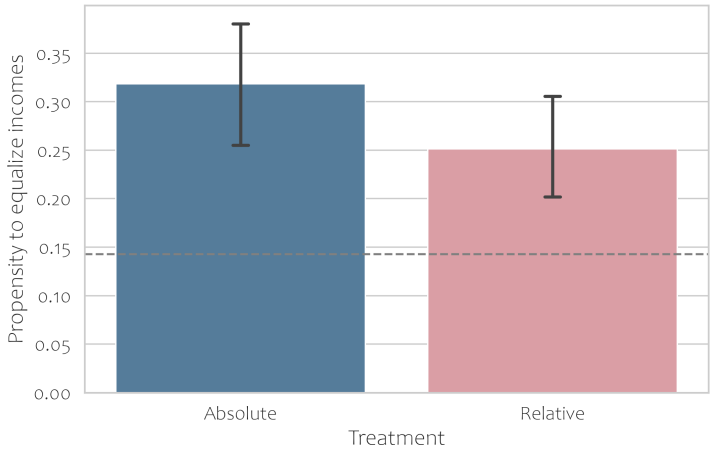

(b) Equal income dummy.

Figure 1.4. Options with most equal incomes and percentage growth. *Panel (a) plots the rate with which the option with (most) equal percentage growth is selected in each treatment. Panel (b) plots the rate with which the option with (most) equal additional incomes is selected in each treatment. The figure is based on the main experiment.*

This analysis suggests that treatment differences in income inequality are primarily driven by the frequent equalization of percentage growth in the *Relative* condition, which is rarely observed in the *Absolute* framing.

Result 1.2. Options that equalize incomes are chosen more frequently in **Absolute** than in **Relative**, while options that equalize percentage growth are chosen more frequently in **Relative** than in **Absolute**. The latter effect is stronger.

Table 1.3. OLS Regressions for Choice Dummies

	Chose equal incomes		Chose equal perc. growth	
	(1)	(2)	(3)	(4)
Constant	0.251***	0.643***	0.187***	−0.039
	(0.023)	(0.103)	(0.027)	(0.096)
Absolute	0.067*	0.099***	−0.134***	−0.149***
	(0.038)	(0.033)	(0.030)	(0.029)
Controls	✗	✓	✗	✓
R^2	0.005	0.096	0.042	0.099
Observations	1552	1552	1552	1552

This table presents OLS regression results. The table is based on the main experiment. In columns (1) and (2), the outcome is a dummy variable that is 1 if the allocation with (most) equal additional incomes was selected. In columns (3) and (4), the outcome is a dummy variable that is 1 if the allocation with (most) equal percentage growth was selected. *Absolute* is a treatment dummy for the Absolute treatment. Stars highlight significant differences from 0 with * for $p < 0.10$, ** for $p < 0.05$, *** for $p < 0.01$. Standard errors are clustered on participant level and in parentheses. Controls consist of age, gender, CRT performance, education level, answers to meritocracy and egalitarian questions as well as the political questionnaire.

1.3.3.3 Classification of Fairness Types

An alternative approach to analyzing the data is to examine how closely participants' actions align with various fairness ideals. In the following analysis, I classify participants into four distinct types:

Income equalizers aim to preserve absolute equality and select the option that induces the most equal additional incomes. *Percentage equalizers* seek to maintain relative equality and choose the option with equal percentage growth values. *Meritocrats* prioritize rewarding high effort and allocate as much of the additional money as possible to the competition winner. *Egalitarians* aim to undo the inequality introduced by the unequal allocation of prizes and give as much money as possible to the competition loser.[2]

Participants are classified based on the type that best matches their behavior across the eight allocation tasks in the second stage. To determine this, I calculate the sum of absolute differences between a participant's chosen allocations and those corresponding to each type. Participants are assigned to the type that minimizes

2. The amount of additional money was never large enough to make the loser's total payment exceed that of the winner.

this sum. If multiple types yield the same minimum, the participant is classified as *Other*.

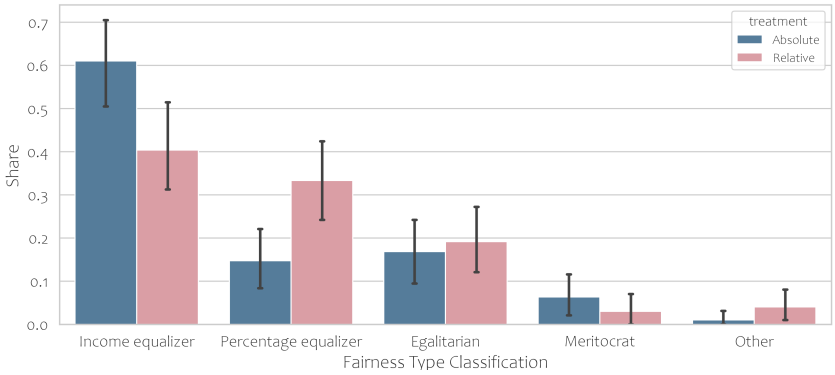

Figure 1.5. Classification of fairness types. *This figure plots the results of classifying participants into four distinct fairness types. Income equalizers aim for equality of incomes, percentage equalizers for equality of percentage growth, meritocrats reward the contest winner, and egalitarians prioritize compensating the loser. Participants are classified by the type that minimizes the sum of absolute differences between their allocations and each type's allocations. If there's a tie, they are classified as Other. The figure is based on the main experiment.*

Figure 1.5 presents the results, revealing significant treatment differences in classified types. Participants in the *Absolute* condition are significantly more likely to be classified as *income equalizers*, while the *Relative* condition includes a larger share of *percentage equalizers* (both $p < 0.01$). For other types, no significant treatment differences are observed.

Notably, the combined share of *income equalizers* and *percentage equalizers* accounts for approximately 75% of participants in both treatments, though their relative distribution differs. This suggests that a certain group of participants—those generally inclined toward balanced and fair allocations—may be particularly susceptible to framing effects, whereas others, such as *egalitarians* and *meritocrats*, may be less affected.

1.3.3.4 Heterogeneity by Meritocracy and Egalitarian Attitudes

Participants with a meritocratic view of the additional money may consistently allocate most or all of it to the contest winner, regardless of framing. In contrast, those without a meritocratic view may prefer a more balanced distribution and, therefore, be more influenced by the treatment manipulation. Similarly, participants with egalitarian motives may seek to undo the initial inequality introduced by the competition, while those without such a view may aim for a more even distribution of the additional money between recipients.

Table 1.A.2 splits the sample based on participants' responses to the meritocracy and egalitarian questions. The statistical power of this analysis is limited. However, the treatment effect appears strongest among participants who hold neither a meritocratic view of the additional money nor an egalitarian motive, as shown in column (3). The second-largest point estimate is observed for participants who lack a meritocratic view but report egalitarian motives, as seen in column (4). For both groups, the estimated treatment coefficient is negative and highly significant ($p < 0.01$). In contrast, for participants with a meritocratic view of the additional money (columns (1) and (2)), the treatment coefficient estimates are smaller and not significant at the 10% level.

1.3.3.5 Within-subject evidence

In the fourth stage of the study, participants could revisit one of their decisions (the *Inbetween* scenario in Appendix 1.A.3). This time, they received information on both the implied absolute incomes and percentage growths of each option. Additionally, they were reminded of their previous allocation and given the opportunity to probabilistically revise their decision from the second stage of the experiment.

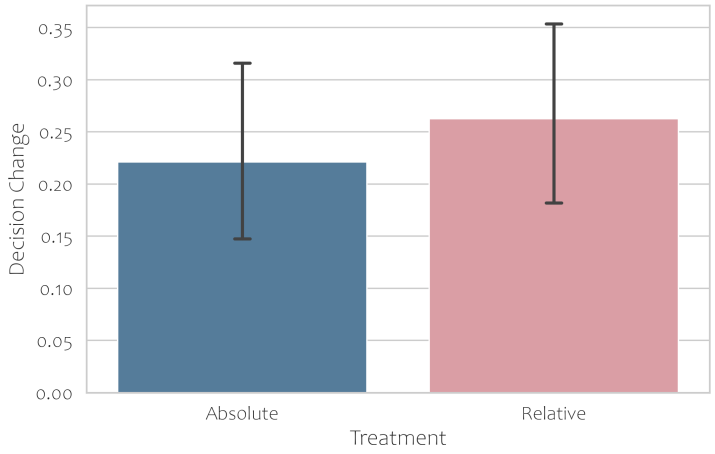

Figure 1.6. Propensity to change decision. *This figure plots the fraction of subjects in each treatment who decide to change their decision in the Inbetween scenario upon being shown both frames simultaneously. The figure is based on the main experiment.*

Figure 1.6 illustrates the propensity to revise decisions after observing both frames in the respective treatment conditions. In both conditions, a significant fraction of participants chooses to make changes. While this fraction is directionally higher in the *Relative* treatment, the difference is not statistically significant at the 10% level.

Figure 1.7 illustrates the impact of the opportunity to revise decisions on income inequality in the *Inbetween* decision scenario. Before this opportunity, the Ad-

ditional Income Gini is significantly lower in *Absolute* than in *Relative* ($p < 0.05$). Allowing participants to revise their decisions directionally increases the Additional Income Gini in *Absolute* ($p < 0.1$) but significantly decreases it in *Relative* ($p < 0.01$). As a result, the treatment difference is no longer significant at the 10% level ($p = 0.34$). Table 1.A.3 confirms these findings using a difference-in-differences regression analysis.

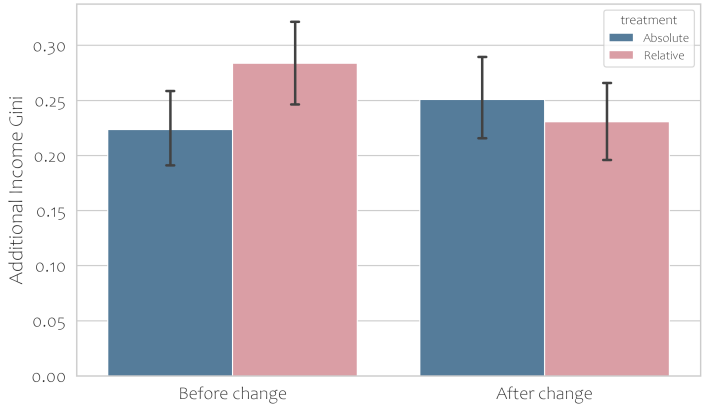

Figure 1.7. Impact of decision changes on additional income Gini. *This figure plots the mean additional income Gini in the Inbetween decision scenario before and after participants were shown both frames simultaneously and had the opportunity to change their decision. The figure is based on the main experiment.*

Result 1.3. The within-subject evidence supports the presence of a framing effect. When exposed to both frames simultaneously, a significant fraction of participants revise their decisions, eliminating the significant treatment difference.

1.3.4 Potential Effect Channels

Framing effects typically influence how people perceive a decision. After completing all allocation tasks, participants were asked whether they had the impression that the competition winner, the loser, or neither was systematically financially advantaged in most of the available options. The instructions remained intentionally vague about the meaning of "financially advantaged." This question was designed to test whether framing influenced perceptions of fairness. While percentage growth values are systematically higher for the loser than for the winner, this pattern does not hold for additional incomes. As a result, participants in the *Relative* treatment may be more likely to perceive the loser as systematically favored.

Figure 1.8 shows a clear treatment difference—the proportion of participants stating that the loser was advantaged is significantly higher in the *Relative* condition than in the *Absolute* condition. The main treatment effect and the difference observed in Figure 1.8 could arise from two primary sources.

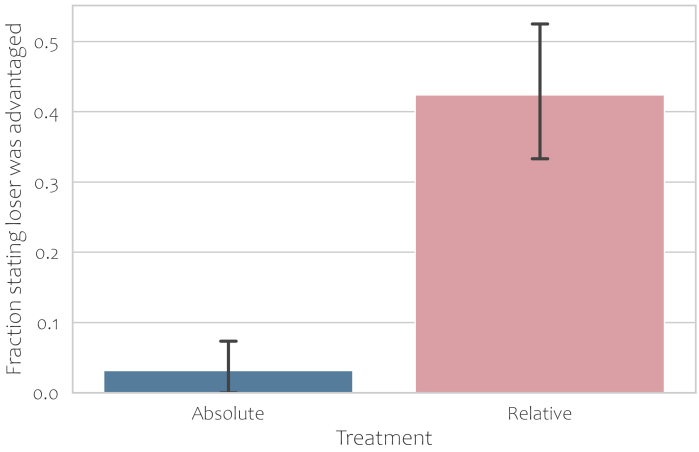

Figure 1.8. Perceived advantage for the loser of the competition. *This figure plots the fraction of participants in each treatment stating that the loser was advantaged in most available allocation options. The figure is based on the main experiment.*

First, framing might confuse participants about the relationship between percentages and additional incomes. For example, participants in the *Relative* condition might mistakenly believe that selecting an option with equal bar lengths equalizes additional incomes, when it actually equalizes percentage growth.[3] Participants in both treatments may apply the same fairness concept (e.g., absolute equality) to their choices, but confusion could lead them to select different options. Past research has shown that people often struggle with percentages (Chen and Rao (2007); Kruger and Vargas (2008)).

Second, framing may influence the salience of different fairness ideals. Participants in the *Relative* treatment may recognize that they are viewing percentage-based information but could be inclined to perceive options with equal percentages as fairest. Conversely, participants in the *Absolute* treatment may consider options with equal incomes to be fairest. Research has shown that fairness attitudes reflecting both absolute and relative inequality are commonly observed (Amiel and F. A. Cowell (1992); Amiel and F. Cowell (1999)). Moreover, Lombardi, Fehr, and Hare (2018) demonstrate that framing effects operate through attention management. Thus, presenting income changes in either relative or absolute terms may nudge individuals toward adopting the corresponding concept of inequality.

In the remainder of this section, I examine the extent to which these channels explain the observed behavior.

3. Another possible misunderstanding in the *Relative* condition is that participants may interpret a percentage value (e.g., 40% for the loser) as indicating that the loser receives 40% of the additional sum. However, this type of confusion is unlikely to play a major role, as every decision scenario included multiple options where the loser's percentage growth exceeded 100%.

1.3.4.1 Does confusion drive the framing effect?

The experimental design includes several features to examine the first channel, which concerns confusion between additional incomes and percentage growth.

It is important to distinguish between whether participants are confused about percentages and incomes and whether this confusion drives the treatment effect. I first examine whether confusion influences the treatment effect before assessing the extent to which participants exhibit signs of confusion in more detail.

After completing all allocation tasks in the second stage, participants were asked to estimate the wealth implications of their choices in the *Baseline* scenario. Specifically, they provided estimates of the total payments received by the winner and the loser, including the competition prize and the allocated additional money or percentage growth.

I classify participants who provided the exact correct estimate as *non-confused* and those who gave any incorrect estimate as *confused*. Notably, this classification of *non-confused* participants is conservative, as it requires them not only to understand the consequences of their actions at the time of decision-making but also to recall these consequences minutes later.[4]

In the *Absolute* condition, 61 participants are classified as *non-confused* and 39 as *confused*, whereas in the *Relative* condition, 37 participants are classified as *non-confused* and 63 as *confused*.

Column (1) of Table 1.4 examines how the confusion variable interacts with the treatment effect by including a dummy for *confused* participants. The treatment effect remains significant among *non-confused* participants. The point estimate for the interaction suggests that confusion, as measured above, has little to no impact on the observed framing effect.[5]

The remaining columns of Table 1.4 explore alternative classifications of participants based on their susceptibility to confusion from framing.

In column (2), the sample is divided based on incorrect answers to control questions. Participants who made at least one mistake on their first attempt are classified as *Low attention*. This applies to 56 out of 100 participants in *Absolute* and 52 out of 100 in *Relative*, effectively creating a median split. If anything, the treatment effect appears stronger among high-attention participants.

In column (3), the sample is divided using a median split based on cognitive reflection. The treatment effect remains present and is marginally significant among participants with above-median cognitive reflection. Although the interaction term

4. Additionally, this definition of confusion is broad, as only participants who are fully aware of the consequences of their actions are classified as *non-confused*. Thus, it also applies to participants experiencing the alternative type of confusion discussed in the previous footnote.

5. Classifying participants with above-median estimation errors as *confused* and others as *non-confused* yields the same conclusion.

Table 1.4. Confusion and the Treatment Effect

	Additional Income Gini			
	(1)	(2)	(3)	(4)
Constant	0.304***	0.349***	0.337***	0.351***
	(0.027)	(0.022)	(0.021)	(0.021)
Absolute	−0.071**	−0.115***	−0.064*	−0.089***
	(0.035)	(0.029)	(0.034)	(0.030)
Confused	0.092***			
	(0.035)			
Absolute × Confused	0.002			
	(0.051)			
Low attention		0.029		
		(0.036)		
Absolute × Low att.		0.062		
		(0.053)		
Low cognitive reflection			0.068*	
			(0.037)	
Absolute × Low cr.			−0.070	
			(0.052)	
Calibration mistake				0.054
				(0.038)
Absolute × Calib. mist.				−0.017
				(0.055)
R^2	0.057	0.043	0.036	0.033
Observations	1552	1552	1552	1552

This table presents OLS regression results for the outcome variable *Additional Income Gini* as defined in Equation 1.1. The table is based on the main experiment. *Absolute* is a treatment dummy for the Absolute treatment. *Confused* is a dummy for giving an incorrect answer when estimating the wealth implications of one's Baseline choice. *Low attention* is a dummy for making at least one mistake in the comprehension checks. *Low cognitive reflection* is a dummy for scoring below median in the cognitive reflection / IQ test. *Calibration mistake* is a dummy for making at least one mistake in the calibration tasks. Stars highlight significant differences from 0 with * for $p < 0.10$, ** for $p < 0.05$, *** for $p < 0.01$. Participant-clustered standard errors in parentheses.

is insignificant, its direction suggests that treatment effects may be stronger among participants with lower cognitive reflection.

Column (4) compares participants who answered the calibration tasks correctly to those who made at least one mistake. The treatment effect is significant among participants who passed the calibration checks. Although the interaction term is insignificant, it should not be overinterpreted, as the vast majority of participants

in both treatments answered both calibration questions correctly (79 in *Relative* and 73 in *Absolute*).

Another way to assess whether confusion influences the treatment effect is to examine within-subject decision changes. If confusion were driving the effect, we would expect primarily confused participants to revise their decisions. Table 1.A.4 analyzes how decision changes differ between confused and non-confused participants. Confused participants revise their decisions more frequently, a difference that is marginally significant. They also make slightly larger adjustments, though the difference between confused and non-confused participants is not statistically significant. Still, a significant number of substantial decision changes occur in both treatment conditions even among non-confused participants.

Although confusion does not appear to drive the treatment effect, it is still present among some participants in the *Relative* treatment. Based on the accuracy of participants' estimates, 39 participants in the *Absolute* treatment were classified as *confused*, compared to 63 in the *Relative* treatment.

In contrast, performance in the calibration exercise presents a different picture. In this task, all participants were explicitly asked to identify options with equal additional incomes and equal percentage growth. As shown in Figure 1.A.2, calibration task performance does not differ significantly between the two conditions. While this does not necessarily imply that participants fully understood the distinction between equal incomes and percentage growth during the allocation task, it suggests that most participants, in principle, have no difficulty differentiating between them.

Finally, after the calibration exercise, participants were explicitly asked whether they had recognized the difference between equalizing incomes and percentage growth while completing the allocation tasks. This question aimed to measure self-reported confusion. The results indicate that 87% of participants in the *Absolute* treatment and 94% in the *Relative* treatment reported no confusion ($p = 0.092$). These findings do not suggest that participants in the *Relative* condition commonly felt misled by the framing into confusing incomes with percentages.

Overall, the evidence suggests that while some participants exhibit signs of confusion about incomes and percentages, this is unlikely to be the primary driver of the observed treatment effect.

Result 1.4. While framing may confuse some participants about percentages and incomes, the treatment effect persists even among non-confused participants.

1.3.4.2 Does the framing affect fairness ideals?

The second way the treatment effect may operate is by nudging participants to apply different concepts of inequality in their choices.

This second channel can be examined by analyzing how participants perceive changes in inequality resulting from different allocations of the additional money, relative to the initial distribution of 5 and 1 euros. Participants in both treatments

rated the inequality changes in three scenarios, each of which was also available during the allocation stage. These ratings were elicited at the end of stage 2.

In Scenario 1, additional incomes were equalized, leading to greater percentage growth for the competition loser. In Scenario 2, percentage growth was equalized, resulting in a larger additional income for the loser. Scenario 3 was an intermediate case, where the winner received a greater additional income, but the loser experienced higher percentage growth. The scenarios were presented using the same framing as in the allocation task—percentage growth in the *Relative* condition and additional incomes in the *Absolute* condition.

Responses were recorded on a 7-point Likert scale, ranging from "*Inequality strongly decreases*" through "*Inequality remains unchanged*" to "*Inequality strongly increases*".

A participant whose views fully align with the concept of *absolute inequality* would be expected to state that inequality remains unchanged in scenario 1, decreases in scenario 2, and increases in scenario 3. Similarly, a participant adhering to the notion of *relative inequality* would report that inequality decreases in scenario 1, remains unchanged in scenario 2, and decreases in scenario 3. Based on these patterns, I classify participants into three groups: those whose responses align with *absolute inequality*, those whose responses align with *relative inequality*, and those whose answers do not clearly correspond to either concept.

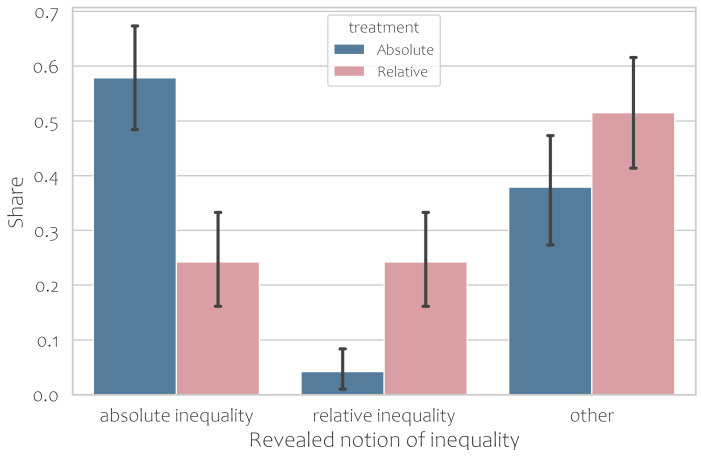

Figure 1.9. Attitudes towards changes in inequality. *This figure plots the share of participants whose answers to the inequality rating questions fully align with the notion of absolute inequality, relative inequality, or neither notion. The figure is based on the main experiment.*

Figure 1.9 illustrates the results of this classification. The concept of *absolute inequality* is significantly more prevalent in the *Absolute* condition, while *relative inequality* is more frequently observed in the *Relative* condition (both $p < 0.01$).

In principle, the observed differences could stem from confusion about percentage growth and additional incomes. However, the same pattern holds within the sub-sample of *non-confused* participants who correctly estimated the wealth implications of their *Baseline* choices (see Figure 1.A.3). This suggests that participants' views on inequality are indeed influenced by the different presentations and that the observed differences are not merely a result of confusion.

Overall, there is evidence that framing affects participants' applied notions of inequality.

Result 1.5. There is suggestive evidence that the treatment effect operates through nudging people toward different concepts of inequality.

1.4 Survey Study

1.4.1 Design

To validate the results in a more natural setting, I conducted an additional survey study.

In the first part of the survey, participants were presented with a hypothetical scenario depicting income growth across the five quintiles of the U.S. household income distribution. They viewed a graph that illustrated income growth in either absolute or relative terms. The scenario was designed so that relative income growth was higher for lower-earning households, while absolute income growth was higher for higher-earning groups. The graphs are shown in Figure 1.10.

After reviewing the figure and answering comprehension questions, participants responded to several questions inspired by Kuziemko et al. (2015). First, they rated the change in inequality induced by the scenario on a Likert scale from 1 ("*Inequality strongly decreases*") to 7 ("*Inequality strongly increases*"). Second, they indicated their agreement with the statement that the depicted development is alarming, using the same 1-to-7 Likert scale. Finally, they answered two policy questions on whether the minimum wage and the income tax for high earners should be increased if the projection were to materialize, with agreement again elicited on a 1-to-7 Likert scale.

In the second part of the survey, participants selected one of six available scenarios for household income growth. The scenarios were presented in the respective treatment framing and ranged from equal absolute income growth to equal relative income growth. The level of average absolute income growth remained constant across all scenarios, mirroring the fixed additional sum in the main study. All scenarios in the *Absolute* framing are reproduced in Figure 1.A.4.

At the end of the survey, participants in both treatments were asked to identify the group with the largest *absolute* income growth in the scenario presented in Part 1. This question aimed to detect whether participants in the *Relative* treatment confused relative growth with absolute growth.

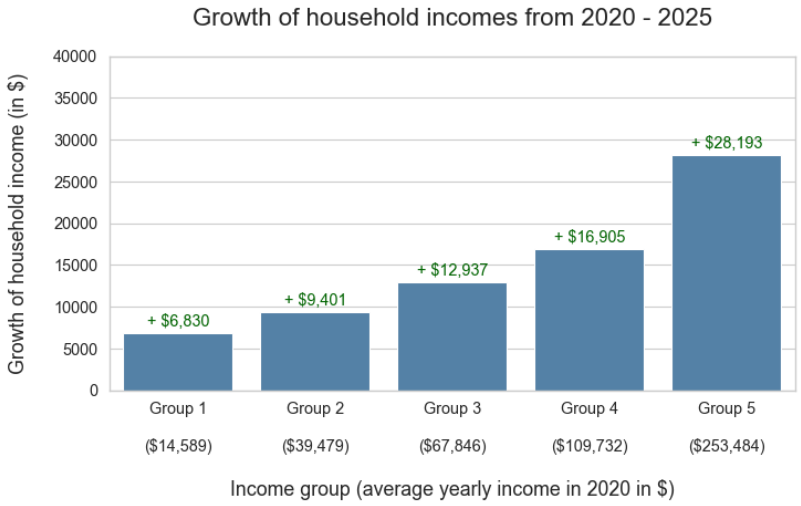

(a) Figure in the Absolute condition.

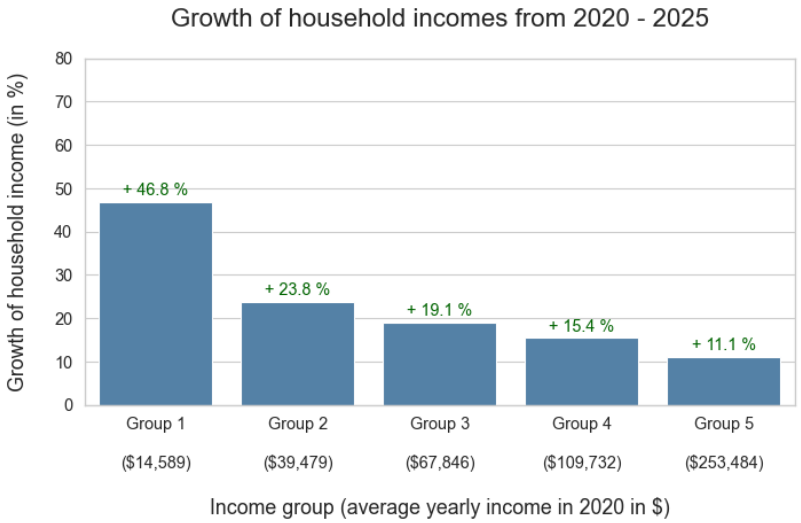

(b) Figure in the Relative condition.

Figure 1.10. Treatment manipulation in the survey study.

The survey was conducted online using Qualtrics and Prolific in September 2022. A total of 300 participants—150 per treatment condition—completed the study.[6] The median completion time was 6 minutes and 30 seconds, and participants received a flat fee of $1.40, corresponding to an hourly wage of approximately $13—above the average reward rate on Prolific.

6. Six participants were excluded due to partially missing demographic information.

1.4.2 Main Results

Table 1.5 presents the results for the first two questions, which assess whether inequality increased and whether the depicted development was perceived as alarming. Mean Likert scale responses to both questions are approximately 1.4 points higher in the *Absolute* treatment than in the *Relative* treatment, indicating that framing significantly influences participants' perceptions of the scenario. The share of participants stating that inequality increases is 42% in the *Relative* condition and 82% in the *Absolute* condition ($p < 0.001$). Similarly, the share of participants reporting that the development is alarming is 43% in the *Relative* condition and 81% in the *Absolute* condition ($p < 0.001$). Table 1.A.5 confirms these findings by incorporating controls for responses to a political questionnaire and baseline demographic variables.

Table 1.5. Perceived Development of Inequality

	Change in Inequality	Alarming Development
Constant	4.223***	4.047**
	(0.118)	(0.139)
Absolute	1.380***	1.453***
	(0.167)	(0.197)
R^2	0.187	0.155
Observations	294	294

This table presents OLS regression results. In column (1) the outcome is *Change in Inequality*, which is the change in inequality induced by the presented scenario on a Likert scale from 1 (*"Inequality strongly decreases"*) to 7 (*"Inequality strongly increases"*). In column (2), the outcome is *Alarming Development*, which is the agreement with the statement that the depicted development is alarming, using a 1-to-7 Likert scale. *Absolute* is a treatment dummy for the Absolute treatment. Stars highlight significant differences from 0 with * for $p < 0.10$, ** for $p < 0.05$, *** for $p < 0.01$. Participant-clustered standard errors in parentheses.

Table 1.6 analyzes responses to the policy questions. Agreement with increasing the minimum wage is examined in column (1), support for raising the income tax on high earners in column (2), and column (3) pools both questions, incorporating question fixed effects and clustering standard errors at the participant level. Support for both policies increases by approximately half a Likert scale point when shifting from the *Relative* to the *Absolute* framing condition.

When controls are added in Table 1.A.6, the point estimates decrease and remain only marginally significant ($p < 0.1$), whereas the pooled regression estimate remains significant at the 5% level. Given that the scenario in this study was purely hypothetical, I interpret these findings as suggestive evidence that support for redis-

tributive policies is influenced by the relative versus absolute framing of changes in the income distribution.

Table 1.6. Policy Support

	Minimum Wage	Income Tax	Pooled Support
Constant	5.642***	5.459***	5.472***
	(0.120)	(0.128)	(0.128)
Absolute	0.427**	0.479***	0.453***
	(0.171)	(0.181)	(0.154)
Question FE	✗	✗	✓
R^2	0.018	0.020	0.021
Observations	294	294	588

This table presents OLS regression results. The table is based on the survey study. In column (1), the outcome is agreement with increasing the *Minimum Wage*. In column (2), the outcome is agreement with increasing the *Income Tax* for the highest income group. Agreement was measured on a Likert scale from 1 (strongly disagree) to 7 (strongly agree). Column (3) pools observations from (1) and (2). *Absolute* is a treatment dummy for the Absolute treatment. Stars highlight significant differences from 0 with * for $p < 0.10$, ** for $p < 0.05$, *** for $p < 0.01$. Participant-clustered standard errors in parentheses.

Finally, Table 1.7 analyzes participants' preferred scenarios in the second part of the survey. In column (1), the dependent variable is the selected scenario, where higher values correspond to more equal distributions of absolute income growth. Consistent with the main experiment, participants in the *Absolute* condition tend to select scenarios with more equal allocations of additional income growth. Column (2) mirrors the analysis from Section 3 by examining the Gini coefficient of absolute income growth in the chosen scenarios. This coefficient is significantly lower in the *Absolute* condition than in the *Relative* condition. Table 1.A.7 repeats the analysis with controls, confirming that the conclusions and significance levels remain unchanged.

Overall, the results indicate that the conclusions from the main experiment extend to the more natural setting of the survey.

1.4.3 Heterogeneity and Size of the Treatment Effect

As part of a political questionnaire, survey participants indicated their political orientation on a Likert scale from 1 (*very liberal*) to 5 (*very conservative*). Based on these responses, I classify participants as liberal (1 or 2), conservative (4 or 5), or neutral (3). Figure 1.11 presents the Gini coefficient of income growth for all combinations of treatment and political attitude.

Table 1.7. Scenario Choice

	Chosen Scenario	Additional Income Gini
Constant	4.412***	0.139***
	(0.121)	(0.010)
Absolute	0.677***	-0.059***
	(0.170)	(0.015)
R^2	0.047	0.048
Observations	294	294

This table presents OLS regression results. The table is based on the survey study. In column (1), the outcome variable is the number of the *Chosen Scenario* (between 1 and 6), where higher scenarios correspond to more equal allocations of absolute income growth. In column (2), the outcome is the *Additional Income Gini* implied by the scenario choice, as defined in Equation 1.1. *Absolute* is a treatment dummy for the Absolute treatment. Stars highlight significant differences from 0 with * for $p < 0.10$, ** for $p < 0.05$, *** for $p < 0.01$. Participant-clustered standard errors in parentheses.

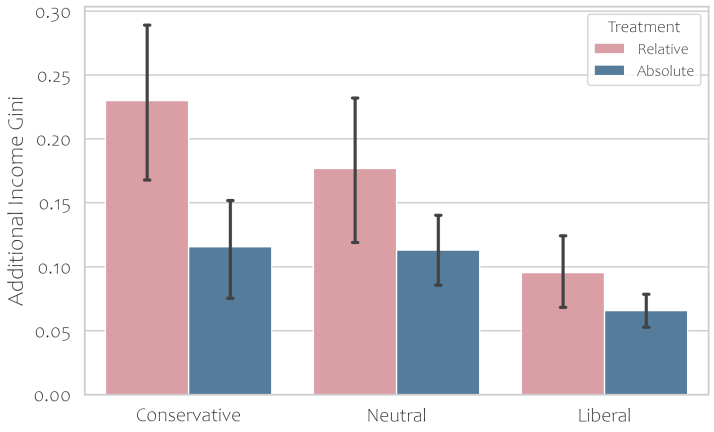

Figure 1.11. Survey treatment effect by political opinion. *This figure plots the additional income Gini implied by participants' chosen scenario in the survey study, by treatment condition and political affiliation.*

The treatment effect is stronger and more significant among conservative participants than among liberal participants. Since conservatives select less equal distributions in the *Relative* treatment, the framing effect has greater scope compared to liberals, who choose relatively equal distributions even when presented with percentage-based information.

To assess the magnitude and relevance of the treatment effect, I compare the distributions chosen by conservative participants in the *Absolute* treatment to those chosen by liberal participants in the *Relative* treatment. The implied Gini coefficients for these two groups do not differ significantly, suggesting that absolute and relative framing can have a substantial impact on distributional choices.

1.4.4 Results on Confusion

At the end of the survey, participants in both treatments were asked to identify the income group with the largest absolute income growth in the scenario presented in Figure 12. While this task was trivial in the *Absolute* treatment, it served to identify participants who confused relative and absolute growth in the *Relative* treatment.

Participants were classified as *confused* if they did not provide the correct answer (Group 5). In the *Absolute* treatment, 2% of participants were confused, compared to 24% in the *Relative* treatment ($p < 0.001$). Thus, while some participants in the *Relative* treatment exhibited signs of confusion, the overwhelming majority answered correctly.

Table 1.8 replicates various analyses from the previous section, restricting the sample to non-confused participants. None of the results change in a meaningful way, suggesting that confusion is not the primary driver of the treatment effect observed in the survey—consistent with findings from the main experiment.

Table 1.8. Results for non-confused participants

	Inequality Change	Alarming Dev.	Pooled Pol. Supp.	Add. Inc. Gini
Constant	4.223***	3.884***	5.414***	0.127***
	(0.133)	(0.156)	(0.148)	(0.011
Absolute	1.371***	1.613***	0.556***	-0.048***
	(0.178)	(0.208)	(0.171)	(0.015)
R^2	0.187	0.188	0.030	0.038
Observations	255	255	510	255

This table presents OLS regression results for the sample of non-confused subjects. The table is based on the survey study. In column (1) the outcome is *Change in Inequality*, which is the change in inequality induced by the presented scenario on a Likert scale from 1 ("*Inequality strongly decreases*") to 7 ("*Inequality strongly increases*"). In column (2), the outcome is *Alarming Development*, which is the agreement with the statement that the depicted development is alarming, using a 1-to-7 Likert scale. In column (3), the outcome is the *Pooled Policy Support* for the minimum wage and income tax questions, measured on a 1-to-7 Likert scale. In column (4), the outcome is the *Additional Income Gini* as defined in Equation 1.1. Column (3) includes Question FE. Stars highlight significant differences from 0 with * for $p < 0.10$, ** for $p < 0.05$, *** for $p < 0.01$. Participant-clustered standard errors in parentheses.

1.5 Conclusion

Economic inequality can be defined in two main ways: absolute inequality and relative inequality. These concepts can lead to fundamentally different conclusions about trends in inequality. This chapter shows that framing changes in inequality as absolute versus relative systematically shapes distributional choices. In a laboratory experiment, participants allocate additional income between the winner and the loser of a contest, with the same options presented either as absolute amounts or as percentage growth relative to an initial contest prize.

Framing strongly affects choices: participants in the *Absolute* condition allocate additional income more equally, whereas those in the *Relative* condition tend to equalize percentage growth.

I examine two potential mechanisms for this effect: confusion about how allocations translate into final incomes versus framing-induced differences in fairness concepts. While some participants show signs of confusion, the treatment effect remains significant even among those who accurately anticipated the consequences of their choices. Instead, the evidence suggests that framing shifts what participants view as fair, rather than simply misleading them.

A complementary survey study replicates these findings in the real-world context of U.S. income growth, showing that such framing impacts both the perceived rise in inequality as well as support for redistributive policies.

Taken together, the results demonstrate that the way information about inequality is presented can meaningfully shape fairness preferences and attitudes toward redistribution. Even when the underlying facts remain the same, framing income changes in absolute or relative terms can shift perceptions of inequality and the policies people support, highlighting the political relevance of how economic information is communicated.

Appendix 1.A Appendix

1.A.1 Additional Tables and Figures: Main Experiment

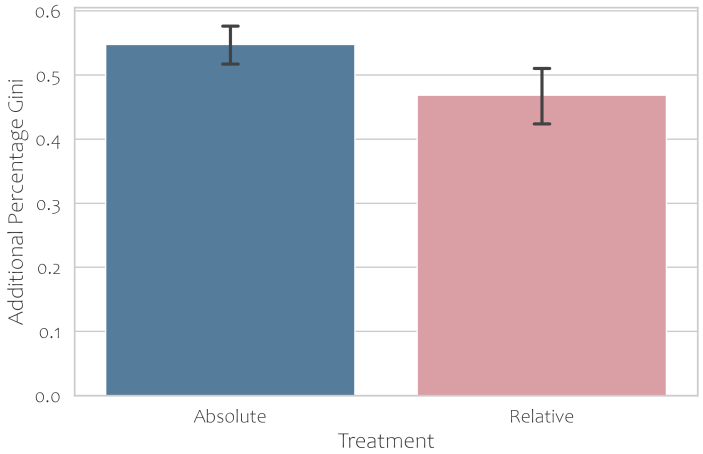

Figure 1.A.1. Mean additional percentage Gini in both treatments. *This figure plots the average additional percentage Gini as defined in Equation 1.2 for both treatment conditions.*

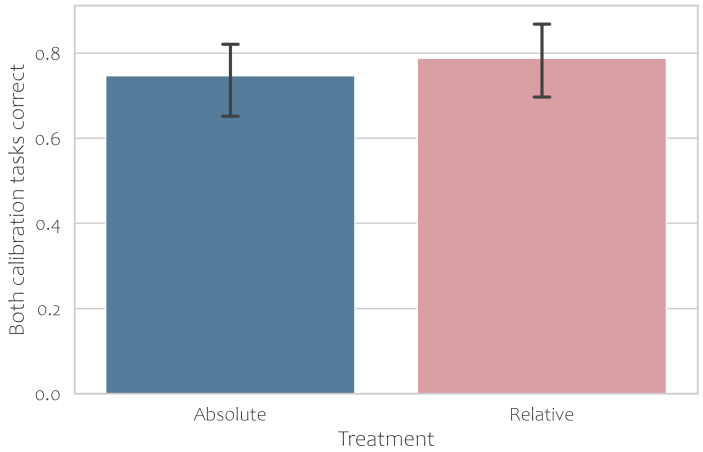

Figure 1.A.2. Calibration task performance. *This figure plots the share of participants who were able to correctly identify both the option with equal additional incomes and equal additional percentage growth in their respective treatment framing.*

Table 1.A.1. Gini coefficient for percentage growth

	Additional Percentage Gini		
	(1)	(2)	(3)
Constant	0.469***	0.508***	0.686***
	(0.022)	(0.072)	(0.068)
Absolute	0.079***	0.080***	0.100***
	(0.026)	(0.026)	(0.022)
Participant characteristics	✗	✓	✓
Participant answers	✗	✗	✓
R^2	0.024	0.039	0.169
Observations	1552	1552	1552

This table presents OLS regression results for the outcome variable *Additional Percentage Gini* as defined in Equation 1.2. The table is based on the main experiment. *Absolute* is a treatment dummy for the Absolute treatment. Stars highlight significant differences from 0 with * for $p < 0.10$, ** for $p < 0.05$, *** for $p < 0.01$. Clustered standard errors in parentheses. All coefficients are estimated using OLS. Participant characteristics controls include age, male dummy, education level and a dummy for an above median CRT score. Participant answer controls include answers to the questions on meritocracy and egalitarianism as well as answers to the political questionnaire.

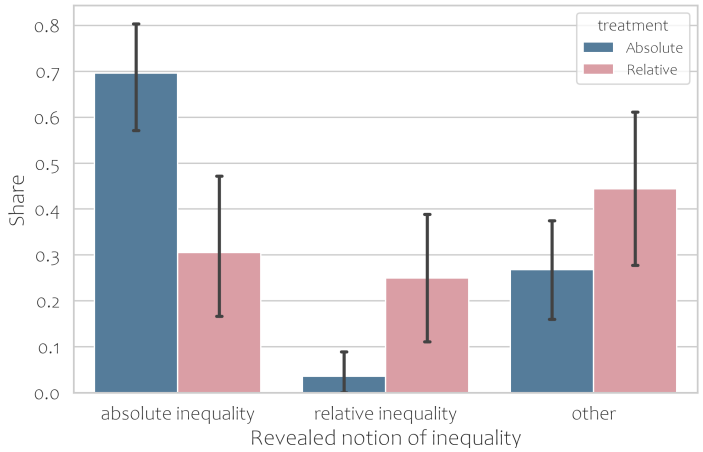

Figure 1.A.3. Revealed inequality attitudes of *non-confused* participants. *This figure plots the share of participants whose answers to the inequality rating questions fully align with the notion of absolute inequality, relative inequality, or neither notion. The figure is the analogue to Figure 1.9, based only on participants who correctly estimated the wealth implications of their Baseline choice.*

Table 1.A.2. Treatment Effect by Attitudes Towards the Additional Money

	Additional Income Gini			
	(1)	(2)	(3)	(4)
Meritocracy dummy	1	1	0	0
Egalitarian dummy	0	1	0	1
Constant	0.456***	0.414***	0.323***	0.308***
	(0.035)	(0.049)	(0.037)	(0.024)
Absolute	−0.072	−0.083	−0.154***	−0.108***
	(0.057)	(0.061)	(0.045)	(0.031)
R^2	0.017	0.029	0.071	0.042
Observations	400	264	280	600

This table presents OLS regression results for the outcome variable *Additional Income Gini* as defined in Equation 1.1. The table is based on the main experiment. *Absolute* is a treatment dummy for the Absolute treatment. *Meritocracy dummy* equals 1 if and only if the participant agrees with the statement 'The winner should be awarded the larger share of the additional money, on top of receiving the greater prize.' *Egalitarian dummy* equals 1 if and only if the participant disagrees with the statement 'Holding everything else constant, the winner of a contestant pair should earn more money than the loser in today's study.' Stars highlight significant differences from 0 with * for $p < 0.10$, ** for $p < 0.05$, *** for $p < 0.01$. Clustered standard errors in parentheses.

Table 1.A.3. Within-subject Regression Analysis

	Additional Income Gini
Constant	0.284***
	(0.019)
Absolute	−0.060**
	(0.026)
After change	−0.053***
	(0.014)
Absolute × after change	0.080***
	(0.019)
R^2	0.009
Observations	388

This table presents OLS regression results for the outcome variable *Additional Income Gini* in the Inbetween decision scenario as defined in Equation 1.1. The table is based on the main experiment. *Absolute* is a treatment dummy for the Absolute treatment. *After change* is a dummy that equals 1 for the participants' revised decisions after being shown both the absolute and relative framing simultaneously. Stars highlight significant differences from 0 with * for $p < 0.10$, ** for $p < 0.05$, *** for $p < 0.01$. Clustered standard errors in parentheses. Sample includes all *Inbetween* scenario decisions before and after the decisions could be changed.

Table 1.A.4. Confusion and Decision Changes

	Changed Decision	Income diff. change winner - loser
Constant	0.167***	−0.194**
	(0.063)	(0.078)
Absolute	0.030	0.328***
	(0.083)	(0.109)
Confused	0.151*	−0.194
	(0.086)	(0.149)
Absolute × *Confused*	−0.091	0.266
	(0.124)	(0.223)
R^2	0.004	0.077
Observations	194	194

This table presents OLS regression results. The table is based on the main experiment. In column (1), the outcome is a dummy for whether a participant changed their decision after being shown both the absolute and relative framing in the Inbetween decision scenario. In column (2), the outcome is *Income diff. change winner - loser*, which is defined as the wealth difference in the *Inbetween* scenario after the decision change opportunity minus the difference in the original choice. Stars highlight significant differences from 0 with * for $p < 0.10$, ** for $p < 0.05$, *** for $p < 0.01$. Clustered standard errors in parentheses.

1.A.2 Additional Tables and Figures: Survey Study

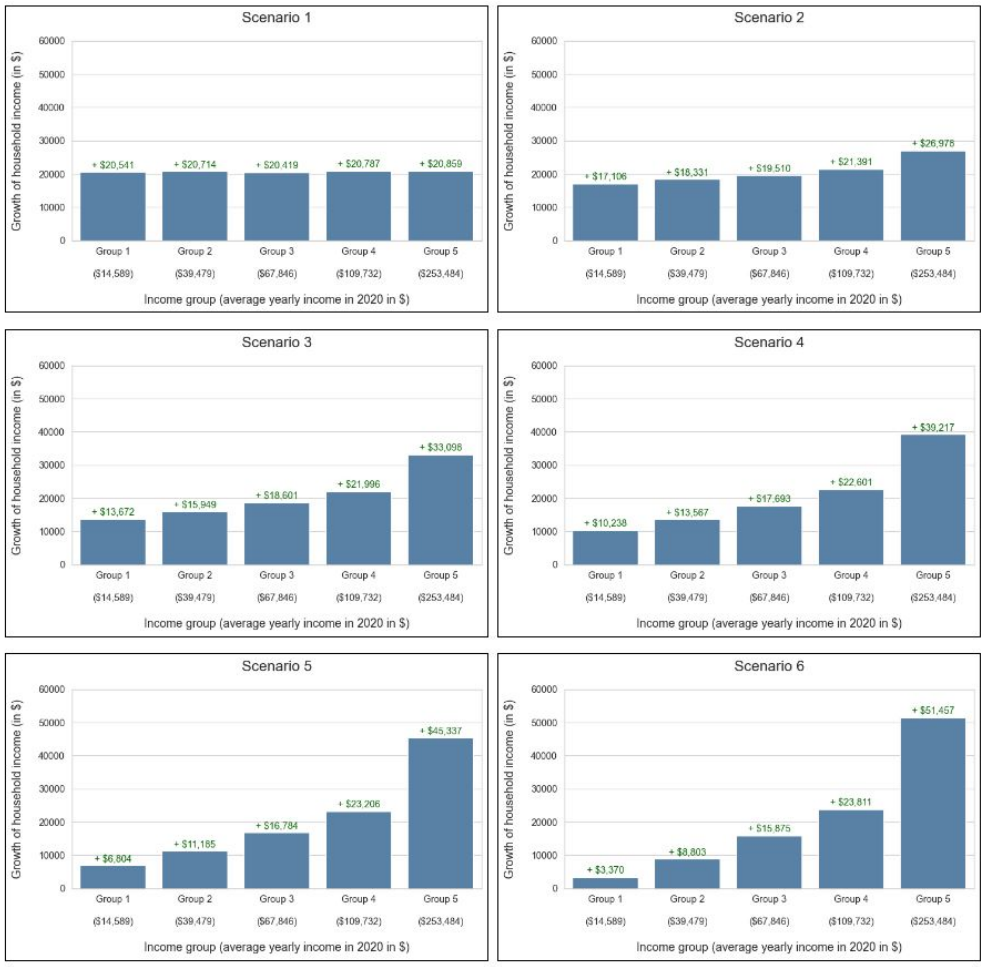

Figure 1.A.4. Scenarios to choose from in the second part of the the survey study. *This figure presents the scenarios in the framing of the Absolute condition.*

Table 1.A.5. Perceived Development of Inequality (with controls)

	Change in Inequality	Alarming Development
Constant	4.471***	4.970***
	(0.365)	(0.420)
Absolute	1.465***	1.478***
	(0.172)	(0.198)
Controls	✓	✓
R^2	0.186	0.193
Observations	294	294

This table presents OLS regression results. The table is based on the survey study. In column (1) the outcome is *Change in Inequality*, which is the change in inequality induced by the presented scenario on a Likert scale from 1 (*"Inequality strongly decreases"*) to 7 (*"Inequality strongly increases"*). In column (2), the outcome is *Alarming Development*, which is the agreement with the statement that the depicted development is alarming, using a 1-to-7 Likert scale. Stars highlight significant differences from 0 with * for $p < 0.10$, ** for $p < 0.05$, *** for $p < 0.01$. Robust standard errors in parentheses. Controls include answers to a political questionnaire and baseline demographics such as gender, age, income and estimated income group.

Table 1.A.6. Policy Support (with controls)

	Minimum Wage	Income Tax	Pooled Support
Constant	7.007***	6.519***	6.684***
	(0.320)	(0.318)	(0.223)
Absolute	0.260*	0.266*	0.263**
	(0.151)	(0.150)	(0.122)
Controls	✓	✓	✓
Question FE	✗	✗	✓
R^2	0.280	0.368	0.331
Observations	294	294	588

This table presents OLS regression results. The table is based on the survey study. In column (1), the outcome is agreement with increasing the *Minimum Wage*. In column (2), the outcome is agreement with increasing the *Income Tax* for the highest income group. Agreement was measured on a Likert scale from 1 (strongly disagree) to 7 (strongly agree). Column (3) pools observations from (1) and (2). Stars highlight significant differences from 0 with * for $p < 0.10$, ** for $p < 0.05$, *** for $p < 0.01$. Robust standard errors in parentheses. Controls include answers to a political questionnaire and baseline demographics such as gender, age, income and estimated income group.

Table 1.A.7. Scenario Choice (with controls)

	Chosen Scenario	Additional Income Gini
Constant	4.834***	0.104***
	(0.347)	(0.030)
Absolute	0.634***	-0.054***
	(0.164)	(0.014)
Controls	✓	✓
R^2	0.186	0.187
Observations	294	294

This table presents OLS regression results. The table is based on the survey study. In column (1), the outcome variable is the number of the *Chosen Scenario* (between 1 and 6), where higher scenarios correspond to more equal allocations of absolute income growth. In column (2), the outcome is the *Additional Income Gini* implied by the scenario choice, as defined in Equation 1.1. *Absolute* is a treatment dummy for the Absolute treatment. Stars highlight significant differences from 0 with * for $p < 0.10$, ** for $p < 0.05$, *** for $p < 0.01$. Robust standard errors in parentheses. Controls include answers to a political questionnaire and baseline demographics such as gender, age, income and estimated income group.

1.A.3 Decision Scenarios in the Main Experiment

Table 1.A.8. Baseline Scenario

Option	€ Winner	€ Loser	% Winner	% Loser	Dummies
1	3.00	0.00	60	0	
2	2.50	0.50	50	50	Percentage
3	2.00	1.00	40	100	
4	1.50	1.50	30	150	Income
5	1.00	2.00	20	200	
6	0.50	2.50	10	250	
7	0.00	3.00	0	300	

Table 1.A.9. Shifted Baseline Scenario

Option	€ Winner	€ Loser	% Winner	% Loser	Dummies
1	3.50	0.25	70	25	
2	3.00	0.75	60	75	Percentage
3	2.50	1.25	50	125	
4	2.00	1.75	40	175	Income
5	1.50	2.25	30	225	
6	1.00	2.75	20	275	
7	0.50	3.25	10	325	

Table 1.A.10. Extreme Scenario

Option	€ Winner	€ Loser	% Winner	% Loser	Dummies
1	3.75	0.75	75	75	Percentage
2	3.50	1.00	70	100	
3	3.25	1.25	65	125	
4	3.00	1.50	60	150	
5	2.75	1.75	55	175	
6	2.50	2.00	50	200	
7	2.25	2.25	45	225	Income

Table 1.A.11. Inbetween Scenario

Option	€ Winner	€ Loser	% Winner	% Loser	Dummies
1	4.75	1.25	95	125	Percentage
2	4.50	1.50	90	150	
3	4.25	1.75	85	175	
4	4.00	2.00	80	200	
5	3.75	2.25	75	225	
6	3.50	2.50	70	250	
7	3.25	2.75	65	275	Income

Table 1.A.12. Equal Incomes Only Scenario

Option	€ Winner	€ Loser	% Winner	% Loser	Dummies
1	4.00	2.00	80	200	Percentage
2	3.75	2.25	75	225	
3	3.50	2.50	70	250	
4	3.25	2.75	65	275	
5	3.00	3.00	60	300	Income
6	2.75	3.25	55	325	
7	2.50	3.50	50	350	

Table 1.A.13. Equal Incomes Only Scenario

Option	€ Winner	€ Loser	% Winner	% Loser	Dummies
1	4.50	0.00	90	0	
2	4.25	0.25	85	25	
3	4.00	0.50	80	50	
4	3.75	0.75	75	75	Percentage
5	3.50	1.00	70	100	
6	3.25	1.25	65	125	
7	3.00	1.50	60	150	Income

Table 1.A.14. Equal Incomes Extreme Scenario

Option	€ Winner	€ Loser	% Winner	% Loser	Dummies
1	3.00	0.00	60	0	
2	2.75	0.25	55	25	
3	2.50	0.50	50	50	Percentage
4	2.25	0.75	45	75	
5	2.00	1.00	35	100	
6	1.75	1.25	30	125	
7	1.50	1.50	25	150	Income

Table 1.A.15. Equal Percentages Extreme Scenario

Option	€ Winner	€ Loser	% Winner	% Loser	Dummies
1	2.50	0.50	50	50	Percentage
2	2.25	0.75	45	75	
3	2.00	1.00	40	100	
4	1.75	1.25	35	125	
5	1.50	1.50	30	150	Income
6	1.25	1.75	25	175	
7	1.00	2.00	20	200	

1.A.4 Instructions for the Main Experiment

The experiment was conducted in German. The instructions here are a translation.

1.A.4.1 Stage 1 Instructions: Page 1

Now part 1 of the study begins.

For this part of the study, you will be assigned a random other participant as your opponent in a competition.

The competition will proceed as follows: You and your opponent will each be given 2 minutes to solve as many tasks as possible. Each task involves counting the number of zeros in a number combination consisting of zeros and ones. More information about the tasks of the competition follows on the next page.

You will receive prize money for your performance in the competition. If you have solved more tasks than your opponent, you will receive a winner's prize of 5 euros. If you have solved fewer tasks, you will receive a loser prize of 1 euro. In case of a tie, a coin flip will determine the outcome.

All in all, winning the competition results in a higher variable compensation from part 1 of the study than losing the competition.

Since today's study takes place online, the pairings will only be drawn and the competition described above will only be evaluated once all participants have completed the experiment.

1.A.4.2 Stage 1 Instructions: Page 2

Below is an example of the type of number combinations you will see in the competition.

00010101101011110110

The opponents in each pair of participants each work on the same combinations in the same order to ensure comparability of the tasks.

A task is considered correctly solved if you specify the correct number of zeros (in the example: 9).

You win the competition if you solve more tasks than your opponent within 2 minutes. You cannot skip any tasks. If you have solved a task correctly and

confirmed it, the next task will appear automatically.

1.A.4.3 Stage 1 Instructions: Comprehension Questions

Question 1: How is your personal variable payout determined from this part of the study?

- You will receive the winner prize if you solve more tasks than your opponent. If you solve fewer tasks, you receive the loser prize. In case of a tie, a draw will be made.
- In any case, your variable payout from this part of the study is independent of your performance in the competition.

Question 2: Which statement about the prize money from the contest is correct?

- It makes no difference to the amount of prize money whether you win the contest or not. So effort does not pay financially.
- The prize for winning is higher than the prize for losing. So it is worth making an effort.

Question 3: Which statement is true?

- You and your opponent are working on different tasks.
- You and your opponent work on the exact same tasks in the exact same order.

1.A.4.4 Stage 2 Instructions: Page 1

Now part 2 of the study begins.

This part will be about extra money that you and the other participants can get besides the prize money from part 1.

For each pair of participants there is a total amount of extra money. This money will be divided between the loser and winner of each pair by a third participant. You will also make such decisions for another pair of participants.

The other participants, like you, did not know about this extra money during the contest.

We are first interested in your opinion regarding the following statements about the distribution of the extra money.

Statement 1: The winner of the contest should receive the larger share of the

extra money in addition to the higher prize money. (Likert scale elicitation)

Statement 2: The winner of the contest should earn more money in the experiment than the loser of the contest, all other decisions being equal. (Likert scale elicitation)

1.A.4.5 Stage 2 Instructions: Page 2

In Part 2 of the study, we ask you to make decisions about the potential payoffs of other study participants.

As mentioned before, your performance from the previous Part 1 will be compared to that of a randomly selected other study participant. The same is true for all other participants in today's study.

The competition will determine the prize money: The winner in each pairing receives a winner prize, and the loser receives a loser prize.

In treatment *Absolute*:
In this part you decide on the amount of possible additional income, which will be paid for another pair of participants in addition to the prize money for the competition. The total payments to the players in the contest will be:

Winner: Total Payment = Winner Prize + Additional Income for the Winner.

Loser: Total Payment = Loser Prize + Additional Income for the Loser.

In treatment *Relative*:
In this part you decide on the amount of possible percentage increases by which the prize money of another pair of participants will be multiplied. The total payments to the players in the competition will then be:

Winner: Total Payment = Winner Prize x (1 + Increase in percentage points for the winner).

Loser: Total Payment = Loser Price x (1 + Increase in percentage points for the loser).

1.A.4.6 Stage 2 Instructions: Page 3

The structure of part 2 of the study is as follows:

In treatment *Absolute*:

You will determine the potential additional income of another pair of participants in 8 situations.

In each situation, you will see seven different proposals for the distribution of additional income between the winner and loser of the contest. The total amount of money distributed is the same for all proposals within a situation.

In treatment *Relative*:
You will determine in 8 situations about the potential percentage increases of another pair of participants.

In each situation, you will see seven different proposals for distributing the percentage increases to the winner and loser of the contest. The total amount of money distributed is the same for all proposals within a situation.

In both treatments:
In the following, you can select one of these proposals.

The 8 situations differ in the amount of additional money as well as in the proposals available for selection. Each decision situation is unique. So please choose carefully the proposal you prefer in each decision.

Participants were then shown an example of how options were visualized, as in Figure 1.

1.A.4.7 Stage 2 Instructions: Page 4

The consequences of your decisions today are the following:

For each pair of participants, a third participant is drawn at the end of the experiment, called the decision maker.

In treatment *Absolute*:
A decision by this decision maker from today's trial then determines the additional income for the winner and loser from each pair of participants.

There is an 80% probability that this decision comes from the current part of the study. In this case, one of the 8 decisions of the decision maker from part 2 is randomly selected and implemented. With a probability of 20%, the relevant decision comes from a later study part.

Thus, each of your decisions today can be implemented and thus determine

parts of the payoffs of other study participants. You should therefore make these decisions carefully.

All decisions are made completely anonymously. If you are selected as a decision maker for a pair of participants, they will only know the result of your decision, but cannot attribute it to you.

Your personal variable payout from this part of the study is determined as follows:

An uninvolved third participant is also selected as the decision maker for the pair of participants to which they belong.

The actions of this decision maker then determine the additional income of you and your opponent as described above.

In treatment *Relative*:
A decision by this decision maker from today's study then determines the percentage growth for the winner and loser from each pair of participants.

This decision comes from the current study part with a probability of 80%. In this case, one of the 8 decisions of the decision maker from part 2 is randomly selected and implemented. With a probability of 20%, the relevant decision comes from a later study part.

You can be selected as the decision maker for any other pair of participants.

Thus, each of your decisions today can be implemented and thus determine parts of the payoffs of other study participants. You should therefore make these decisions carefully.

All decisions are made completely anonymously. If you are selected as a decision maker for a pair of participants, they will only know the result of your decision, but cannot attribute it to you.

Your personal variable payout from this part of the study is determined as follows:

An uninvolved third participant is also selected as the decision maker for the pair of participants to which they belong.

The actions of this decision maker then determine the percentage increases for you and your opponent as described above.

1.A.4.8 Stage 2 Instructions: Comprehension questions

Question 1: Which of the following statements about your choices in this part of the study is correct?

- All presented decision scenarios are hypothetical.
- Your decisions can affect the payoffs of other participants.

Question 2 (*Absolute*): The total variable payout is made up of the prize for the competition and the additional income. Which components do you decide on in this part of the study?

- Additional income.
- Additional income and prize for the competition.

Question 2 (*Relative*): The total variable payout is composed of the prize for the competition as well as the percentage increase. Which components do you decide on in this part of the study?

- Percentage increase.
- Percentage increase and prize for the competition.

Question 3: Did you (and the other contestants) know before the contest in Part 1 that you can receive additional money besides the prize money, the amount of which is decided by another contestant?

- Yes.
- No.

Question 4: Here participants were shown an example for an option as in Figure 1, and were asked to state the additional incomes / values for percentage growth as well as the implied total payments.

1.A.4.9 Additional questions at the end of part 2

Question 1
On the previous pages, you made decisions about how to divide extra money between the winner and loser of a contest.
Please describe briefly below how you made these decisions. For example, were there rules of thumb that you used to guide your proposal selection?

Question 2

On the previous pages, they divided additional money between the winner and loser of a contest, which they receive in addition to the prize money for the contest. The total payment in the different scenarios was made up as follows:
(Descriptions of the total payments in the treatment congruent framing)
In one of the decisions, the conditions were as follows:
Winner prize: 5.00 Euro.
Loser prize: 1,00 Euro.
Additional sum: 3,00 Euro.
All possible distributions of the additional sum to both players were selectable in increments of 50 cents.
We now want to know from you: What were the total payments (in euros) to the winner and loser of the contest resulting from your decision in this situation?
Reward: If your estimate for both players differs from the truth by no more than 0.50 euros each, you will receive a payout bonus of 1.00 euro.

Question 3
On the previous pages, you made decisions about how to divide extra money between the winner and loser of a contest.
In doing so, you were able to choose from 7 proposals in each of 8 decisions, so you saw a total of 56 proposals.
Did you feel that either the winner or the loser of the contest was financially favored in the majority of these proposals, in terms of how the additional money was divided?
Important: This question is not about who you personally favored through your choices. Instead, we are asking for your impression on who was favored in the majority of the proposals up for selection.

- The loser was favored in most proposals.
- The winner was favored in most of the proposals.
- The proposals were balanced.

Question 4:
On the previous pages, you made decisions about how to allocate additional money to the winner and loser of a contest.
In the following, we ask you to estimate the inequality generated by selected proposals.
In particular, we are interested in your assessment of the change in inequality: Do you think that the inequality between the winner and loser increases, decreases, or stays the same compared to the baseline situation of prize money of 5 euros and 1 euro?
There are no right and wrong answers here, we are just asking for your personal assessment.

Option 1
The payout of the winner increases *by 30% / 1.50 € (depending on treatment)* from 5.00 € to 6.50 €.
The payout of the loser increases *by 150% / 1.50 €* from 1.00 € to 2.50 €.

How do you evaluate the change in inequality if this proposal is chosen? Inequality...

- decreases strongly.
- decreases.
- decreases slightly.
- remains unchanged.
- increases slightly.
- increases.
- increases strongly.

Option 2
The payout of the winner increases *by 40% / 2.00 €* from 5.00 € to 7.00 €.
The payout of the loser increases *by 100% / 1.00 €* from 1.00 € to 2.00 €.

Option 3
The payout of the winner increases *by 50% / 2.50 €* from 5.00 € to 7.50 €.
The payout of the loser increases *by 50% / 0.50 €* from 1.00 € to 1.50 €.

Additionally, graphs visualizing the options as in Figure 1 were displayed.

1.A.4.10 Stage 3 Instructions: Page 1

Now part 3 of the study begins.

In this part you will again see suggestions for the distribution of *percentage increases / additional incomes (depending on treatment*) among different players.

However, all situations in this part of the study are purely hypothetical. This means that your decisions in this part have no consequences for the payouts of other study participants.

To make this fact clear, we will call the players in this part Robot A and Robot B.

We will ask you to identify a particular distribution of *percentage increases / additional incomes (depending on treatment*) in two situations.

In Decision 1, we are looking for the split that leads to equal absolute income gains for both robots. This means that both robots should receive the same additional amount of euros in addition to their different prize money. An example would be an additional income of €2.50 each for both robots in addition to their prize money.

In decision 2, we are looking for the split that results in equal percentage income increases for both robots. This means that the payoff from the additional income should increase by the same factor for both robots. An example of this would be a doubling for both robots relative to their prize money, or a percentage increase in income of 100%.

Your answer in a randomly selected task will be rewarded. If you choose the proposition you are looking for in this one, you will receive a reward of €1.00 in addition to all your previous income in the experiment.

1.A.4.11 Stage 3 Instructions: Comprehension questions

Question 1: Which of the statements about this part of the experiment is correct?

- Your decisions in this part can have consequences for the payoffs of other study participants.
- Your decision in this part can only affect your own payoff.

Question 2: Which split should you choose in the first task?

- The split that you personally think is fairest.
- The distribution that leads to equal additional euro amounts.

1.A.4.12 Additional questions at the end of part 3

Question 1: In the second part of the study, you made decisions about how to allocate additional money to the winner and loser of the contest.
Were you aware at this point that a proposal with equal bar lengths for the winner and loser (and thus *equal percentage increases / equal additional incomes (depending on treatment)*) would not result in *equal additional absolute incomes / equal percentage growth (depending on treatment)* for both players?

- Yes
- No

Question 2: Looking back, would you want to change your decisions from part two in this regard?

- Yes

- No

Question 3: If you would like to explain your answers above or make any other comments, please feel free to do so here.

1.A.4.13 Stage 4 Instructions: Main Instructions

Now part 4 of the study begins.

In this part of the study, we will once again show you a decision situation from Part 2 of the study. In addition, we will show you which suggestion you chose in this situation in part 2 of the study.

You can subsequently decide to change this answer.

As in Part 2 of the study, we show you *the percentage increases / additional incomes (depending on treatment* for the winner and loser of the competition for each option. These are highlighted in blue.

In addition, for each option, we now show in red *the absolute incomes in euros / percentage increases (depending on treatment)* for the winner and loser of the contest.

Below is an example of how the options are shown on the next page.

Example proposal (in treatment *Relative*):
The winner's payout increases by 80% or 4.00 € from 5.00 € to 9.00 €.
The loser's payout increases by 200% or 2.00 € from 1.00 € to 3.00 €.

Example proposal (in treatment *Absolute*):
The winner's payout increases by 4.00 € or 80% from 5.00 € to 9.00 €.
The loser's payout increases by 2.00 € or 200% from 1.00 € to 3.00 €.

If you are selected as the decision maker:
With a probability of 80%, one random of the eight decisions from part 2 of the study will be implemented regardless of your change decision. With the counter probability of 20%, your (possibly changed) decision from the current Part 4 will be implemented.

1.A.4.14 Stage 4 Instructions: Decision Page

Below you see a decision situation from the second part of the study.

The proposal you selected in the second part is marked with "(original selection)".

Participants then saw the seven options.

Do you want to change your decision?

- Yes.
- No.

If participants answered yes, they were forwarded to a decision screen where they could change their decision.

1.A.4.15 Additional questions at the end of part 4

Question 1: On the previous pages, income changes of the two players were presented partly as absolute income increases (in euros) and partly as percentage income increases (in percent, relative to prize money).
Which of the two representations do you find more helpful in making your personal decisions?

- The representation as absolute income.
- The representation as percentage increases.

Question 2: In the first part of the study, you competed against a random other participant. Do you think you solved more tasks than your opponent and thus won the competition?

- No.
- Rather not.
- I am unsure.
- Rather yes.
- Yes.

1.A.4.16 Stage 5 Instructions: General Instructions

Now the last part of the study begins.

In this part of the study, we ask you to do two things.

On the next page, we ask you to answer six quiz questions. Each question has one correct answer. Your answer in a random question will be rewarded 50 cents if it is correct. You will have 60 seconds to answer each question.

Then we ask you to evaluate four theses around the topic of economic inequality and redistribution in Germany. Please indicate how much you agree with each thesis.

On the next page you will see the first quiz question.

1.A.4.17 Stage 5 Instructions: Cognitive reflection / IQ test

Question 1: Five bakers bake five breads in five minutes. How many breads do 100 bakers bake in 100 minutes?

Question 2: Ivy is growing on the facade of a house. The ivy grows so fast that its area doubles every day. Suppose it takes 40 days for the ivy to cover the entire facade. After how many days does the ivy cover a quarter of the facade?

Question 3: Robin received both the fourth highest and fourth lowest grade in his course. How many students (including Robin) are in the course?

Question 4: Suppose you have invested 3000 euros in an account with an annual interest rate of 10 percent. How many euros are in total in the account two years later?

Question 5: *This was a multiple choice raven matrix question.*

Question 6: The following series of letters follows a certain pattern. Which letter comes next in the series? Q S N P L ?

1.A.4.18 Stage 5 Instructions: Political questionnaire

Participants were asked whether they strongly disagree / rather disagree / are neutral about / rather agree / strongly agree with the following statements.

Statement 1: The problem of social inequality in Germany is exaggerated.

Statement 2: The hourly minimum wage should be increased to at least 12 euros as soon as possible.

Statement 3: It should remain possible to reduce Hartz 4 / unemployment benefits by the means of sanctions.

Statement 4: There should be a wealth tax in Germany.

1.A.4.19 Stage 5 Instructions: Demographics

Participants were asked about their age, gender assigned at birth, highest educational attainment, current main occupation, field of study or occupation and monthly after-tax income.

1.A.5 Instructions for the Survey Study

1.A.5.1 General Instructions

Welcome to today's study.

The study will take about 8 minutes, and the questions will begin on the next page. You will receive 1.20 GBP (roughly 1.40 USD) for completing the survey!

After an initial political survey and short introduction, there will be 2 comprehension questions. They can be answered correctly with the information that is provided on the question page. You need to answer these questions correctly in order to get a payment. If you fail a question twice, the survey stops.

If you fail the comprehension questions or do not agree with these terms, please return your submission on Prolific.

After successful completion, you will be automatically redirected to Prolific!

1.A.5.2 Political Questionnaire

Please answer the following questions.

Which party are you most likely to vote for in the next election?

- Democratic Party
- Republican Party
- Libertarian Party
- A different party
- I will not vote

What best describes your political opinion?

- Very liberal
- Leaning liberal
- Moderate
- Leaning conservative

- Very conservative

Where do you stand on the following issues?

- Scale from 1-11 between 'I am pro choice' and 'I am pro life'
- Scale from 1-11 between 'I am for stricter gun control' and 'I am against stricter gun control'
- Scale from 1-11 between 'The government should assist the poor' and 'The government should not be involved in redistributional issues'

1.A.5.3 Instructions

This part of the study concerns income growth in the United States between the years 2020 and 2025.

Economic think tanks regularly try to predict how the incomes of different population groups will develop in the future. These predictions are frequently used by politicians as a guide toward which economic policies to implement.

The figure at the bottom of this page depicts one potential scenario for income growth between the years 2020 and 2025.

To create the figure, the US population was divided into five groups of equal size.

Group 1 contains the 20% of households with the lowest income in the year 2020. Group 2 contains the next 20% of households in terms of income, and so forth, while Group 5 contains the 20% of households that earned the highest incomes in 2020. At the bottom of the figure, the average yearly income in dollars in 2020 is shown for each group in parentheses.

The bars visualize the potential income growth *in dollars / in percentage points* between 2020 and 2025 for each group, that is, how much more money *(in dollars / in percentage points)* the average household in that group would earn in 2025 *on top of their yearly income in 2020 / relative to their yearly income in 2020*.

[Here, the corresponding image from Figure 12 in the main text was shown.]

1.A.5.4 Comprehension

Please answer the following questions to ensure that you fully understand the figure. You need to answer these correctly in order to receive your payment.

Which of the groups consists of the households that had the highest incomes in 2020?

- Group 1
- Group 2
- Group 3
- Group 4
- **Group 5**

According to the figure, which income growth between 2020 and 2025 is projected for the group with the lowest earning households in 2020?

- ***6,830 USD / 46.8 %***
- *28,193 USD / 23.8 %*
- *14,589 USD / 19.1 %*
- *253,484 USD / 15.4 %*
- *9,401 USD / 13.1 %*

1.A.5.5 Main Outcomes

Now, we will ask you some questions. For your answers, please imagine that the figure shows a correct projection of the income growth from 2020 to 2025.

For each question in this section, we are interested in your personal opinion. There are no wrong and right answers.

[Projection displayed on each page]

How would you describe the development of inequality in the United States as projected by the figure?

- strongly decrease
- decrease
- somewhat decrease
- remain constant
- somewhat increase
- increase
- strongly increase

What is your level of agreement with the following statement:'The scenario depicted in the figure implies an alarming development of income inequality.'

- strongly disagree

- disagree
- somewhat disagree
- neither agree nor disagree
- somewhat agree
- agree
- strongly agree

Please specify what you find alarming (if anything).

- Income growth is too high for low earning households, compared to high earning households.
- Income growth is too high for high earning households, compared to low earning households.
- Something else (please specify):
- Nothing is alarming.

The minimum wage is a political instrument that sets a lower bound for the hourly wage that workers must at least be paid. The federal minimum wage is currently at 7.25$. Some people argue in favor of increasing the minimum wage to decrease economic inequality.
What is your level of agreement with the following statement?
'If the above projection were true, the minimum wage should be increased.'

- strongly disagree
- disagree
- somewhat disagree
- neither agree nor disagree
- somewhat agree
- agree
- strongly agree

The income tax is a political instrument that determines how much of their income workers have to pay to the government. Some people argue in favor of raising income taxes for high earners (while lowering them for low earners) to decrease economic inequality.
What is your level of agreement with the following statement?
'If the above projection were true, taxes for the highest income group should be increased while decreasing the tax burden for the lower earning groups.'

- strongly disagree
- disagree

- somewhat disagree
- neither agree nor disagree
- somewhat agree
- agree
- strongly agree

1.A.5.6 Scenario Choice

Now, we are showing you six possible scenarios for how incomes in the US could develop in the years between 2020 and 2025.

Importantly, the average income growth (in $) over all groups is held constant in all of the scenarios. The income growth is only distributed differently between the 5 groups in the different scenarios.

[Here, the six scenarios from Figure A.3 were presented in the respective treatment framing. The order was randomized as either poor increasing or poor decreasing.]

Which of the six scenarios do you consider to be the most desirable?

- Scenario 1
- Scenario 2
- Scenario 3
- Scenario 4
- Scenario 5
- Scenario 6

1.A.5.7 Questions about Confusion and Own Income Group

The next question refers to the initial projection of incomes that you saw. It will be displayed again on the next page.

[Initial figure was displayed.]

Please answer the following question about the initial figure.
In the initial projection, which of the groups has the largest income gain in dollars?

- Group 1
- Group 2
- Group 3
- Group 4

- Group 5

Which income group do you think you belong to?
Recall that Group 1 consists of the 20% of households with the lowest income, Group 2 contains the next 20% in terms of incomes, and Group 5 contains the 20% of households that have the highest incomes. The mean incomes of the groups in 2020 are stated in the graphic above.

- Group 1
- Group 2
- Group 3
- Group 4
- Group 5

1.A.5.8 Demographics

Participants were asked about sex at birth, age, occupation and monthly income before taxes.

References

Amiel, Yoram, and Frank Cowell. 1999. *Thinking about inequality: Personal judgment and income distributions.* Cambridge University Press. [19]

Amiel, Yoram, and Frank A Cowell. 1992. "Measurement of income inequality: Experimental test by questionnaire." *Journal of Public Economics* 47 (1): 3–26. [5, 19]

Andreoni, James, and B Douglas Bernheim. 2009. "Social image and the 50–50 norm: A theoretical and experimental analysis of audience effects." *Econometrica* 77 (5): 1607–36. [5]

Bandyopadhyay, Sanghamitra. 2018. "The absolute Gini is a more reliable measure of inequality for time dependent analyses (compared with the relative Gini)." *Economics Letters* 162: 135–39. [4]

Bock, Olaf, Ingmar Baetge, and Andreas Nicklisch. 2014. "hroot: Hamburg registration and organization online tool." *European Economic Review* 71: 117–20. [10]

Brekke, Kjell Arne, James Konow, and Karine Nyborg. 2017. "Framing in a threshold public goods experiment with heterogeneous endowments." *Journal of Economic Behavior & Organization* 138: 99–110. [5]

Bruckmüller, Susanne, Gerhard Reese, and Sarah E Martiny. 2017. "Is higher inequality less legitimate? Depends on How You Frame it!" *British Journal of Social Psychology* 56 (4): 766–81. [4]

Chen, Daniel L, Martin Schonger, and Chris Wickens. 2016. "oTree—An open-source platform for laboratory, online, and field experiments." *Journal of Behavioral and Experimental Finance* 9: 88–97. [10]

Chen, Haipeng, and Akshay R Rao. 2007. "When two plus two is not equal to four: Errors in processing multiple percentage changes." *Journal of Consumer Research* 34 (3): 327–40. [5, 19]

Chow, Rosalind M, and Jeff Galak. 2012. "The effect of inequality frames on support for redistributive tax policies." *Psychological science* 23 (12): 1467–69. [4]

Cruces, Guillermo, Ricardo Perez-Truglia, and Martin Tetaz. 2013. "Biased perceptions of income distribution and preferences for redistribution: Evidence from a survey experiment." *Journal of public Economics* 98: 100–112. [3]

Dietze, Pia, and Maureen A Craig. 2021. "Framing economic inequality and policy as group disadvantages (versus group advantages) spurs support for action." *Nature human behaviour* 5 (3): 349–60. [3, 4]

Exley, Christine L, and Judd B Kessler. 2024. "Equity concerns are narrowly framed." *American Economic Journal: Microeconomics* 16 (2): 147–79. [5]

Glaser, Markus, Zwetelina Iliewa, and Martin Weber. 2019. "Thinking about prices versus thinking about returns in financial markets." *Journal of Finance* 74 (6): 2997–3039. [5]

Hauser, Oliver P, and Michael I Norton. 2017. "(Mis) perceptions of inequality." *Current Opinion in Psychology* 18: 21–25. [3, 4]

Jun, Sora, et al. 2022. "Chronic frames of social inequality: How mainstream media frame race, gender, and wealth inequality." *Proceedings of the National Academy of Sciences* 119 (21): e2110712119. [3]

Kruger, Justin, and Patrick Vargas. 2008. "Consumer confusion of percent differences." *Journal of Consumer Psychology* 18 (1): 49–61. [5, 19]

Kuziemko, Ilyana, et al. 2015. "How elastic are preferences for redistribution? Evidence from randomized survey experiments." *American Economic Review* 105 (4): 1478–508. [3, 4, 24]

Lembregts, Christophe, and Mario Pandelaere. 2014. "" A 20% income increase for everyone?": The effect of relative increases in income on perceived income inequality." *Journal of Economic Psychology* 43: 37–47. [5]

Lombardi, Gaia, Ernst Fehr, and Todd Hare. 2018. "The attentional foundations of framing effects." Working paper. [19]

Lorenz, Jan, Fabian Paetzel, and Markus Tepe. 2017. "Just don't call it a tax! Framing in an experiment on voting and redistribution." *Journal of Experimental Political Science* 4 (3): 183–94. [3, 5]

Messick, David M, and Terry Schell. 1992. "Evidence for an equality heuristic in social decision making." *Acta Psychologica* 80 (1-3): 311–23. [5]

Niño-Zarazúa, Miguel, Laurence Roope, and Finn Tarp. 2017. "Global inequality: Relatively lower, absolutely higher." *Review of Income and Wealth* 63 (4): 661–84. [4]

Starmans, C, M Sheskin, and P Bloom. 2017. *Why people prefer unequal societies. Nature Human Behaviour, 1 (4), 1–7.* [5]

Yan, Dengfeng. 2019. "Subtraction or division: Evaluability moderates reliance on absolute differences versus relative differences in numerical comparisons." *Journal of Consumer Research* 45 (5): 1103–16. [5]

Chapter 2

Model Uncertainty*

Joint with Florian Zimmermann

2.1 Introduction

People rely on mental models to interpret and navigate the complexities of external reality. These models are mental frameworks that shape how individuals process information, form expectations, and make decisions. In contexts of economic decision-making, mental models play a critical role, as individuals attempt to simplify and make sense of complex environments. However, in many situations, economic agents face not only uncertainty about future outcomes, but also model uncertainty—uncertainty regarding which mental model best captures the true dynamics of the environment.

Model uncertainty is prevalent in many areas of economics. For instance, when forming expectations about future inflation or returns on investments, individuals may be uncertain about which underlying model of the economy is most appropriate to use. Similarly, when assessing the effectiveness of government policies, people may lack clarity on which mental model best explains how government actions translate into economic outcomes. In these cases, individuals face a dual problem: not only must they process new information, but they must also decide how to map this information into a decision, given uncertainty about competing models of the world. This problem arguably is especially severe in environments where work-

* We are grateful to Chiara Aina, Kai Barron, John Conlon, Benjamin Enke, Tilman Fries, Nicola Gennaioli, Alex Imas, Botond Koszegi, Alexander Laubel and Chris Roth for very helpful comments and suggestions. Funding by the Deutsche Forschungsgemeinschaft (DFG) through CRC TR 224 (Project A01) and under Germany's Excellence Strategy - EXC 2126/1 - 390838866 is gratefully acknowledged. This project has received funding from the European Research Council (ERC) under the European Union's Horizon 2020 research and innovation programme (ERC Starting Grant: 948424 - MEMEB). Proofreading assistance for this chapter was provided by the LLM ChatGPT; the research design, analyses, and all arguments remain solely the work of the authors.

ing with models is complex. The question hence arises how people function under model uncertainty, in particular when model complexity is high. The uncertainty about which model is correct can have profound implications for decision-making, beliefs, and ultimately, economic outcomes.

Existing research on mental models has largely focused on understanding the determinants of model selection, how mental models can be used to persuade others, and the consequences of relying on misspecified models. For example, Schwartzstein and Sunderam (2021) and Barron and Fries (2024) examine how individuals might be persuaded to adopt specific mental models under different contexts, while Heidhues, Kőszegi, and Strack (2018) explore how biased models can lead to systematic mistakes in decision-making. Similarly, recent work by Frick, Iijima, and Ishii (2022) investigates the strategic use of mental models in shaping economic beliefs. However, less is known about how individuals cope with model uncertainty in the first place.

In this chapter, we aim to fill this gap by studying how individuals operate under model uncertainty. We hypothesize that the complexity of mental models determines whether people neglect model uncertainty. Our experimental results indeed reveal that, when complexity is high, people simplify the world by operating as if one model of the world is correct, hence fully neglecting model uncertainty. This is echoed by hovering data which suggest that when complexity is high, people have a tendency to predominantly attend to one specific model of the world. These results are robust to variations in the signal space as well as variations in the specific task people need to perform. Turning to implications of this complexity-induced simplification, we document two key results: (i) neglect of model uncertainty does not translate into a distorted view of the world. When we directly elicit beliefs about which model of the world is correct, these beliefs do reflect model uncertainty despite model neglect in actions, creating a wedge between actions and beliefs; (ii) model neglect creates an illusion of certainty. The complexity-induced simplification leads to higher levels of confidence in the optimality of own actions such that, perhaps counter to intuitions, complexity increases rather than decreases decision confidence in our setting.

We design an experiment that allows us to infer from their actions how people deal with model uncertainty. We implement our experiment in a financial decision-making context where participants task is to provide valuations for fictitious companies. There is a set of variables that are potentially relevant to determine company values. Participants face uncertainty about which of two models of the world makes correct use of these variables to estimate the value of a company. Participants are endowed with a 50-50 prior and then receive a noisy signal about which model is correct. In a between-subjects design, we manipulate the complexity of these models, keeping all other aspects of the decision environment constant. In the low-complexity condition, each model provides a direct estimate of the company's value using the variables as input, requiring no computations from participants to get at

each model's estimate. In the high-complexity condition, we add a layer of computational complexity, as participants need to compute the value estimates of the models themselves using the variables as inputs. Participants provide value estimates for eight different companies. Afterwards, we also measure beliefs about which model participants think is correct. This allows us to examine the extent to which complexity influences both actions (i.e., the valuation) and beliefs (i.e., the perception of which model is correct).

Our findings suggest that complexity significantly alters how people deal with model uncertainty. Compared to the low-complexity condition, participants in the high-complexity condition have a more pronounced tendency to simplify the world by neglecting model uncertainty in their actions, behaving as if the more likely model is definitely correct. This tendency to act as if there is no model uncertainty implies an overreaction to the signal about the correct model. In other words, when complexity is high, participants place too much weight on the value estimate provided by the more likely model. Data on hovering times indicate that complexity affects decision-making through an attention channel (Bordalo, Gennaioli, and Shleifer (2012), Bordalo et al. (2023b), Bordalo et al. (2023a), Ba, Bohren, and Imas (2022)). In the high-complexity condition, participants spend significantly more time focusing on the signal-congruent model, whereas in the low-complexity condition, their attention is more evenly distributed between the two models.

In another experiment, we verify that simplification does not require one model to be more likely than the other. When both models of the world are equally likely, complexity nonetheless leads to a neglect of model uncertainty. Additional experiments further underscore the robustness of our results. When we replace the value estimation task with an investment task, we continue to see that complexity leads to a neglect of model uncertainty.

Next, we turn to implications of this complexity-induced simplification. We first document that neglect of model uncertainty in actions does not translate into beliefs. Participants' stated beliefs continue to reflect model uncertainty. This disconnect between beliefs and actions creates a systematic wedge between what individuals believe and how they behave under model uncertainty. This echoes earlier work by Giglio et al. (2021), Ameriks et al. (2020), Beutel and Weber (2023), Laudenbach et al. (forthcoming) who have identified a gap between subjective beliefs and economic behavior in different contexts.[1] Notably, this gap arises directly after participants stated their valuations, and persists even after a 1 day delay, as we verify in a separate experiment that spans over 2 days.

Second, using additional experiments where we measure participants' confidence in the optimality of their actions, we document that confidence in the optimality of own actions is higher in the high-complexity conditions compared to the

1. Yang (2023) provides an explanation for the attenuated relation between beliefs and actions based on a specific form of cognitive uncertainty.

low-complexity conditions. This holds for an unincentivized confidence measure, where participants state the probability that their guesses were optimal, as well as an incentivized measure, where participants place a bet on the optimality of their actions. Prior literature in contexts different from ours has shown that complexity tends to increase cognitive uncertainty (Enke and Graeber (2023), Enke, Graeber, and Oprea (forthcoming), Enke, Graeber, and Oprea (2023)). In contrast, it seems that in the context of model uncertainty, the possibility to respond to complexity by simplifying the world through full neglect of model uncertainty leads to an illusion of certainty and hence increased confidence in action optimality. In the final part of the chapter we present a simple model that formalizes this intuition and can generate the key results from our experiments. The model is an augmented and simplified version of Bordalo et al. (2025).

Our work directly relates to a growing literature on (misspecified) mental models (Schwartzstein and Sunderam (2021), Montiel Olea et al. (2022), Gagnon-Bartsch, Rabin, and Schwartzstein (2023), Mailath and Samuelson (2019), Heidhues, Kőszegi, and Strack (2018) Heidhues, Kőszegi, and Strack (2023) Frick, Iijima, and Ishii (2022), Aina (2024), Barron and Fries (2024)).[2] Until now, this literature has largely focused on an analysis of the determinants of model selection, how mental models can be used to persuade others, and the consequences of relying on misspecified models. In contrast, our focus is on how people deal with model uncertainty. A common assumption in the literature is that people work with a single (possibly misspecified) model when making decisions, rather than entertaining multiple weighted models simultaneously (e.g. Schwartzstein (2014), Schwartzstein and Sunderam (2021), Montiel Olea et al. (2022)).[3] We empirically document this kind of simplification, show that it increases with decision complexity and study its implications.

Our work also ties to a literature that studies how limited attention shapes how people react to information.[4] Bordalo et al. (2023b) and Bordalo et al. (2023a)

2. Relatedly a growing theoretical and empirical literature studies the role of stories and narratives in economics (e.g., Shiller (2017), Eliaz and Spiegler (2022), Andre et al. (2024), Graeber, Roth, and Zimmermann (2024), Graeber, Roth, and Schesch (2024)).

3. Aina and Schneider (2025) study how people update their beliefs in the presence of competing models that could explain observed data. They provide evidence that the majority of people select the model that explains the observed data best. While our focus is on the role of complexity and the implications of model uncertainty neglect, consistent with our finding of full neglect of model uncertainty, they also find that many participants in their experiments build their belief formation process on only one model of the world.

4. More broadly, Bordalo, Gennaioli, and Shleifer (2012), Bushong, Rabin, and Schwartzstein (2021), Kőszegi and Szeidl (2013) formalize how contextual features steer attention and focus and hence influence behavior. (See Dertwinkel-Kalt et al. (2022) as well as Somerville (2022) for experimental tests.) Gabaix (2014), Kőszegi and Matějka (2020), Caplin, Dean, and Leahy (2019) formalize attention as a "top-down" process where decision-makers decide to limit attention to reduce the complexity of a problem.

provide formal frameworks as well as experimental evidence of how attention (and memory) patterns shape belief formation and information processing. Ba, Bohren, and Imas (2022) show that attention processes can lead to overreaction to information. Esponda, Oprea, and Yuksel (2023) show that a form of representativeness heuristic has important implications in contexts of statistical discrimination. Enke and Zimmermann (2019), Enke (2020), and Graeber (2023) provide evidence that people systematically fail to attend to key aspects of the information environment when processing new information.[5] We document a related phenomenon in the context of model uncertainty where people selectively attend to only one model of the world.

Our work also relates to a literature that studies the effect of complexity on attention. Recent research on complexity has made substantial progress in defining and quantifying complexity and studying implications in different decision contexts (e.g., Oprea (2020), Kendall and Oprea (2024), Enke and Shubatt (2024), Shubatt and Yang (2024)), Arrieta and Nielsen (2024). A common finding is that limited attention on a subset of relevant decision parameters is a simplification response to complexity (Enke (2024), Ba, Bohren, and Imas (2022), Enke (2020), Enke and Zimmermann (2019), Graeber (2023)). We show that in the presence of model uncertainty, complexity induces people to fully neglect model uncertainty. Furthermore, we document that due to the simplification response of uncertainty neglect, confidence in the optimality of own actions is higher in the high-complexity conditions compared to the low complexity conditions.

The rest of the chapter is structured as follows. In Section 2.2, we describe the baseline experimental design. Section 2.3 presents the results on how model complexity leads to the simplification of model uncertainty in actions. Section 2.4 then delves into implications of this simplification, studying beliefs and cognitive uncertainty. In Section 2.5 we present a short model that can generate the observed pattern of results, before concluding in Section 2.6.

2.2 Baseline Experimental Design

We designed our experiment with the following goals in mind: (i) implement an economically meaningful and somewhat natural decision environment that allows us to study how people deal with model uncertainty in the face of model complexity; (ii) achieve a well-defined notion of model uncertainty that allows us to exogenously vary model complexity in a straightforward way and (iii) being able to infer the weighting of competing models through the measurement of actions.

5. Hartzmark, Hirshman, and Imas (2021) show that ownership leads to overreaction to information, an effect that is driven by channeled attention on information. Augenblick, Lazarus, and Thaler (forthcoming) and Fan, Liang, and Peng (2024) instead focus on the role of cognitive noise and similarity patterns, respectively, in explaining over- and underreaction to information.

The Task and Model Uncertainty. We chose a financial decision-making task. In the experiment, respondents had to estimate the value of 8 fictitious companies. The correct company value was determined by one of two models, "The CEO is key" or "Products are crucial". The two models used different variables as inputs. One of the models was correct, meaning that it provided the correct company values for all 8 companies, while the other model produced uninformative values.

Each of the two models consists of a formula to calculate the proposed values. "The CEO is key" had the variables CEO competence C and Supporting Staff S as inputs, which could be used to calculate the company value as $C \times S - C - S + 10$. "Products are crucial" had the inputs number of products P and research cost R, and the company value was given by $P \times (10 - R) + R - P$.

To implement model uncertainty, the correct rule was determined in secret by the computer with a simulated coin flip. Therefore, without any additional information, the probability that either rule produced the correct company values was 50%. Respondents then received a noisy but informative signal about the correct model. This signal corresponded to the truth with a probability of 65%.[6] This was visualized using a ball drawn from an urn containing 65 balls with the correct model and 35 balls with the incorrect model. Afterwards, to measure actions, respondents were asked to provide value estimates for the 8 companies, each featuring different sets of variable realizations.[7]

Company value estimates were incentivized through a binarized scoring rule.[8] In this way, the chance of respondents to win a bonus was maximized by stating their best guess. Danz, Vesterlund, and Wilson (2022) document empirically that the binarized scoring rule can lead to systematic bias. Notice that such bias (if present in our setting) would not compromise our identification which relies on the comparison of value estimates between conditions of high and low complexity. Furthermore, our results are robust to using investment behavior rather than value estimates as an outcome, which features a different incentive scheme (see Section 2.3.3).

Complexity Manipulation. We varied the implementation complexity of the two models between-subjects. In treatment condition *HighComplexity*, working with the models was complex. Specifically, respondents had to calculate the company values under both models themselves. To obtain the Bayesian company value guess,

6. We did not provide any feedback between rounds. Hence, the only information respondents receive about which model is correct is the signal described above.

7. The variable realizations were integers between 0 and 10 and all implemented configurations were selected to yield company values between 10 and 100.

8. Every tenth participant was eligible to receive a bonus payment, in which case a random decision from the survey was incentivized. If a company value guess was incentivized, respondents received \$10 with a probability (in percent) of $100 - 100 \times ({}^{\text{Truth}}/_{100} - {}^{\text{Guess}}/_{100})^2$, where Truth is the true company value and Guess the company value guess.

they then needed to weight both values by the respective probabilities of 65% and 35%. In treatment condition *LowComplexity*, they were provided with the calculated company values for both models. Hence, respondents only needed to combine and weight the two values to make their guess, which significantly reduced the complexity of working with the models.[9]

Cognitive Uncertainty and Beliefs about Model Uncertainty. We measured cognitive uncertainty, i.e., people's confidence in the optimality of their value estimates similar to Enke and Graeber (2023). Specifically, after the series of 8 value estimation tasks, we asked people: *'How certain are you that, on average, your guesses were no more than 10 points away from the best possible guess given the information you received?'* Respondents indicated their answer on a scale from 0 to 100 percent. The cognitive uncertainty measure was only implemented in a subset of experiments (see Table 2.A.1 and Section 2.4.2) and was not incentivized. We conducted an additional experiment that features an incentivized confidence measure where participants place a bet on the optimality of their actions (see Section 2.4.2)).

In the last part of the experiment, respondents had to guess the probabilities that either of the two models generated the correct company values. Since respondents received a noisy signal during the company valuation task, this corresponds to the standard bookbag and chips belief updating task with a signal precision of 65% (cf. Benjamin (2019)). The stated belief was incentivized using a binarized scoring rule.[10] Participants were also asked an un-incentivized direct recall question where they were asked to state which rule was indicated to be more likely by the signal.

Design Details and Procedures. Respondents in both treatment conditions initially received the same set of instructions, explaining both treatments. Afterwards, they went through compulsory comprehension checks and two test runs, where they could practice applying each of the model formulas. After they completed the test runs, they were randomly assigned their treatment, received the noisy signal, and subsequently completed the 8 rounds of the task.

On the decision screens, respondents had to hover their mouse over the name of the respective model to uncover the variables and formula or calculated company value (cf. Ba, Bohren, and Imas (2022)). This allows us to study the attention paid to both of the models and signal realizations.

The experiments in this chapter were pre-registered on the platform AsPredicted. The pre-registrations include the experimental design, hypotheses, analy-

9. Indeed, average guessing times in the Baseline experiment were significantly shorter at 17.09 seconds in *LowComplexity*, compared to 49.11 seconds in *HighComplexity* ($p<0.001$).

10. If a probability guess was incentivized, respondents received \$10 with a probability of $100 - 100 \times (\text{Truth}/100 - \text{Guess}/100)^2$ %, where Truth is either 100 or 0, and Guess is the stated probability between 0 and 100.

ses, sample sizes, and exclusion criteria. Table 2.A.1 provides links to each preregistration.

We conducted the experiments online using the survey provider Prolific. Respondents were recruited from the United States and restricted to be fluent in English. To qualify for the survey, participants had to pass comprehension checks after reading the instructions. The experimental instructions can be found in Appendix 2.A.4. The median completion time in the Baseline Experiment was 22 minutes. Respondents received a fixed payment of $4 for the initial study. Every tenth respondent had the chance of winning an additional bonus of up to $10.

As preregistered, we focus our analysis on two different samples. After reading the instructions, but before treatment assignment, we presented respondents with two test runs for how to calculate the estimates of the models under high complexity (as described above). There, they could familiarize themselves with how to calculate the estimates under both decision models. In order to ensure that we have a respondent pool that is in principle able to solve the formulas in the high-complexity condition, we pre-registered to restrict our sample to respondents who correctly answered both of the two example decisions. The restricted sample of the Baseline Experiment includes 230 respondents. All figures and tables in the main text and Appendix 2.A.1 are based on the restricted sample. Appendix 2.A.2 reproduces all exhibits using the also pre-registered more lenient sample that only requires one of the two questions to be answered correctly, featuring 319 respondents for the Baseline Experiment.

2.3 Complexity and Simplification of Model Uncertainty

2.3.1 Framework and Hypothesis

Take a decision-maker whose task is to state a value estimate g_{actual}. Recall that in each decision scenario, the correct company value corresponds to one of the two values given by the models "The CEO is key" and "Products are crucial". The noisy signal received by the participant then indicates the model that is more likely to deliver the correct value. We call the signal-consistent value $v_{consistent}$, and the signal-inconsistent value proposed by the other model $v_{inconsistent}$. Since the signal reveals the correct rule with a probability of 65%, the rational guess for the company value is given by

$$g_{rational} = 0.65 \times v_{consistent} + 0.35 \times v_{inconsistent}.$$

Now take a decision-maker who seeks to simplify the world by neglecting model uncertainty. Such a decision-maker will base their value estimates exclusively on the more likely model. The naive benchmark that fully neglects model uncertainty

and takes the signal at face value is hence given by

$$g_{naive} = v_{consistent}.$$

As pre-registered, the main statistical measure we employ is the naive weight λ implicitly defined by

$$g_{actual} = \lambda \times g_{naive} + (1 - \lambda) \times g_{rational}, \tag{2.1}$$

which can be rearranged to obtain

$$\lambda = \frac{g_{actual} - g_{rational}}{g_{naive} - g_{rational}}. \tag{2.2}$$

The naive weight λ equals 1 if a participant states the naive guess (full neglect of model uncertainty) and 0 if they state the rational guess.

We expect that decision-makers will be more prone to simplify the world if complexity is high. Therefore we state the following, pre-registered, hypothesis:

Hypothesis 2.1. Decision-makers are more likely to neglect model uncertainty when complexity is high, compared to when complexity is low.

2.3.2 Results

Figure 2.1a plots histograms of decision-level naive weights in our main sample for both treatments. The distribution in *HighComplexity* features more mass at 1 than the one in *LowComplexity*, indicating a larger amount of fully naive guesses and overreaction to information in the former condition. Conversely, the distribution for *LowComplexity* has more mass around 0 than for *HighComplexity*, implying more guesses in close proximity to the rational benchmark. The prevalence of fully naive guesses is 80% in *HighComplexity*, compared to 59% in *LowComplexity* ($p < 0.01$). Hence, the vast majority of guesses is fully naive in *HighComplexity*.[11] Figure 2.1b confirms this pattern by comparing the average naive weight across treatments, finding a significantly larger average naive weight under complexity ($p < 0.01$).

Table 2.1 complements this analysis by regressing the company value guesses on the rational and naive benchmarks. This can be interpreted as estimating Equation 2.1 without the restriction that the weights on the rational and naive benchmarks add up to 1. We can see that the fully naive benchmark is relatively more predictive of participants' guesses in treatment *HighComplexity* compared to *LowComplexity*.

Result 2.1. There is substantially more neglect of model uncertainty when complexity is high, compared to when it is low.

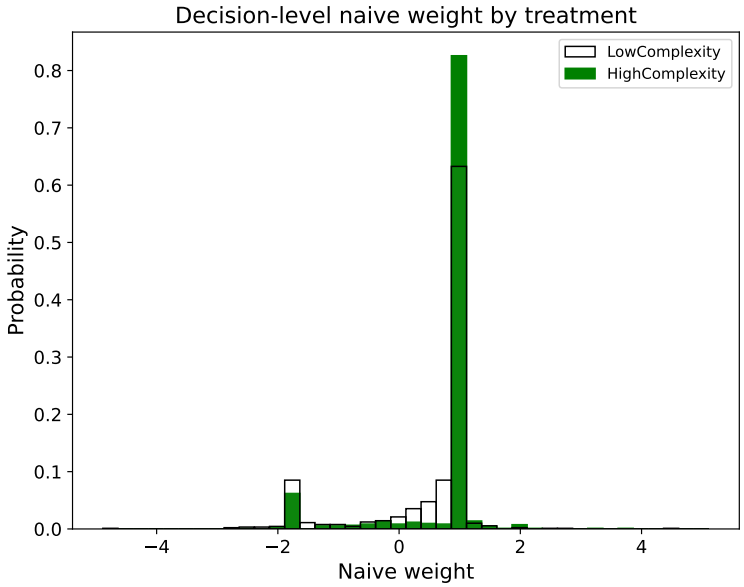

(a) Distribution of decision-level naive weights

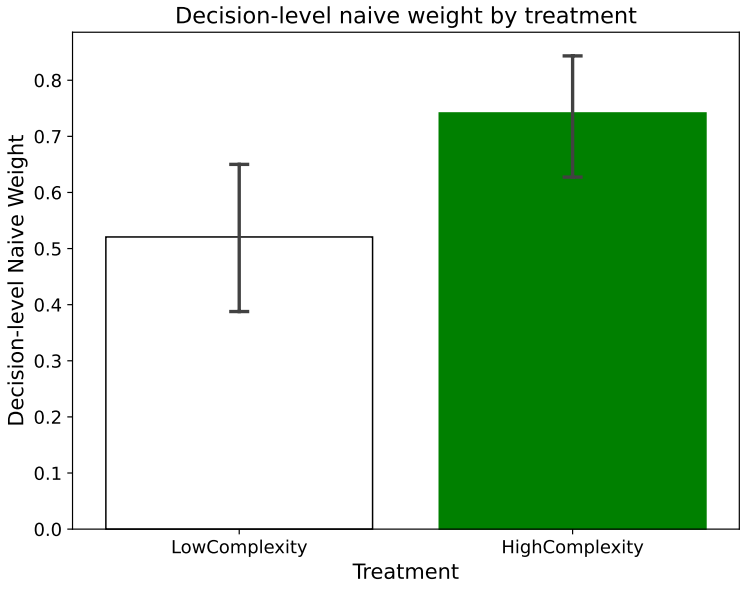

(b) Mean decision-level naive weights

Figure 2.1. Decision-level naive weights in LowComplexity and HighComplexity conditions. *Panel (a) plots the distribution of naive weights λ calculated as specified in Equation 2.2, using the restricted sample of the Baseline Experiment with 230 participants. Panel (b) plots average naive weights.*

Table 2.1. Company Value Guesses

Dependent variable:	Company Value Guess		
Sample:	LowComplexity (1)	HighComplexity (2)	Pooled (3)
Rational Benchmark	0.593***	0.343***	0.593***
	(0.079)	(0.065)	(0.079)
Naive Benchmark	0.481***	0.703***	0.481***
	(0.070)	(0.057)	(0.070)
Rational B. × HighComplexity			-0.250**
			(0.102)
Naive B. × HighComplexity			0.222**
			(0.090)
R^2	0.881	0.918	0.900
Observations	904	936	1840

The table presents OLS regressions of respondents' company value guesses on the rational and naive benchmarks as detailed in Section 2.3.1. The table is based on the restricted sample of the Baseline Experiment. Column (1) uses observations from the LowComplexity treatment, column (2) from the HighComplexity treatment, and column (3) from both. Stars highlight significant differences from 0 with * for $p < 0.10$, ** for $p < 0.05$, *** for $p < 0.01$. Standard errors are clustered on the subject level.

We note that the naive weight defined in Equation (2) is quite sensitive to outliers. To ensure that our results are not driven by outliers, we look at the 8 decision scenarios for each subject and compute the median naive weight. In an analysis that was not preregistered, we plot the histogram of median naive weights for each subject in Figure 2.A.1. We again observe a higher mass around naive weights of 1 in *HighComplexity*, and more mass between 0 and 1 in *LowComplexity*, confirming the earlier results that full neglect of model uncertainty is more prevalent under complexity.

Appendix 2.A.2.1 replicates all exhibits using the also pre-registered more lenient sample that only requires one of the two questions to be answered correctly. In Appendix 2.A.1.7 we show that we replicate all the findings for the Baseline Confidence Experiment (see Section 2.4.2 and Table 2.A.1).

Role of Attention. The above results indicate that people simplify complexity by neglecting model uncertainty. To further corroborate this finding, we study the role of attention. Recall that respondents had to hover their mouse over either "The CEO is key" or "Products are crucial" to observe the parameters needed to make company value guesses. In *LowComplexity*, the company value was displayed only when hovering over the respective rule. In *HighComplexity*, the respective variable realizations needed to calculate the company value were displayed. The resulting data on hover times allows us to study how much attention participants paid to either rule.

Figure 2.2 shows the distribution of the median consistent hovering shares per subject separately for both treatments. In treatment *HighComplexity*, an overwhelming majority of participants only considers the variables of the signal-consistent model. In treatment *LowComplexity*, most participants consider the proposed values by both models, with the mode being at equal hovering times for both the signal-consistent and inconsistent model.

Table 2.2 analyzes more formally how attention differs between the two treatments. The first column shows that the average hover time for the signal-consistent rule more than triples with higher complexity. The second column shows that the hover time for the signal-inconsistent rule also increases, but by much less. The third column confirms that the share of consistent hover time increases in *HighComplexity*, meaning that participants in this condition pay relatively more attention to the signal-consistent rule.

2.3.3 Robustness

Investment Behavior. In the baseline experiment, our main outcome measure is respondents' value estimates for the hypothetical companies. We ran two additional

11. A large fraction of guesses in *LowComplexity* also reveal full naivete, which may be caused by the residual complexity of the general experimental set-up or more specifically of the need to combine the signal-consistent and signal-inconsistent values into a value estimate.

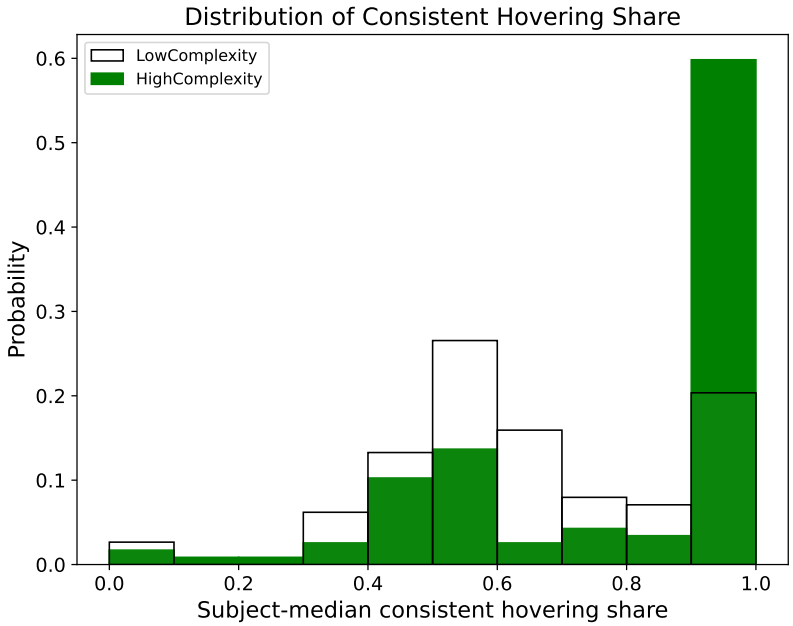

Figure 2.2. Distribution of subject-medians of the consistent hovering shares. *The figure plots the distribution of the share of time that respondents spent looking at the values of the signal-consistent rule, using the restricted sample of the Baseline Experiment with 230 participants. Only the median consistent share for each participant is plotted.*

Table 2.2. Hover Times

Dependent variable:	Consistent Hover Time	Inconsistent Hover Time	Consistent Share
Sample:	Pooled (1)	Pooled (2)	Pooled (3)
Constant	2.641***	1.533***	0.635***
	(0.232)	(0.153)	(0.018)
HighComplexity	11.611***	1.588***	0.153***
	(0.954)	(0.392)	(0.027)
R^2	0.262	0.033	0.079
Observations	1840	1840	1840

The table presents OLS regressions using the restricted sample of the Baseline Experiment. Hover times were winsorized at the top at the 97.5% quantile. All columns use observations from both the HighComplexity and LowComplexity conditions. In column (1), the time that respondents spent looking at the values for the signal-consistent rule is regressed on a constant and a treatment dummy for the HighComplexity condition. In column (2), the dependent variable is the time spent looking at the signal-inconsistent rule. In column (3), it is the share of time that was spent looking at the signal-consistent rule. Stars highlight significant differences from 0 with * for $p < 0.10$, ** for $p < 0.05$, *** for $p < 0.01$. Standard errors are clustered on the subject level.

preregistered versions of the Baseline Experiment that replace the company value estimate with an investment decision. The design of the experiments was exactly as described in Section 2.2, with only one change. Instead of providing a guess for the value of the company, participants were asked to submit an investment bid for each company. For each decision, they received a budget of 100 cents and could subsequently decide how much to bid for the company. The bids were incentivized through a random auction mechanism: a random price between 10 and 100 cents was drawn for the company. If the bid was greater than or equal to the price, the company was bought by the participant, paying the drawn price, and receiving the true company value. If the bid was lower than the price, there was no transaction.

Both experiments were pre-registered (see Table 2.A.1). To ensure comparability with the other experiments, for both experiments we focus on the results from the restricted sample. For experiment 1 this yields 193 participants, for experiment 2 this sample features 323 participants.[12]

Taken together, the results go in the same direction as for the Baseline Experiment. The overall prevalence of fully naive bids in experiment 1 is 73% in *HighComplexity* and 47% in *LowComplexity* ($p<0.001$). Figure 2.A.5b shows that the mean naive weight is significantly higher under complexity. Hover times show the same patterns as in the Baseline Experiment, with there being significantly more consistent hovering in *HighComplexity*. In experiment 2, the share of fully naive guesses is 70% in *HighComplexity* and 41% in *LowComplexity*, and the mean naive weight is also significantly higher under complexity as can be seen in Figure 2.A.9b, confirming the results from the first experiment. Again, hover times also show the same pattern as before (see Appendix 2.A.1.5).[13]

Result 2.2. Our results are robust to using a different outcome, namely investment behavior. People's investment bids reflect substantially more full neglect of model uncertainty when complexity is high, compared to when it is low.

Equally Likely Models. In the baseline experiment, we endow respondents with a natural candidate for simplification, namely the objectively more likely model. Hence, the question arises whether simplification also occurs if both models are equally likely. To address this, we conducted an additional preregistered study, where participants completed a version of the Baseline Experiment that excluded the informative signal. Instead, they were only endowed with a 50-50 prior when

12. Notice that for experiment 1, we deviate from the specification in the pre-registration where we pre-registered the use of the lenient sample and the full sample. Appendix 2.A.2.3 and Appendix 2.A.2.4 produce the corresponding results. Also notice that for experiment 1 we pre-registered a smaller sample size than for the other experiments.

13. Appendices 2.A.2.3 and 2.A.2.5 produce results for the more lenient sample. Overall, results are similar to the restricted sample, although the treatment difference in mean naive weights in experiment 1, while directionally present, fails to be significant in the pre-registered lenient sample ($p=0.207$).

estimating company values. The link to the preregistration can be found in Table 2.A.1.

The experiment followed the design described in Section 2.2 but omitted both the signal and the belief elicitation regarding the rule determining company values. To maintain the experiment's length and incentive structure, we replaced the belief elicitation with an unrelated belief-updating task.

In the absence of a noisy signal, there is no clear reason to expect participants to simplify model uncertainty in a specific direction when faced with complexity. Therefore, our hypothesis for this study was that participants in the *HighComplexity* condition would more frequently state guesses equal to the values proposed by either rule, whereas more intermediate guesses would be observed in the *LowComplexity* condition.

To test this, we examine the implicit decision weight, γ, assigned to the company value proposed by the "The CEO is key" rule, defined as

$$g_{actual} = \gamma \times v_{CEO} + (1-\gamma) \times v_{Products}. \tag{2.3}$$

Additionally, we define weight extremity as

$$\text{Weight extremity} = |\gamma - \frac{1}{2}|. \tag{2.4}$$

Our pre-registered hypothesis then is that weight extremity is higher in the *HighComplexity* condition compared to the *LowComplexity* condition.

The pre-registered restricted sample consists of 348 participants who correctly solved both HighComplexity example decision screens.

Figure 2.3a displays the distribution of decision-level CEO weights for both treatments. In the *LowComplexity* condition, the distribution has greater mass at intermediate values, particularly near the rational CEO weight of 0.5. In contrast, the *HighComplexity* distribution shows more mass at weights corresponding to simplified guesses, especially at a CEO weight of 1. Figure 2.3b further confirms that the average weight extremity in *HighComplexity* is significantly higher than in *LowComplexity*.

Figure 2.A.2 displays the distribution of the share of time participants spent hovering over the "The CEO is key" model in both treatments. Participants in *HighComplexity* are more likely to focus on a single model, whereas those in *LowComplexity* typically attend to both models.

In Appendix 2.A.2.2, we replicate the analysis using the more lenient sample of 452 participants who answered at least one of the two HighComplexity example screens correctly. All results continue to hold in this sample.[14]

Result 2.3. People also neglect model uncertainty as a response to complexity when both models are equally likely.

14. A perhaps interesting follow-up question is whether people that simplify model uncertainty as a response to complexity do so in a consistent fashion, always focusing on the same model. To

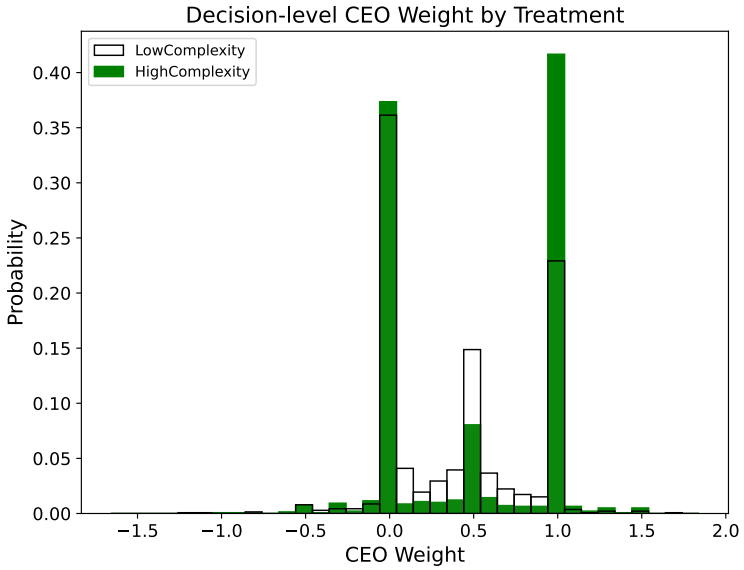

(a) Distribution of decision-level CEO weights

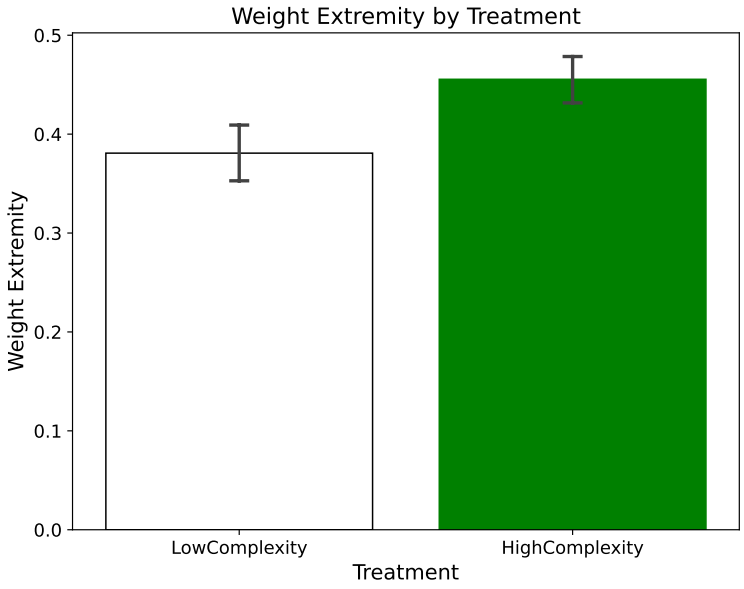

(b) Mean decision-level weight extremity

Figure 2.3. Decision-level CEO weights and weight extremity in LowComplexity and HighComplexity conditions of the Equally Likely Models Experiment. *Panel (a) plots the distribution of CEO weights γ calculated as specified in Equation 2.3, using the restricted sample of the Equally Likely Models study with 348 participants. Panel (b) plots the average weight extremity* $|\gamma - \frac{1}{2}|$ *calculated as specified in Equation 2.4.*

2.4 Results: Implications of Simplification

2.4.1 Beliefs about Models

We have shown that model complexity leads to a simplification response: people tend to fully neglect model uncertainty in their actions. As pre-registered, we now ask whether this simplification also causes a distorted view of reality, namely that people misperceive model uncertainty when directly asked.[15]

For this purpose, beliefs about which model is correct were elicited immediately after all company value guesses had been submitted.

Figure 2.4a shows histograms of signed probability guesses (signed in the direction of the received signal), separately for treatments *LowComplexity* and *HighComplexity* in the Baseline experiment. There are no visible differences in beliefs. This is confirmed by Figure 2.4b which plots average guesses for the probability that "The CEO is key" is the correct model, stratified by the signal participants received and their assigned treatment. There is no significant treatment difference in average beliefs for either signal.

Table 2.3 presents regressions for signed probability guesses and correct recall of the signal. The regressions confirm that there are no significant treatment differences in beliefs. Similarly, there are no differences in the accuracy of recall of the received signal, as column (2) reveals.

We also investigate belief patterns in the Investment 1, Investment 2 and Baseline Confidence experiments. In the restricted samples of all of these studies, we find the same result of no significant treatment differences in beliefs and recall, as can be seen in the respective sections of Appendix 2.A.1.[16] Taken together, these results suggest that the complexity-induced neglect of model uncertainty in actions does not translate into corresponding belief patterns.

Result 2.4. The neglect of model uncertainty in actions does not translate to beliefs about model uncertainty. These beliefs appear to reflect model uncertainty and are largely unaffected by complexity.

---

shed some light on this, in an analysis that was not preregistered, Figure 2.A.3 restricts the sample to participants in the *HighComplexity* condition who made only fully naive guesses—that is, who consistently guessed a value corresponding to either the "The CEO is key" or "Products are crucial" rule for every decision. This applies to 91 out of all 174 participants in *HighComplexity*. The figure suggests a large fraction of respondents (about half) consistently simplify in the same direction, always choosing values from the same model, while a sizeable fraction alternates between models.

15. Notice that, while we pre-registered this analysis, we did not pre-register a specific hypothesis.

16. In Appendix 2.A.2 we present the results for the lenient samples. Here, we find a (marginally) stronger belief response under complexity in the Baseline and Baseline Confidence studies, and no significant treatment differences in Investment Experiment 1 and Investment Experiment 2.

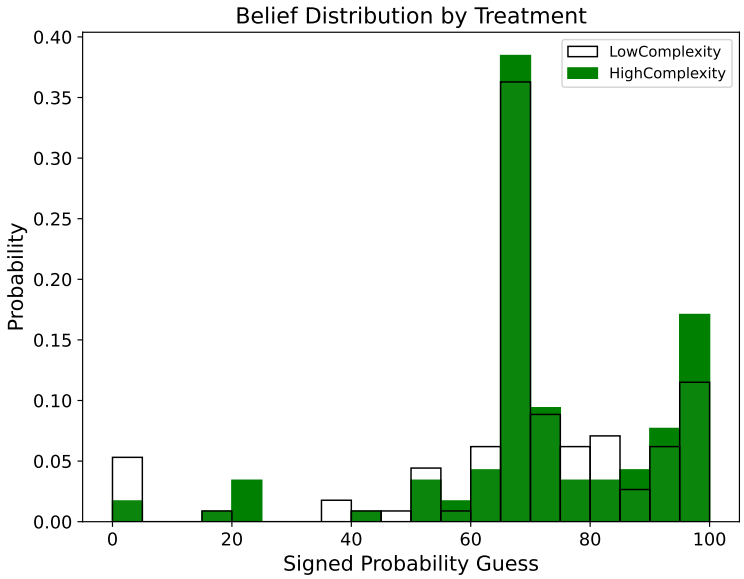

(a) Histogram of signed probability guesses.

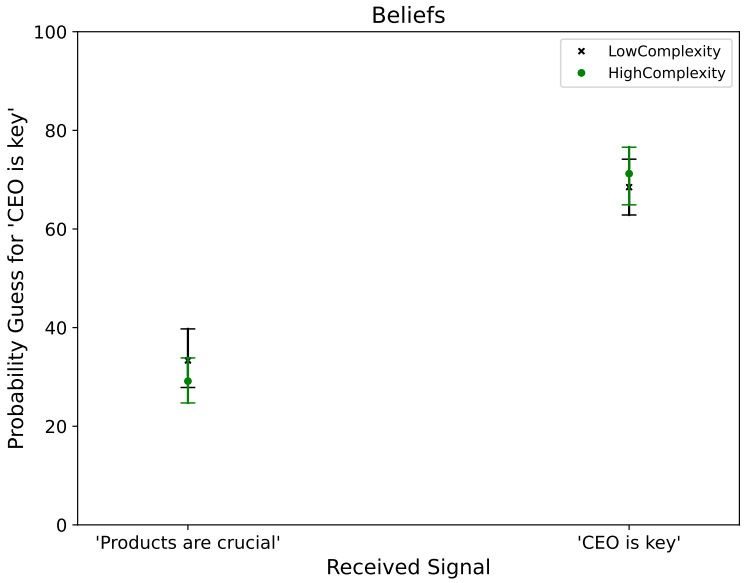

(b) Mean probability guesses for "The CEO is key".

Figure 2.4. Beliefs in the Baseline Experiment. *In Panel (a) beliefs are converted into the direction of the more likely model, so that 65 corresponds to the Bayesian probability. Panel (b) plots mean probability guesses for the model "The CEO is key". The figure is based on the restricted sample of the Baseline Experiment with 230 participants.*

Table 2.3. Beliefs and Recall in the Baseline Experiment

Dependent variable:	Probability Guess	Correct Recall
Sample:	Pooled	Pooled
	(1)	(2)
Constant	67.679***	0.965***
	(2.078)	(0.018)
HighComplexity	3.351	-0.007
	(2.819)	(0.026)
R^2	0.006	0.000
Observations	230	230

The table presents OLS regressions using the restricted sample of the Baseline Experiment. In column (1), the dependent variable is the probability guess for the likely state. Column (2) uses a dummy for whether respondents correctly recall the more likely model. All columns use observations from both the HighComplexity and LowComplexity conditions. Stars highlight significant differences from 0 with * for $p < 0.10$, ** for $p < 0.05$, *** for $p < 0.01$. Standard errors are clustered on the subject level.

Robustness: Delayed Belief Elicitation. In an additional pilot study that was not preregistered (see Table 2.A.1), we investigated whether adding a delay between actions (value guesses) and belief elicitation about model uncertainty could potentially induce treatment differences in beliefs. The rationale was that it might take time for simplification-induced misperceptions to form. The design was exactly as described in Section 2.2, with an additional second part that took place one day after the initial survey. Respondents in the *Immediate* condition stated their beliefs immediately after the company value guesses as in the Baseline Survey, and completed unrelated tasks during the second survey. Respondents in the *Delay* condition completed the company value guesses and unrelated tasks during the first part of the survey, and the belief elicitation during the second survey one day later.

We first note that also in this experiment, we see that complexity induces people to neglect model uncertainty in actions. However, as Appendix 2.A.1.6 reveals, both with and without delay, simplification does not induce biased beliefs about model uncertainty.

To summarize, adding a one day delay does not induce treatment differences both in stated beliefs and correct recall of the received signal.

Results on Belief-Action Link. Notice that our results imply a complexity-induced wedge between actions and beliefs. A recent literature has investigated the link between beliefs and actions (Giglio et al. (2021), Ameriks et al. (2020), Beutel

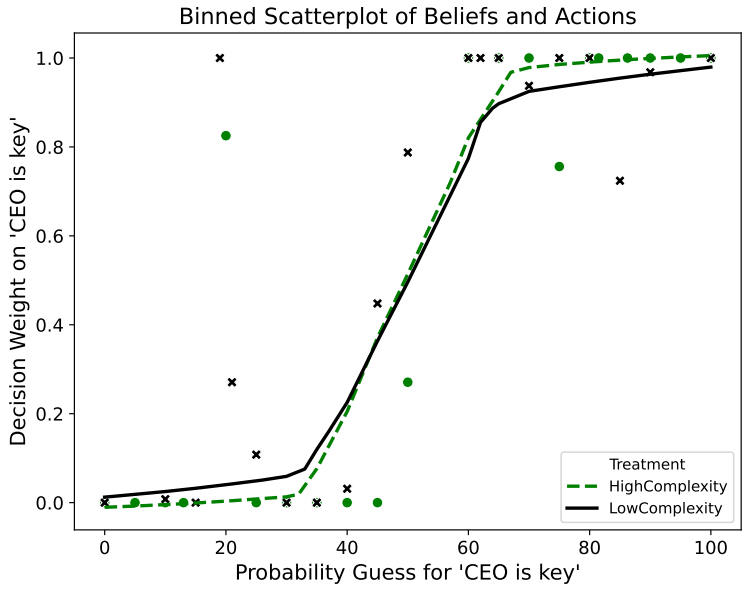

Figure 2.5. The relationship between decisions and beliefs. *The figure shows a binned scatterplot using the restricted sample of the Baseline Experiment with 230 participants. It has the probability guess that "The CEO is key" is the more likely model on the horizontal axis, and the decision weight γ as defined in Equation 2.3 on the vertical axis. The lines show LOWESS regressions based on all datapoints.*

and Weber (2023), Laudenbach et al. (forthcoming)), and the determinants of how strongly beliefs translate into actions. Charles, Frydman, and Kilic (2024) find that increased complexity of forming a belief weakens the transmission of beliefs into actions. Similarly, Yang (2023) as well as Enke et al. (2024) highlight a link between information processing constraints and a weak elasticity of decisions with respect to economic fundamentals.

To investigate this wedge in our experiment more formally, Figure 2.5 presents a binned scatterplot with beliefs on the horizontal axis, and corresponding actions on the vertical axis, using data from the Baseline Experiment. Beliefs are given by the probability guess for the "The CEO is key" model. To ensure comparability, actions are represented by the implicit decision weight γ on the company value proposed by the "The CEO is key" model, as defined in Equation 2.3.

The plot illustrates the overreaction in actions caused by the complexity of calculating the optimal guess: when moving away from a probability guess of 50%, the decision weights in *HighComplexity* quickly move to extreme values, while the response is more muted in *LowComplexity*. This implies that the effect of complexity on the belief-action link is not straightforward in our setting. In a neighborhood around 50%, complexity increases responsiveness of actions to beliefs. However, when moving to more extreme beliefs, complexity renders the action response

rather flat, since naive guessing is already triggered starting from fairly moderate beliefs.

We show analogous plots for the Investment 1, Investment 2 and Baseline Confidence studies in Appendix 2.A.1, as well as results for the lenient samples in Appendix 2.A.2, each showing qualitatively identical results.

Result 2.5. Complexity induces a wedge between beliefs and actions. When complexity is high, beliefs continue to reflect model uncertainty, actions tend to be based on neglect of model uncertainty.

2.4.2 Confidence

Here we investigate whether the complexity-induced simplification of neglect of model uncertainty in actions affects how people view the optimality of their actions. This type of confidence has been shown to explain a broad range of behavioral anomalies and also mediates to what extent individual biases matter for aggregate outcomes (Enke and Graeber (2023), Enke, Graeber, and Oprea (forthcoming), Enke, Graeber, and Oprea (2023)). A typical and highly intuitive finding in the literature is that complexity reduces confidence in action optimality.

Baseline Confidence. In a replication of the Baseline Experiment, we elicited cognitive uncertainty, i.e. people's confidence in the optimality of their value estimates. The experiment was pre-registered, including the analysis of the cognitive uncertainty measure (see Table 2.A.1). As pre-registered, we use the same sample restriction as for the Baseline study, resulting in a restricted sample with 336 participants.[17]

Figure 2.6 shows the average confidence levels by treatment condition. We find that participants in the *HighComplexity* condition are significantly *more* confident in their guesses than those in the *LowComplexity* condition—despite exhibiting greater simplification and a lower rate of rational guesses under complexity, a seemingly contradictory pattern.

We also elicited the same measure in the Equally Likely Models study and featured its analysis in the preregistration. Figure 2.A.4 confirms that the finding replicates in this study. Additional results presented in Appendix 2.A.2 show that the results in both studies also hold true when using the more lenient sample restrictions.

Incentivized Confidence. In another replication of the Baseline Confidence experiment, we further added an incentivized version of the confidence elicitation. Here, participants could bet on the optimality of their guesses. They received an endowment of $10 and could choose how much to bet (between $0 and $10) on the

17. Similar to the analysis of beliefs, while we pre-registered this analysis, we did not pre-register a specific hypothesis.

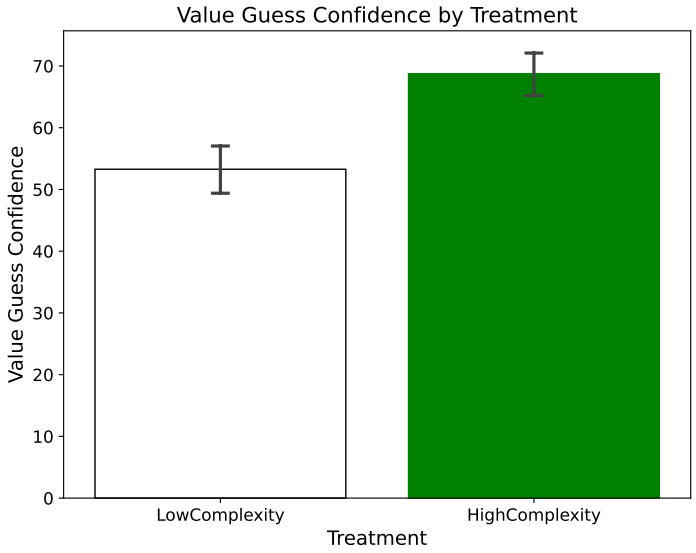

Figure 2.6. Average value guess confidence in the Baseline Confidence Experiment. *The figure plots the average confidence that respondents had in their company value guesses, using the restricted sample of the Baseline Confidence study with 336 participants.*

event that their guesses had been, on average, no more than 10 points away from the best possible guess. The bet was multiplied by 3 if their guesses had indeed been accurate and was lost otherwise. After completing all company value guesses, participants answered the non-incentivized (probability) and incentivized (bet) versions on the same screen. The experiment was pre-registered, including the analysis of the incentivized and non-incentivized cognitive uncertainty measures (see Table 2.A.1). As pre-registered, we use the same sample restriction as for the Baseline study, resulting in a restricted sample with 203 participants. We report the results on confidence for this experiment in the main text and refer the reader to Appendix 2.A.1.8 for the remaining analyses.

Figure 2.7 plots the average confidence measures. The lower panel replicates the previous result that participants in the high-complexity condition stated a higher confidence level in the optimality of their company value guesses than their counterparts in the low-complexity condition. The upper panel presents results from the incentivized betting task. While these data are naturally more noisy due to the betting context, the Figure again confirms our result that higher complexity yields higher confidence. Appendix 2.A.2.7 produces these results under the more lenient sample restrictions, yielding qualitatively identical results.

Taken together, our results indicate that confidence in the optimality of own actions is higher in the high-complexity conditions compared to the low complexity conditions. It seems that in contrast to prior results from different contexts, in the context of model uncertainty, the possibility to respond to complexity by greatly

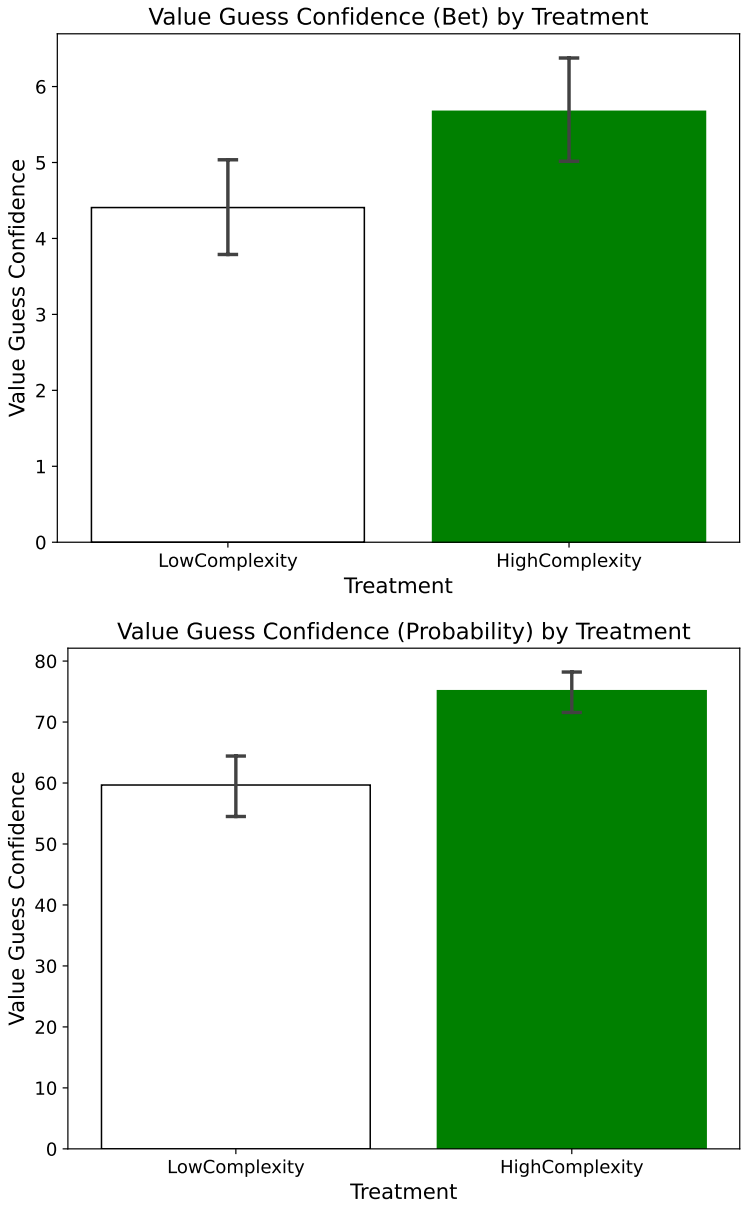

Figure 2.7. Average value guess confidence in the Incentivized Confidence Experiment. *The top figure plots the average incentivized confidence measure, while the bottom figure plots the non-incentivized measure, both using the restricted sample of the Incentivized Confidence study with 203 participants.*

simplifying the world through full neglect of model uncertainty leads to an illusion of certainty and hence increased confidence in action optimality.

Result 2.6. Complexity leads to higher confidence in the optimality of own actions.

2.5 A Simple Model of Representations

We have shown that respondents simplify model uncertainty when computational complexity is high. This, however, does not carry over to beliefs about model uncertainty. Finally and perhaps most surprisingly, higher computational complexity increases confidence in the optimality of own actions rather than decreasing it. In the final part of the chapter we present a simple model that can generate this pattern of results. The model is an augmented and simplified version of Bordalo et al. (2025). In the model, when faced with a decision problem, an agent first forms a mental representation of the problem. This process is shaped both by bottom-up and top-down attention. Upon being presented a decision problem, a bottom-up process of cue-dependent memory determines which of the currently stored mental representations is top of mind. Then, in a top-down process, the agent decides whether they want to further simplify this representation. This section contains the basic intuitions, while Appendix 2.A.3 presents the formal model.

This model formalizes economic decision making as a cognitively constrained process operating over a structured internal database of mental representations. Initially, this database comprises only a broad representation, which encodes a probabilistic assessment of the decision environment (e.g., a 0.65 likelihood for the more probable state) along with abstract contextual features (e.g., task framing, informational structure, and the need to estimate a company's value).

Each decision - whether it concerns an action, confidence judgment, or belief report - invokes a two-step cognitive process to form a mental representation of the decision problem:

(1) Bottom-Up Retrieval: Agents hold a database of mental representations. When faced with a decision, similarity-based recall (Bordalo, Gennaioli, and Shleifer (2020), Bordalo et al. (2023b), Enke, Schwerter, and Zimmermann (2024), Graeber, Roth, and Zimmermann (2024), Jiang et al. (forthcoming)) determines which representation is top of mind. In other words, the representation most similar to the current decision cue is retrieved from memory.

(2) Top-Down Simplification: Once a representation is top of mind, the agent, in a top-down process, evaluates whether the cognitive costs of computing a response using the retrieved representation exceed the benefits. If so, the agent further simplifies the representation to reduce processing demands.

The key insight from our model is that people make decisions within the mental representation they formed for this specific decision problem. Hence, specific

decisions (actions, confidence judgments, or belief reports) in a given underlying environment may be based on very different mental representations of that environment, allowing us to explain our seemingly contradictory pattern of results.

Estimation of Company Values. Upon observing a signal, the agent must compute an optimal guess. Since this is the first decision agents take, the database of representations only contains the broad representation of the problem. Under low complexity, cognitive costs are manageable, so participants compute the optimal guess using the full broad representation of the problem, i.e., fully taking model uncertainty into account. Under high complexity, cognitive costs may exceed the benefits, prompting simplification to a mental representation in which model uncertainty is fully neglected. Once a decision has been taken, the corresponding representation is added to the database, including their contextual features.

Confidence Elicitation. In the low complexity treatment, when faced with the confidence assessment, agents will retrieve the broad representation (their database only consists of broad representations). As confidence judgments are cognitively light, no further simplification occurs. In the high complexity treatment, similarity-based retrieval favors the recall of the simplified representation. This is because the confidence elicitation cues the estimation tasks (it explicitly asks for confidence about the estimation task). As described above, confidence in own action optimality is then assessed by agents within this mental representation. Consequently, confidence can be elevated by complexity if confidence within the simplified high complexity environment is higher than in the full representation of the low complexity environment.[18]

Belief Elicitation. In the low complexity treatment, the broad representation is again retrieved and used without simplification, yielding probabilistically grounded belief reports about model uncertainty. In the high complexity treatment, retrieval of mental representations favors the broad representation over the simple one, since the belief task directly asks about model uncertainty and is hence more similar to the broad representation. Since the cognitive costs of belief elicitation are negligible, no further simplification occurs. People then state the beliefs of their mental representation and therefore beliefs in both treatments will tend to reflect model uncertainty.

18. While this seems plausible, our model does not formalize why confidence within the simplified high complexity environment may be higher than in the full representation of the low complexity environment. The key insight from our model is that confidence is assessed within the mental representation of the problem.

2.6 Conclusion

This chapter explores how individuals navigate model uncertainty in economic decision-making and demonstrates that complexity significantly influences the way people deal with model uncertainty. Through a controlled experimental framework, we investigate whether individuals account for model uncertainty when making decisions or instead simplify the world by implicitly assuming one model is correct. Our findings reveal that when model complexity is high, individuals tend to neglect model uncertainty in their actions, behaving as if one model is definitively correct. However, this neglect of uncertainty does not translate into distorted beliefs, as participants' stated beliefs continue to reflect model uncertainty. This creates a systematic wedge between actions and beliefs. Furthermore, our results show that complexity-induced simplification leads to increased confidence in decision optimality, contradicting prior findings that suggest complexity typically raises cognitive uncertainty.

Our results provide a direct test of the widely held assumption in the theoretical literature on misspecified models that people attend to a single model when making decisions, rather than entertaining multiple weighted models simultaneously.

By systematically manipulating the complexity of models, we provide robust evidence that individuals simplify complex problems by focusing on a single mental model. Data on attention allocation further support this conclusion, showing that participants in high-complexity conditions spend more time attending to one model, rather than considering both models equally. This attention-based mechanism helps explain why complexity amplifies the tendency to neglect model uncertainty in actions. The robustness of these findings is confirmed across different tasks, including an investment decision context, demonstrating that the observed pattern extends beyond the specific valuation task used in our primary experiment.

Our findings contribute to a deeper understanding of how complexity shapes economic cognition. Existing research on mental models has largely focused on selection and persuasion, whereas our study highlights a novel dimension: the impact of complexity on how people handle competing models. Our findings reveal that, when complexity is high, in order to be able to operate and make decisions, people need to simplify the world by acting as if only one model exists. Interestingly, this complexity-induced simplifications can lead to an illusion of certainty, where people are confident about the optimality of their actions.

While we do not study persuasion directly, an intuitive implication of our results may be that, when complexity is high, the presence of model uncertainty might make decision-makers more susceptible to persuasive narratives that present a single, seemingly definitive interpretation of economic realities.

Appendix 2.A Appendix

2.A.1 Additional Results for the Restricted Samples

2.A.1.1 Overview of Data Collections

Table 2.A.1. Overview of Data Collections

Collection	Participants	Description	Link to pre-analysis plan
Baseline	600	As described in Section 2.2. Treatments: HighComplexity and LowComplexity. Outcomes: Company value guesses, hover times, beliefs.	https://aspredicted.org/t79h-2mkj.pdf
Equally Likely Models	600	As Baseline, but without noisy indication, hence 50-50 belief about more likely model, and with additional guess confidence measure. Treatments: HighComplexity and LowComplexity. Outcomes: Company value guesses, guess confidence, hover times.	https://aspredicted.org/92cr-9kwd.pdf
Investment 1	400	As Baseline, but company value guesses replaced by investment decisions. Treatments: HighComplexity and LowComplexity. Outcomes: Investment decisions, hover times, beliefs.	https://aspredicted.org/4c38-7psb.pdf
Investment 2	600	As Investment, but larger sample size. Treatments: HighComplexity and LowComplexity. Outcomes: Investment decisions, hover times, beliefs.	https://aspredicted.org/7vcv-gq8z.pdf
Delayed Belief Elicitation	588	As Baseline, but additional variation the timing of belief elicitation: For half the respondents, beliefs are elicited immediately as in Baseline, for the other half they are elicited with a one-day delay. Treatments: (HighComplexity, LowComplexity) × (Immediate, Recall). Outcomes: Company value guesses, hover times, beliefs.	Not preregistered
Baseline Confidence	600	As Baseline, but with additional guess confidence measure. Treatments: HighComplexity and LowComplexity. Outcomes: Company value guesses, guess confidence, hover times, beliefs.	https://aspredicted.org/v6cv-y9yh.pdf
Incentivized Confidence	600	As Baseline Confidence, but with additional incentivized guess confidence measure on the same page as the non-incentivized measure. Treatments: HighComplexity and LowComplexity. Outcomes: Company value guesses, guess confidence, hover times, beliefs.	https://aspredicted.org/txsk-3hky.pdf

This Table provides an overview of the different data collections. The sample sizes refer to the size of the original data collection, prior to applying exclusion restrictions.

2.A.1.2 Baseline Experiment: Additional Results in Restricted Sample

Here, we present additional results for the restricted sample of the Baseline Experiment, featuring 230 participants who solved both of the example screens.

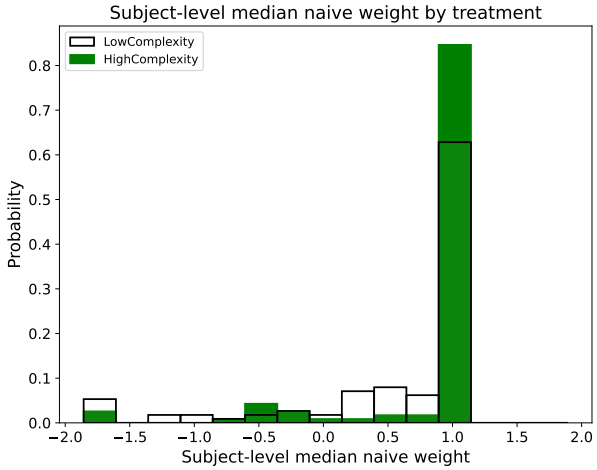

Figure 2.A.1. Distribution of median naive weights, computed for each subject. *The figure plots the distribution of naive weights λ calculated as specified in Equation (2), using the restricted sample of the Baseline Experiment with 230 participants. Only the median naive weight for each participant is plotted.*

2.A.1.3 Equally Likely Models: Additional Results in Restricted Sample

Here, we present additional results for the restricted sample of the Equally Likely Models Experiment, featuring 348 participants who solved both of the example screens.

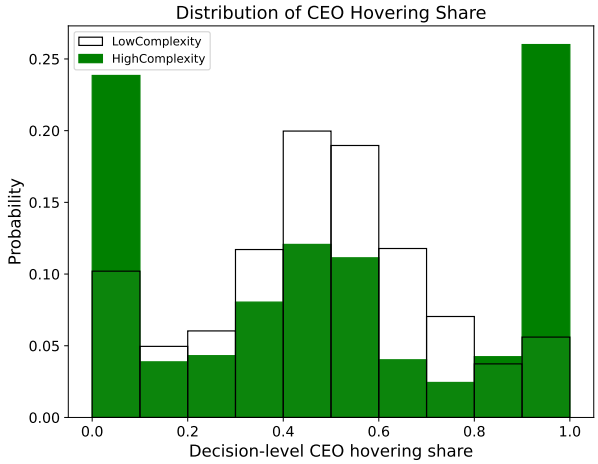

Figure 2.A.2. Distribution of decision-level CEO hovering share in the Equally Likely Models Experiment. *The figure plots the distribution of the share of time that respondents spent looking at the values of the "The CEO is key" model, using the restricted sample of the Equally Likely Models study with 348 participants. Hovering shares are plotted separately for each of the eight guessses made by respondents*

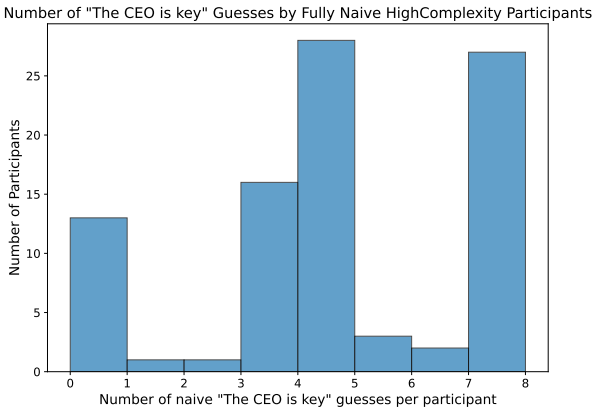

Figure 2.A.3. Distribution of naive decision-making in the Equally Likely Models Experiment. *The figure plots the distribution of the number of times participants selected the value corresponding to the "The CEO is key" model, using the restricted sample of the HighComplexity treatment of the Equally Likely Models study, limited to the 91 participants in the HighComplexity condition who made only fully naive guesses, i.e. who always selected a value corresponding to either the "The CEO is key" or "Products are crucial" model.*

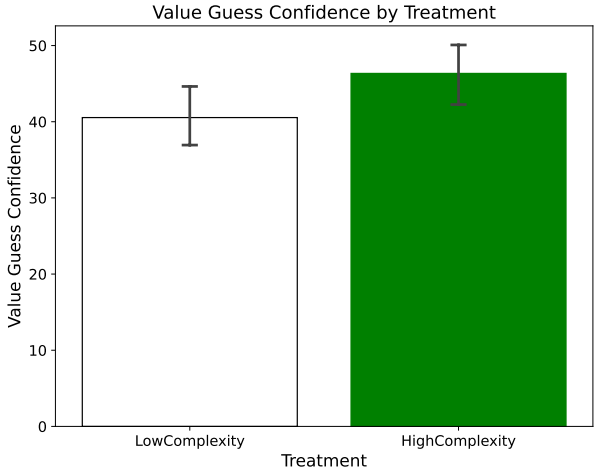

Figure 2.A.4. Average value guess confidence in the Equally Likely Models Experiment. *The figure plots the average confidence that respondents had in their company value guesses, using the restricted sample of the Equally Likely Models study with 348 participants.*

2.A.1.4 Investment Experiment 1: Results in Restricted Sample

Here, we present the results for the restricted sample of the Investment Experiment 1, featuring 193 participants who solved both of the example screens.

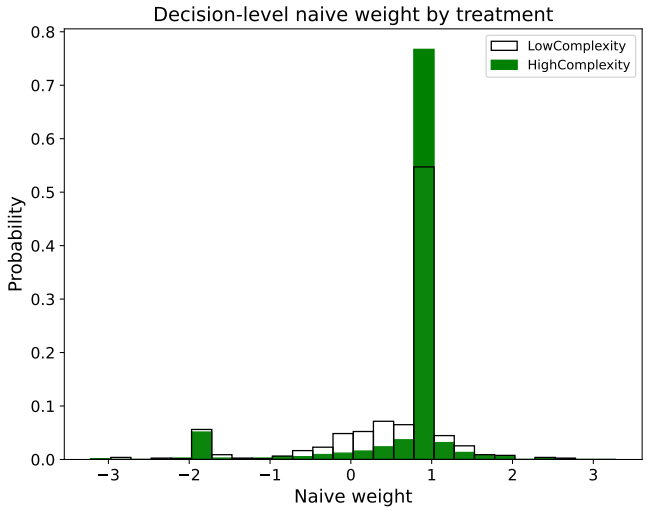

(a) Distribution of decision-level naive weights

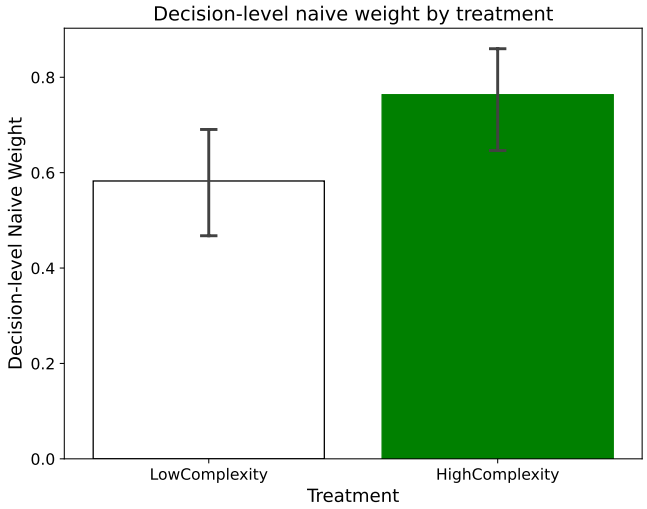

(b) Mean decision-level naive weights

Figure 2.A.5. Decision-level naive weights in the restricted sample of the Investment Experiment 1. *Panel (a) plots the distribution of naive weights λ calculated as specified in Equation 2.2, using the restricted sample of the Investment Experiment 1 with 193 participants. Panel (b) plots average naive weights.*

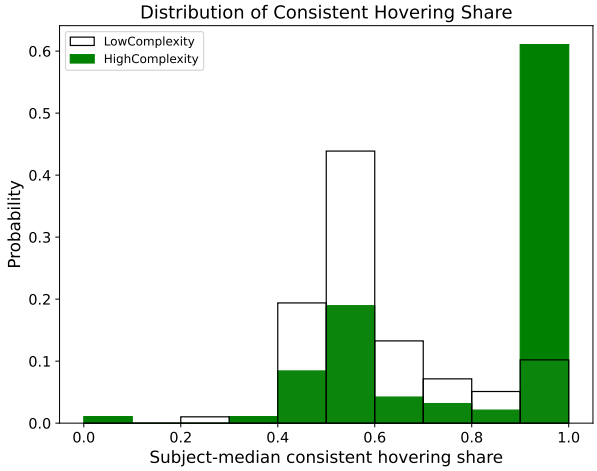

Figure 2.A.6. Distribution of subject-medians of the consistent hovering shares in the restricted sample of the Investment Experiment 1. *The figure plots the distribution of the share of time that respondents spent looking at the values of the signal-consistent model, using the restricted sample of the Investment Experiment 1 with 193 participants. Only the median consistent share for each participant is plotted.*

Table 2.A.2. Company Bids in Restricted Sample of Investment Experiment 1

Dependent variable:	Company Bids		
Sample:	LowComplexity (1)	HighComplexity (2)	Pooled (3)
Rational Benchmark	0.478***	0.311***	0.478***
	(0.065)	(0.065)	(0.064)
Naive Benchmark	0.558***	0.725***	0.558***
	(0.053)	(0.055)	(0.052)
Rational B. × HighComplexity			-0.167*
			(0.092)
Naive B. × HighComplexity			0.167**
			(0.076)
R^2	0.906	0.920	0.913
Observations	784	760	1544

The table presents OLS regressions of respondents' company bids on the rational and naive benchmarks as detailed in Section 2.3.1. The table is based on the restricted sample of the Investment Experiment 1. Column (1) uses observations from the LowComplexity treatment, column (2) from the HighComplexity treatment, and column (3) from both. Stars highlight significant differences from 0 with * for $p < 0.10$, ** for $p < 0.05$, *** for $p < 0.01$. Standard errors are clustered on the subject level.

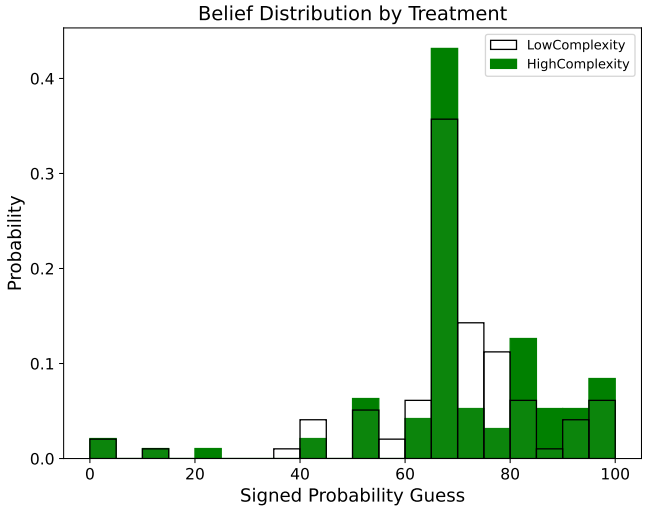

(a) Histogram of signed probability guesses.

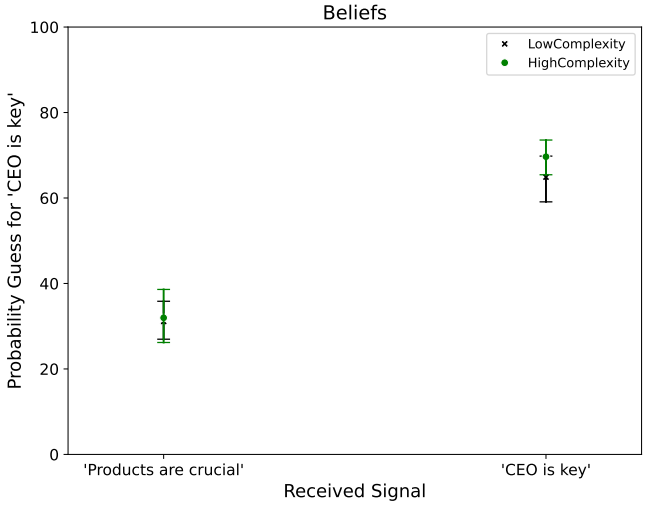

(b) Mean probability guesses for "The CEO is key".

Figure 2.A.7. Beliefs in the restricted sample of the Investment Experiment 1. *In Panel (a) beliefs are converted into the direction of the more likely model, so that 65 corresponds to the Bayesian probability. Panel (b) plots mean probability guesses for the model "The CEO is key". The figure is based on the restricted sample of the Investment Experiment 1 with 193 participants.*

Table 2.A.3. Beliefs and Recall in Restricted Sample of the Investment Experiment 1

Dependent variable:	Probability Guess	Correct Recall
Sample:	Pooled (1)	Pooled (2)
Constant	66.844***	0.959***
	(1.758)	(0.020)
HighComplexity	2.044	0.030
	(2.598)	(0.023)
R^2	0.003	0.009
Observations	193	193

The table presents OLS regressions using the restricted sample of the Investment Experiment 1. In column (1), the dependent variable is the probability guess for the likely state. Column (2) uses a dummy for whether respondents correctly recall the more likely model. All columns use observations from both the HighComplexity and LowComplexity conditions. Stars highlight significant differences from 0 with * for $p < 0.10$, ** for $p < 0.05$, *** for $p < 0.01$. Standard errors are clustered on the subject level.

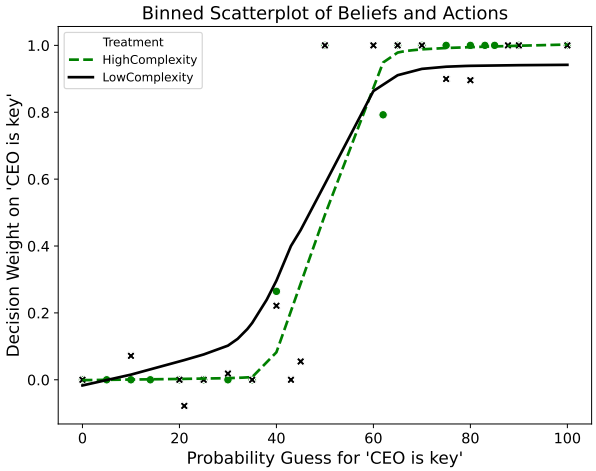

Figure 2.A.8. The relationship between decisions and beliefs in the restricted sample of the Investment Experiment 1. *The figure shows a binned scatterplot using the restricted sample of the Investment Experiment 1 with 193 participants. It has the probability guess that "The CEO is key" is the more likely model on the horizontal axis, and the decision weight γ as defined in Equation 2.3 on the vertical axis. The lines show LOWESS regressions based on all datapoints.*

2.A.1.5 Investment Experiment 2: Results in Restricted Sample

Here, we present the results for the restricted sample of the Investment Experiment 2, featuring 323 participants who solved both of the example screens.

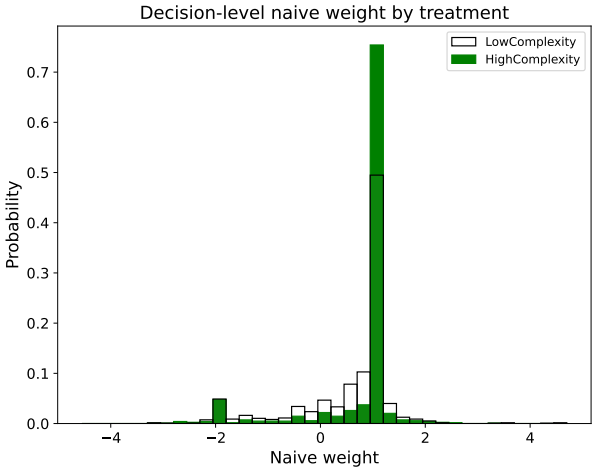

(a) Distribution of decision-level naive weights

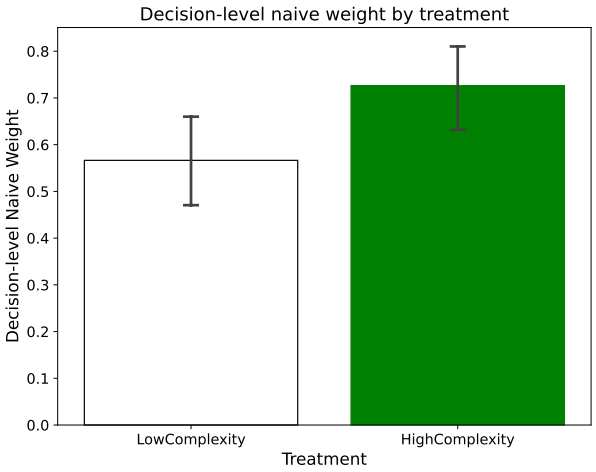

(b) Mean decision-level naive weights

Figure 2.A.9. Decision-level naive weights in the restricted sample of the Investment Experiment 2. *Panel (a) plots the distribution of naive weights λ calculated as specified in Equation 2.2, using the restricted sample of the Investment Experiment 2 with 323 participants. Panel (b) plots average naive weights.*

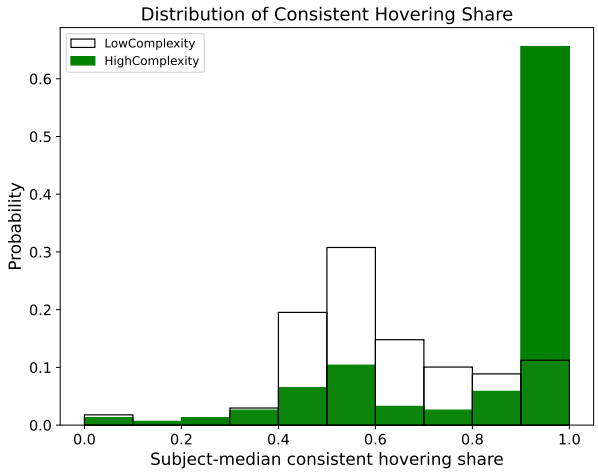

Figure 2.A.10. Distribution of subject-medians of the consistent hovering shares in the restricted sample of the Investment Experiment 2. *The figure plots the distribution of the share of time that respondents spent looking at the values of the signal-consistent model, using the restricted sample of the Investment Experiment 2 with 323 participants. Only the median consistent share for each participant is plotted.*

Table 2.A.4. Company Bids in Restricted Sample of Investment Experiment 2

Dependent variable:	Company Bids		
Sample:	LowComplexity (1)	HighComplexity (2)	Pooled (3)
Rational Benchmark	0.508***	0.281***	0.508***
	(0.052)	(0.050)	(0.052)
Naive Benchmark	0.534***	0.721***	0.534***
	(0.043)	(0.045)	(0.043)
Rational B. × HighComplexity			-0.227***
			(0.072)
Naive B. × HighComplexity			0.187***
			(0.062)
R^2	0.892	0.908	0.900
Observations	1352	1232	2584

The table presents OLS regressions of respondents' company bids on the rational and naive benchmarks as detailed in Section 2.3.1. The table is based on the restricted sample of the Investment Experiment 2. Column (1) uses observations from the LowComplexity treatment, column (2) from the HighComplexity treatment, and column (3) from both. Stars highlight significant differences from 0 with * for $p < 0.10$, ** for $p < 0.05$, *** for $p < 0.01$. Standard errors are clustered on the subject level.

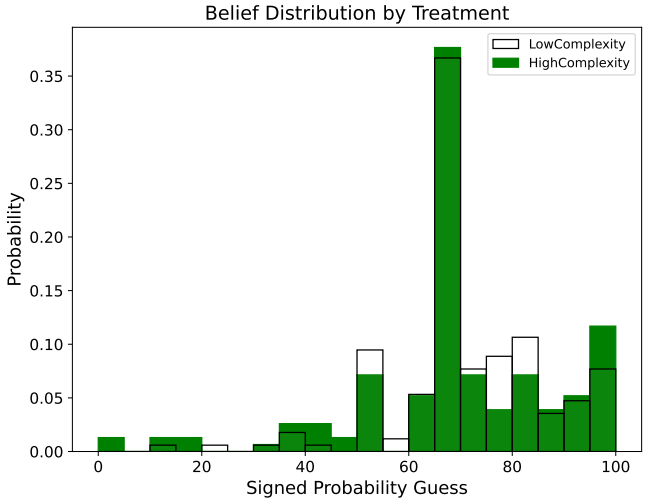

(a) Histogram of signed probability guesses.

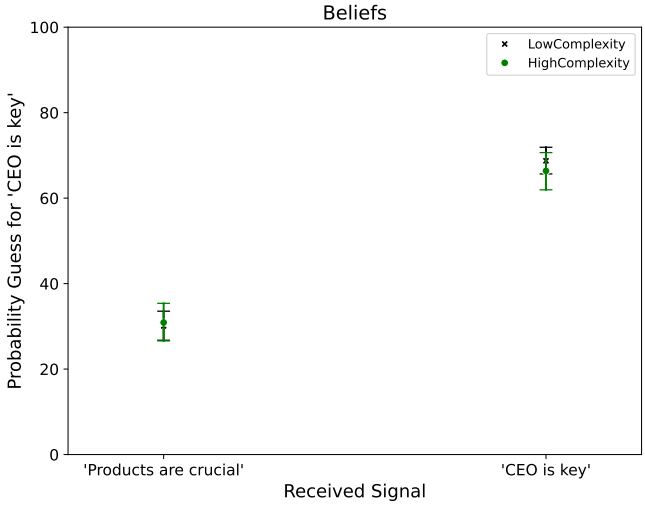

(b) Mean probability guesses for "The CEO is key".

Figure 2.A.11. Beliefs in the restricted sample of the Investment Experiment 2. *In Panel (a) beliefs are converted into the direction of the more likely model, so that 65 corresponds to the Bayesian probability. Panel (b) plots mean probability guesses for the model "The CEO is key". The figure is based on the restricted sample of the Investment Experiment 2 with 323 participants.*

Table 2.A.5. Beliefs and Recall in Restricted Sample of the Investment Experiment 2

Dependent variable:	Probability Guess	Correct Recall
Sample:	Pooled (1)	Pooled (2)
Constant	69.302***	0.982***
	(1.176)	(0.010)
HighComplexity	-1.563	-0.021
	(1.999)	(0.019)
R^2	0.002	0.004
Observations	323	323

The table presents OLS regressions using the restricted sample of the Investment Experiment 2. In column (1), the dependent variable is the probability guess for the likely state. Column (2) uses a dummy for whether respondents correctly recall the more likely model. All columns use observations from both the HighComplexity and LowComplexity conditions. Stars highlight significant differences from 0 with * for $p < 0.10$, ** for $p < 0.05$, *** for $p < 0.01$. Standard errors are clustered on the subject level.

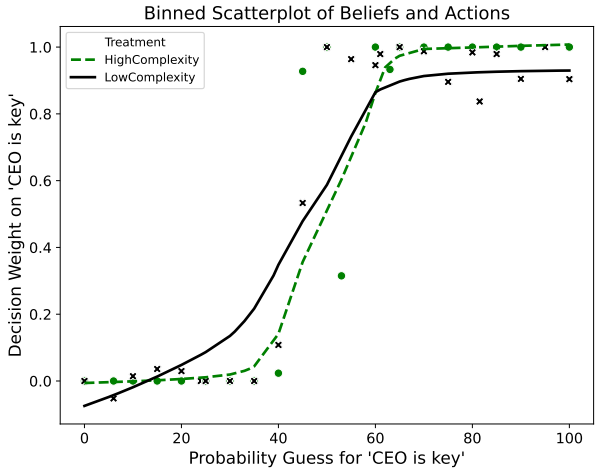

Figure 2.A.12. The relationship between decisions and beliefs in the restricted sample of the Investment Experiment 2. *The figure shows a binned scatterplot using the restricted sample of the Investment Experiment 2 with 323 participants. It has the probability guess that "The CEO is key" is the more likely model on the horizontal axis, and the decision weight γ as defined in Equation 2.3 on the vertical axis. The lines show LOWESS regressions based on all datapoints.*

2.A.1.6 Delayed Belief Elicitation: Results in Restricted Sample

Here, we present the results for the restricted sample of the Delayed Belief Elicitation Experiment, featuring 342 participants who solved both of the example screens.

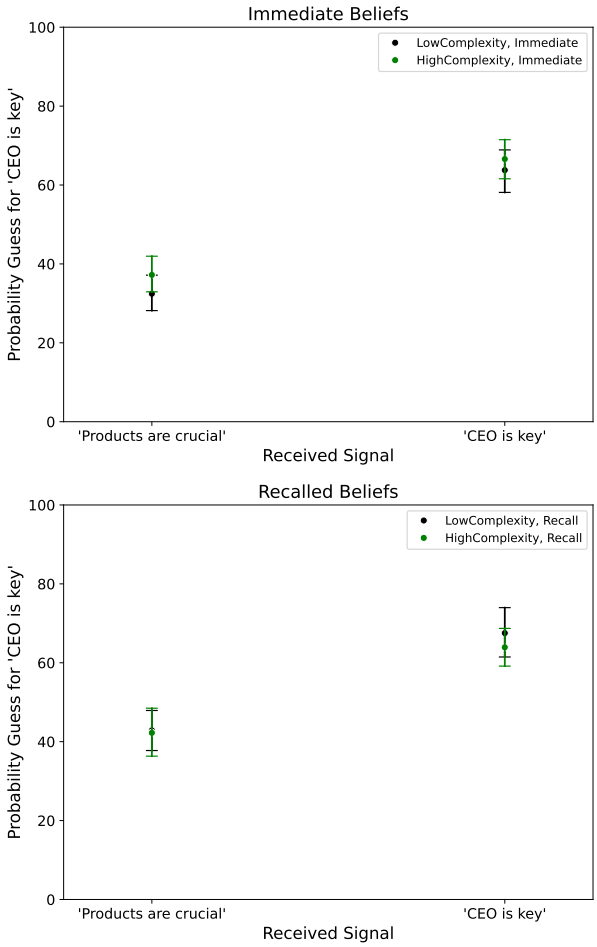

Figure 2.A.13. Immediate and recalled beliefs in the Delayed Belief Elicitation Experiment. *The top panel plots the mean probability guess for the "The CEO is key" model in the Immediate condition by received signal separately for the LowComplexity and HighComplexity condition. The bottom panel does the same for the Recall condition, where beliefs were elicited with a one day delay. This figure is based on the restricted sample of the Delayed Belief Elicitation Experiment with 342 participants.*

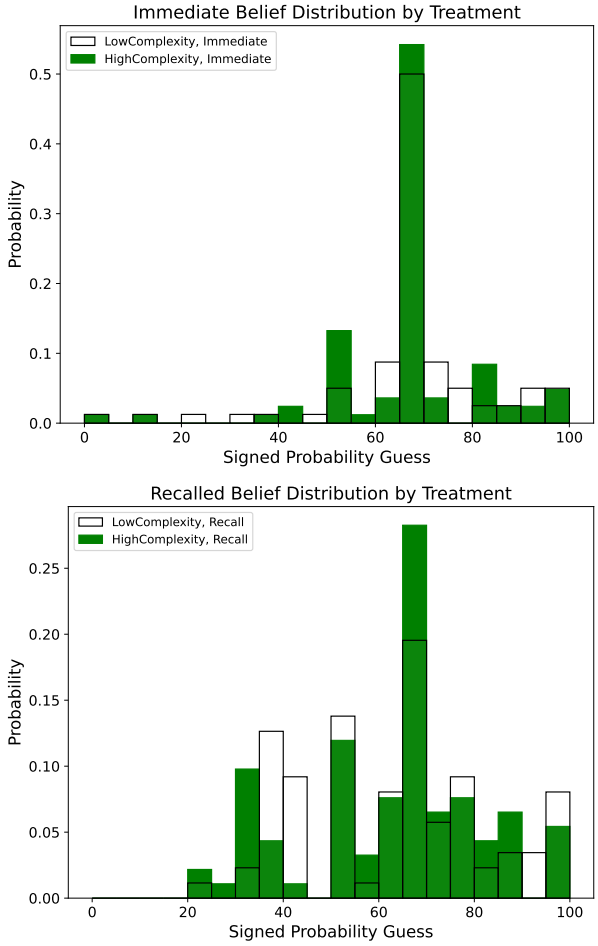

Figure 2.A.14. Distribution of immediate and recalled beliefs in the restricted sample of the Delayed Belief Elicitation Experiment. *The top panel plots the distribution of immediate beliefs and the bottom panel of recalled beliefs. Beliefs are converted into the direction of the more likely model, so that 65 corresponds to the Bayesian probability. This figure is based on the restricted sample of the Delayed Belief Elicitation Experiment with 342 participants.*

Table 2.A.6. Beliefs and Recall in the Restricted Sample of the Delayed Belief Eliciation Experiment

Dependent variable:	Probability Guess		Correct Recall	
Sample:	Immediate (1)	Recall (2)	Immediate (3)	Recall (4)
Constant	65.713***	61.195***	0.950***	0.736***
	(1.865)	(2.106)	(0.025)	(0.048)
HighComplexity	-1.008	0.044	0.026	0.014
	(2.547)	(2.843)	(0.030)	(0.066)
R^2	0.001	0.000	0.005	0.000
Observations	163	179	163	179

The table presents OLS regressions using the restricted sample of the Delayed Belief Elicitation Experiment. In columns (1) and (2), we regress respondents' probability guesses on a constant and a treatment dummy for the HighComplexity condition. Probability guesses are converted in the direction of the more likely model, so that a guess of 65 corresponds to the Bayesian probability. In columns (3) and (4), the dependent variable is a dummy for whether respondents correctly recall the more likely model. Columns (1) and (3) use observations from the Immediate condition, while columns (2) and (4) use observations from the Recall condition. Stars highlight significant differences from 0 with * for $p < 0.10$, ** for $p < 0.05$, *** for $p < 0.01$. Standard errors are clustered on the subject level.

2.A.1.7 Baseline Confidence: Additional Results in Restricted Sample

Here, we present the results for the restricted sample of the Baseline Confidence Experiment, featuring 336 participants who solved both of the example screens.

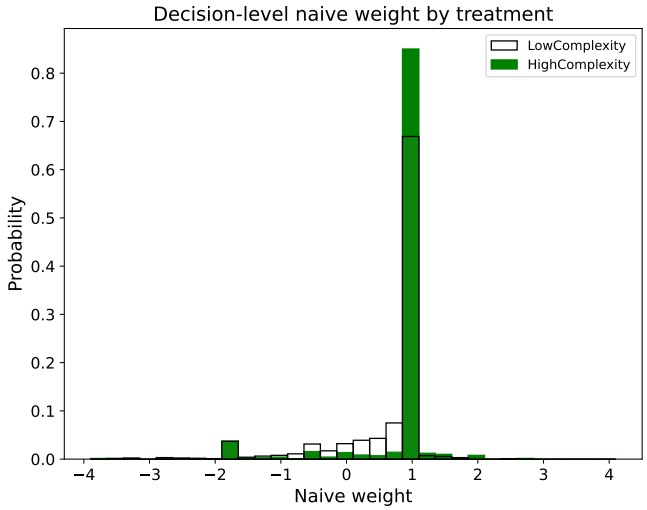

(a) Distribution of decision-level naive weights

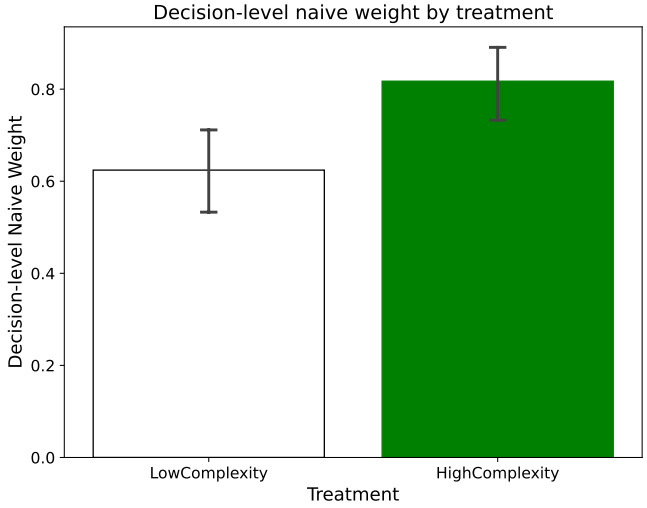

(b) Mean decision-level naive weights

Figure 2.A.15. Decision-level naive weights in the restricted sample of the Baseline Confidence Experiment. *Panel (a) plots the distribution of naive weights λ calculated as specified in Equation 2.2, using the restricted sample of the Baseline Confidence Experiment with 336 participants. Panel (b) plots average naive weights.*

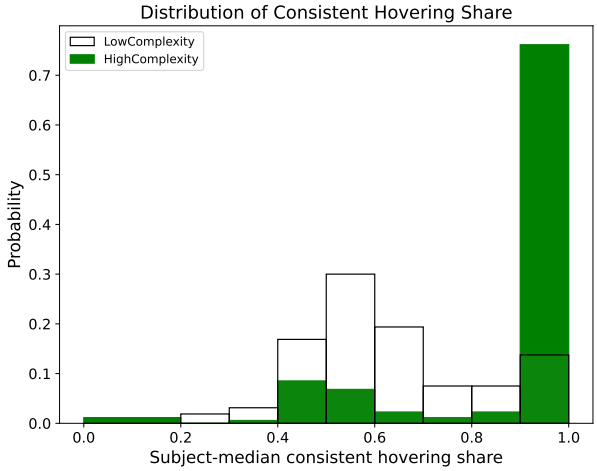

Figure 2.A.16. Distribution of subject-medians of the consistent hovering shares in the restricted sample of the Baseline Confidence Experiment. *The figure plots the distribution of the share of time that respondents spent looking at the values of the signal-consistent model, using the restricted sample of the Baseline Confidence Experiment with 336 participants. Only the median consistent share for each participant is plotted.*

Table 2.A.7. Company Value Guesses in the Restricted Sample of the Baseline Confidence Experiment

Dependent variable:	Company Value Guess		
Sample:	LowComplexity (1)	HighComplexity (2)	Pooled (3)
Rational Benchmark	0.428***	0.226***	0.428***
	(0.052)	(0.042)	(0.052)
Naive Benchmark	0.613***	0.797***	0.613***
	(0.044)	(0.039)	(0.044)
Rational B. × HighComplexity			-0.202***
			(0.067)
Naive B. × HighComplexity			0.184***
			(0.059)
R^2	0.920	0.936	0.928
Observations	1280	1408	2688

The table presents OLS regressions of respondents' company value guesses on the rational and naive benchmarks as detailed in Section 2.3.1. The table is based on the restricted sample of the Baseline Confidence Experiment. Column (1) uses observations from the LowComplexity treatment, column (2) from the HighComplexity treatment, and column (3) from both. Stars highlight significant differences from 0 with * for $p < 0.10$, ** for $p < 0.05$, *** for $p < 0.01$. Standard errors are clustered on the subject level.

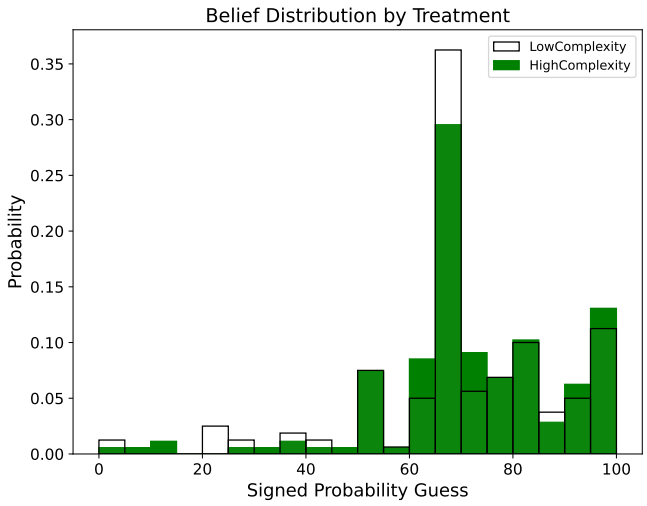

(a) Histogram of signed probability guesses.

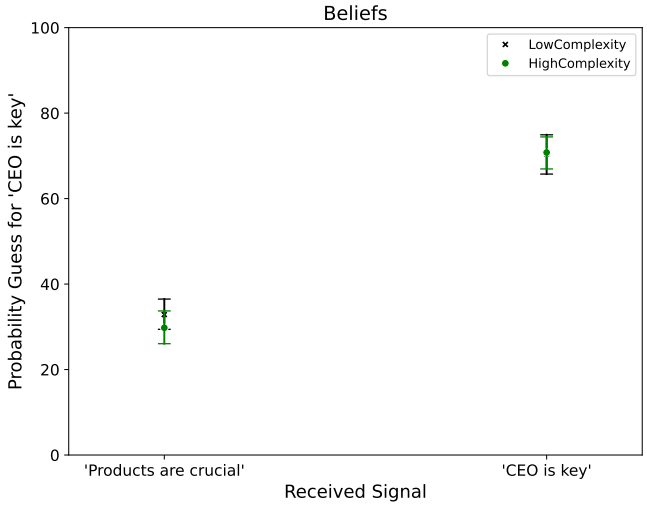

(b) Mean probability guesses for "The CEO is key".

Figure 2.A.17. Beliefs in the restricted sample of the Baseline Confidence Experiment. *In Panel (a) beliefs are converted into the direction of the more likely model, so that 65 corresponds to the Bayesian probability. Panel (b) plots mean probability guesses for the model "The CEO is key". The figure is based on the restricted sample of the Baseline Confidence Experiment with 336 participants.*

Table 2.A.8. Beliefs and Recall in the Restricted Sample of the Baseline Confidence Experiment

Dependent variable:	Probability Guess	Correct Recall
Sample:	Pooled (1)	Pooled (2)
Constant	68.806*** (1.518)	0.944*** (0.018)
HighComplexity	1.716 (2.056)	0.022 (0.023)
R^2	0.002	0.003
Observations	336	336

The table presents OLS regressions using the restricted sample of the Baseline Confidence Experiment. In column (1), the dependent variable is the probability guess for the likely state. Column (2) uses a dummy for whether respondents correctly recall the more likely model. All columns use observations from both the HighComplexity and LowComplexity conditions. Stars highlight significant differences from 0 with * for $p < 0.10$, ** for $p < 0.05$, *** for $p < 0.01$. Standard errors are clustered on the subject level.

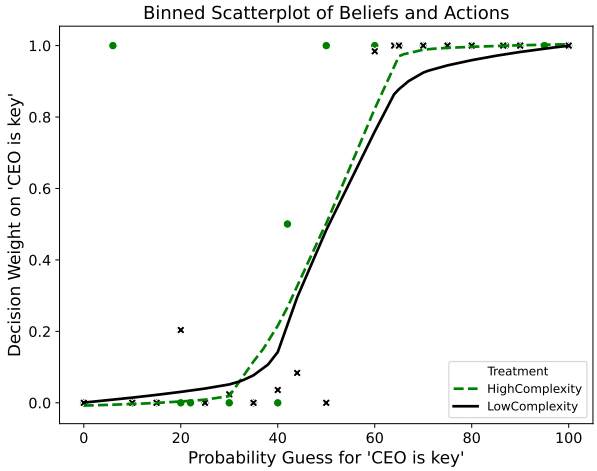

Figure 2.A.18. The relationship between decisions and beliefs in the restricted sample of the Baseline Confidence Experiment. *The figure shows a binned scatterplot using the restricted sample of the Baseline Confidence Experiment with 336 participants. It has the probability guess that "The CEO is key" is the more likely model on the horizontal axis, and the decision weight γ as defined in Equation 2.3 on the vertical axis. The lines show LOWESS regressions based on all datapoints.*

2.A.1.8 Incentivized Confidence: Additional Results in Restricted Sample

Here, we present the results for the restricted sample of the Incentivized Confidence Experiment, featuring 203 participants who solved both of the example screens.

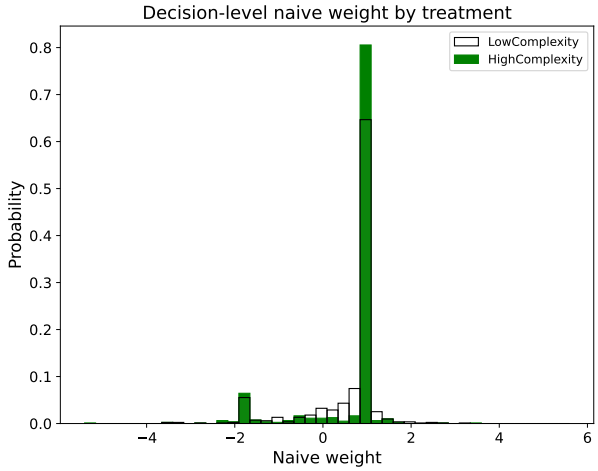

(a) Distribution of decision-level naive weights

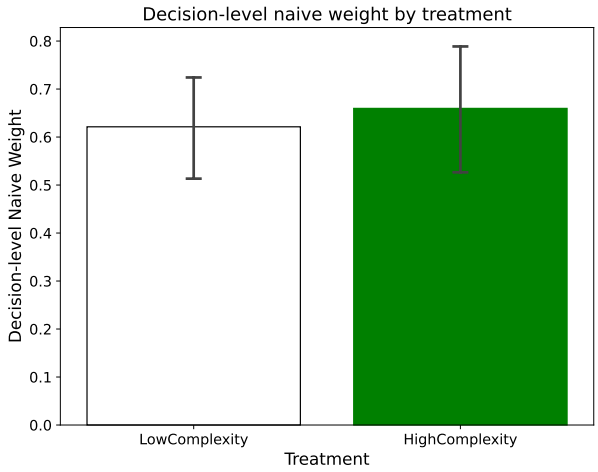

(b) Mean decision-level naive weights

Figure 2.A.19. Decision-level naive weights in the restricted sample of the Incentivized Confidence Experiment. *Panel (a) plots the distribution of naive weights λ calculated as specified in Equation 2.2, using the restricted sample of the Incentivized Confidence Experiment with 203 participants. Panel (b) plots average naive weights.*

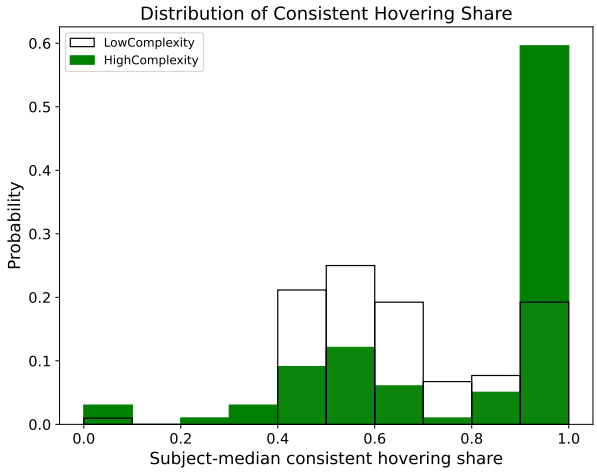

Figure 2.A.20. Distribution of subject-medians of the consistent hovering shares in the restricted sample of the Incentivized Confidence Experiment. *The figure plots the distribution of the share of time that respondents spent looking at the values of the signal-consistent model, using the restricted sample of the Incentivized Confidence Experiment with 203 participants. Only the median consistent share for each participant is plotted.*

Table 2.A.9. Company Value Guesses in the Restricted Sample of the Incentivized Confidence Experiment

Dependent variable:	Company Value Guess		
Sample:	LowComplexity (1)	HighComplexity (2)	Pooled (3)
Rational Benchmark	0.488***	0.404***	0.488***
	(0.069)	(0.080)	(0.069)
Naive Benchmark	0.593***	0.638***	0.593***
	(0.054)	(0.069)	(0.054)
Rational B. × HighComplexity			-0.084
			(0.106)
Naive B. × HighComplexity			0.045
			(0.087)
R^2	0.912	0.898	0.905
Observations	832	792	1624

The table presents OLS regressions of respondents' company value guesses on the rational and naive benchmarks as detailed in Section 2.3.1. The table is based on the restricted sample of the Incentivized Confidence Experiment. Column (1) uses observations from the LowComplexity treatment, column (2) from the HighComplexity treatment, and column (3) from both. Stars highlight significant differences from 0 with * for $p < 0.10$, ** for $p < 0.05$, *** for $p < 0.01$. Standard errors are clustered on the subject level.

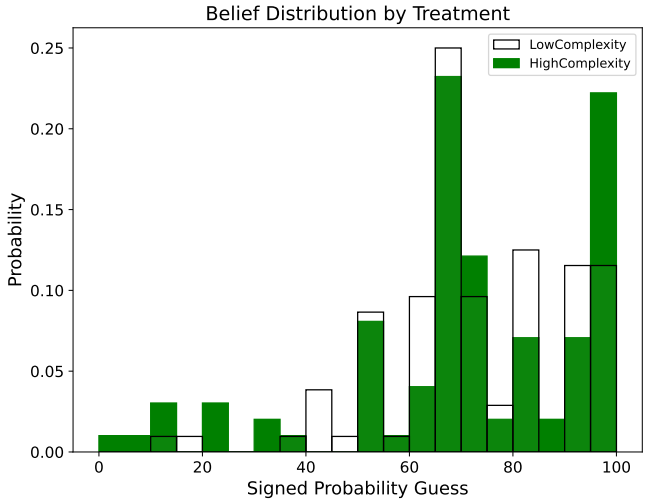

(a) Histogram of signed probability guesses.

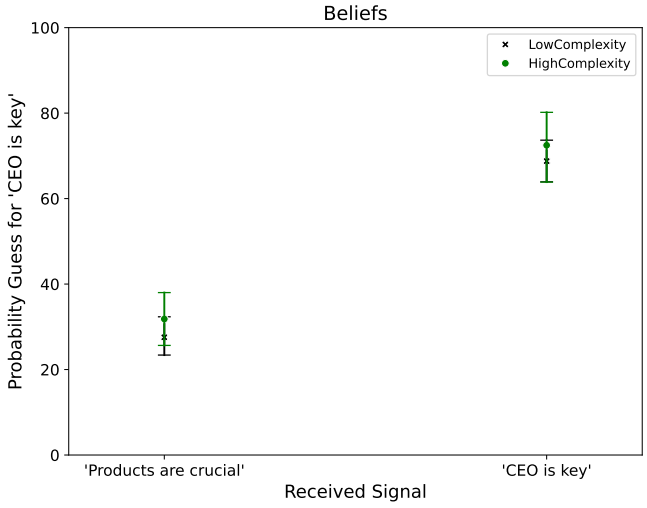

(b) Mean probability guesses for "The CEO is key".

Figure 2.A.21. Beliefs in the restricted sample of the Incentivized Confidence Experiment. *In Panel (a) beliefs are converted into the direction of the more likely model, so that 65 corresponds to the Bayesian probability. Panel (b) plots mean probability guesses for the model "The CEO is key". The figure is based on the restricted sample of the Incentivized Confidence Experiment with 203 participants.*

Table 2.A.10. Beliefs and Recall in the Restricted Sample of the Incentivized Confidence Experiment

Dependent variable:	Probability Guess	Correct Recall
Sample:	Pooled (1)	Pooled (2)
Constant	70.327***	0.962***
	(1.766)	(0.019)
HighComplexity	-0.238	0.018
	(3.036)	(0.024)
R^2	0.000	0.003
Observations	203	203

The table presents OLS regressions using the restricted sample of the Incentivized Confidence Experiment. In column (1), the dependent variable is the probability guess for the likely state. Column (2) uses a dummy for whether respondents correctly recall the more likely model. All columns use observations from both the HighComplexity and LowComplexity conditions. Stars highlight significant differences from 0 with * for $p < 0.10$, ** for $p < 0.05$, *** for $p < 0.01$. Standard errors are clustered on the subject level.

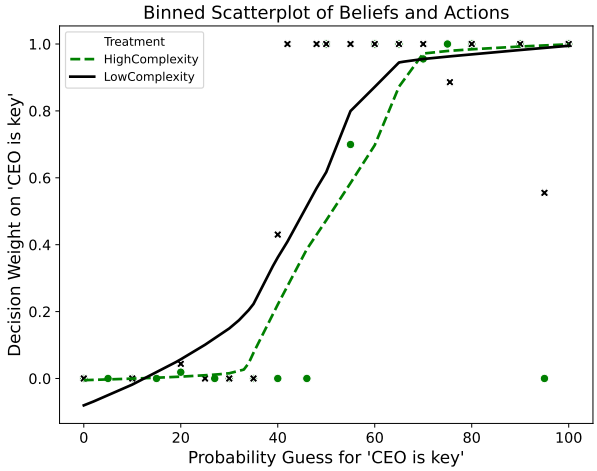

Figure 2.A.22. The relationship between decisions and beliefs in the restricted sample of the Incentivized Confidence Experiment. *The figure shows a binned scatterplot using the restricted sample of the Incentivized Confidence Experiment with 203 participants. It has the probability guess that "The CEO is key" is the more likely model on the horizontal axis, and the decision weight γ as defined in Equation 2.3 on the vertical axis. The lines show LOWESS regressions based on all datapoints.*

2.A.2 Additional Results for the Robustness Samples

2.A.2.1 Baseline Experiment: Results in Lenient Sample

Here, we present the results for the lenient sample of the Baseline Experiment, featuring 319 participants who solved at least one of the example screens.

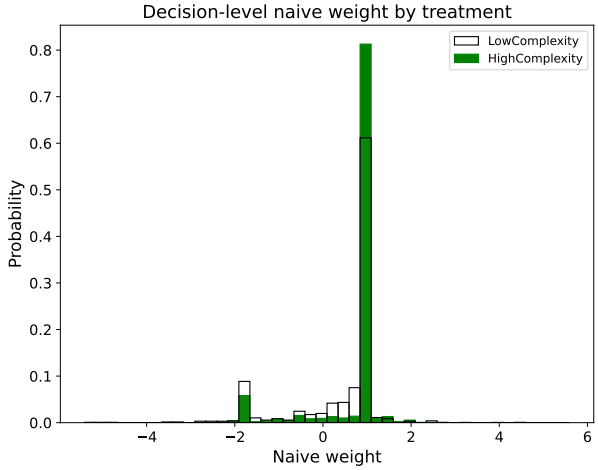

(a) Distribution of decision-level naive weights

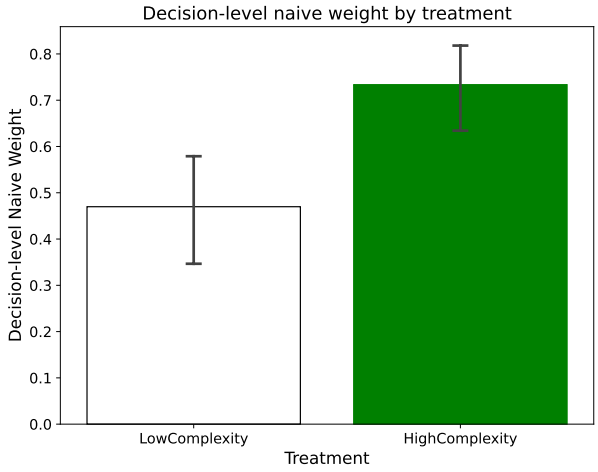

(b) Mean decision-level naive weights

Figure 2.A.23. Decision-level naive weights in the lenient sample of the Baseline Experiment. *Panel (a) plots the distribution of naive weights λ calculated as specified in Equation 2.2, using the lenient sample of the Baseline Experiment with 319 participants. Panel (b) plots average naive weights.*

Table 2.A.11. Company Value Guesses in the Lenient Sample of the Baseline Experiment

Dependent variable:	Company Value Guess		
Sample:	LowComplexity (1)	HighComplexity (2)	Pooled (3)
Rational Benchmark	0.658***	0.334***	0.658***
	(0.070)	(0.053)	(0.070)
Naive Benchmark	0.432***	0.702***	0.432***
	(0.060)	(0.047)	(0.060)
Rational B. × HighComplexity			-0.323***
			(0.088)
Naive B. × HighComplexity			0.270***
			(0.076)
R^2	0.875	0.915	0.894
Observations	1264	1288	2552

The table presents OLS regressions of respondents' company value guesses on the rational and naive benchmarks as detailed in Section 2.3.1. The table is based on the lenient sample of the Baseline Experiment. Column (1) uses observations from the LowComplexity treatment, column (2) from the HighComplexity treatment, and column (3) from both. Stars highlight significant differences from 0 with * for $p < 0.10$, ** for $p < 0.05$, *** for $p < 0.01$. Standard errors are clustered on the subject level.

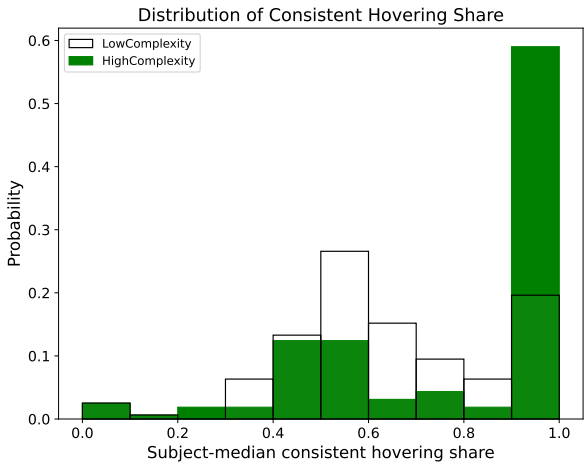

Figure 2.A.24. Distribution of subject-medians of the consistent hovering shares in the lenient sample of the Baseline Experiment. *The figure plots the distribution of the share of time that respondents spent looking at the values of the signal-consistent model, using the lenient sample of the Baseline Experiment with 319 participants. Only the median consistent share for each participant is plotted.*

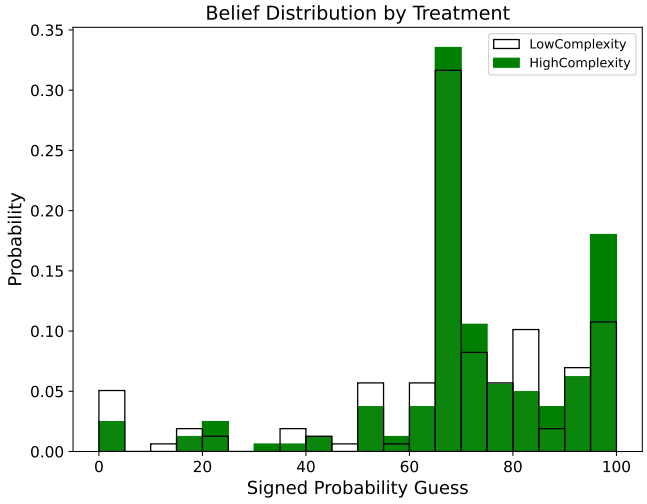

(a) Histogram of signed probability guesses.

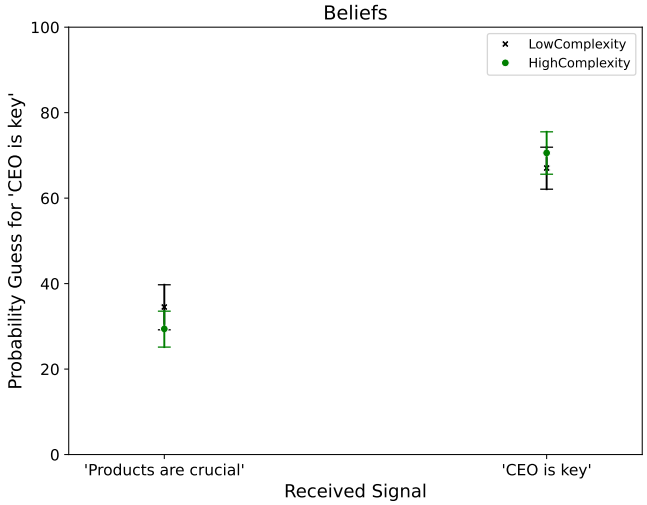

(b) Mean probability guesses for "The CEO is key".

Figure 2.A.25. Beliefs in the Lenient Sample of the Baseline Experiment. *In Panel (a) beliefs are converted into the direction of the more likely model, so that 65 corresponds to the Bayesian probability. Panel (b) plots mean probability guesses for the model "The CEO is key". The figure is based on the lenient sample of the Baseline Experiment with 319 participants.*

Table 2.A.12. Beliefs and Recall in the Lenient Sample of the Baseline Experiment

Dependent variable:	Probability Guess	Correct Recall
Sample:	Pooled (1)	Pooled (2)
Constant	66.321***	0.943***
	(1.871)	(0.019)
HighComplexity	4.266*	0.020
	(2.536)	(0.024)
R^2	0.009	0.002
Observations	319	319

The table presents OLS regressions using the lenient sample of the Baseline Experiment. In column (1), the dependent variable is the probability guess for the likely state. Column (2) uses a dummy for whether respondents correctly recall the more likely model. All columns use observations from both the HighComplexity and LowComplexity conditions. Stars highlight significant differences from 0 with * for $p < 0.10$, ** for $p < 0.05$, *** for $p < 0.01$. Standard errors are clustered on the subject level.

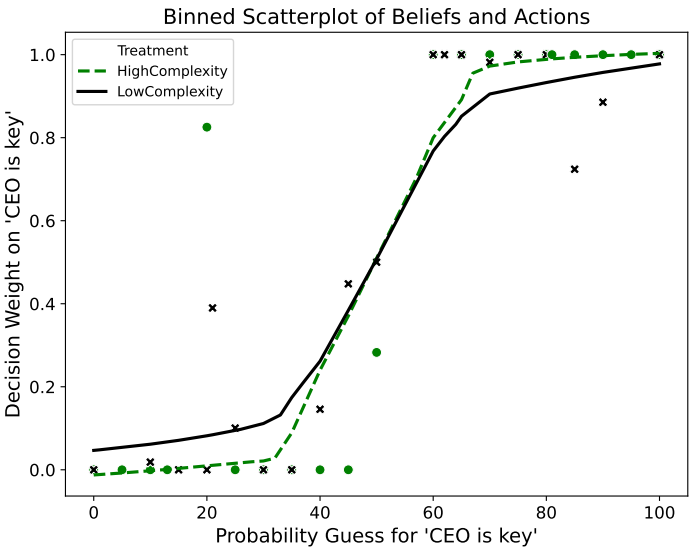

Figure 2.A.26. The relationship between decisions and immediate beliefs in the lenient sample of the Baseline Experiment. *The figure shows a binned scatterplot using the lenient sample of the Baseline Experiment with 319 participants. It has the probability guess that "The CEO is key" is the more likely model on the horizontal axis, and the decision weight γ as defined in Equation 2.3 on the vertical axis. The lines show LOWESS regressions based on all datapoints.*

2.A.2.2 Equally Likely Models: Results in Lenient Sample

Here, we present the results for the lenient sample of the Equally Likely Models Experiment, featuring 452 participants who solved at least one of the example screens.

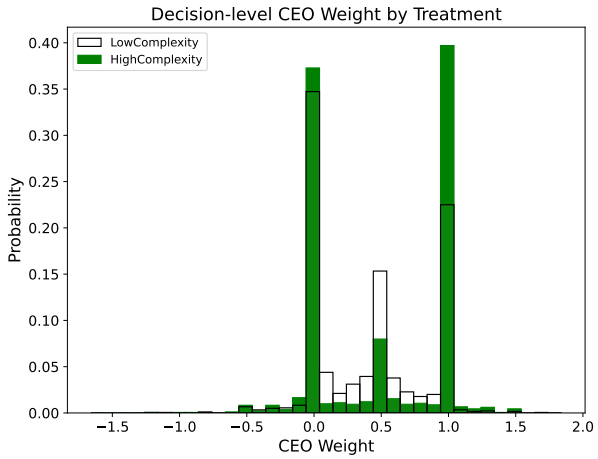

(a) Distribution of decision-level CEO weights

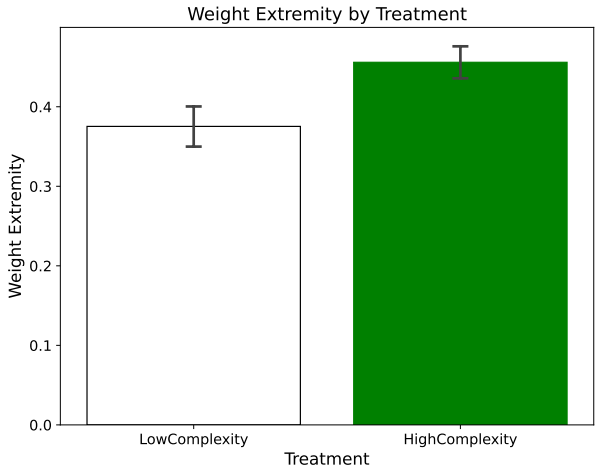

(b) Mean decision-level CEO weight extremity

Figure 2.A.27. Decision-level CEO weights and CEO weight extremity in the lenient sample of the Equally Likely Models Experiment. *Panel (a) plots the distribution of CEO weights γ calculated as specified in Equation 2.3, using the lenient sample of the Equally Likely Models study with 452 participants. Panel (b) plots the average CEO weight extremity* $|\gamma - \frac{1}{2}|$ *calculated as specified in Equation 2.4.*

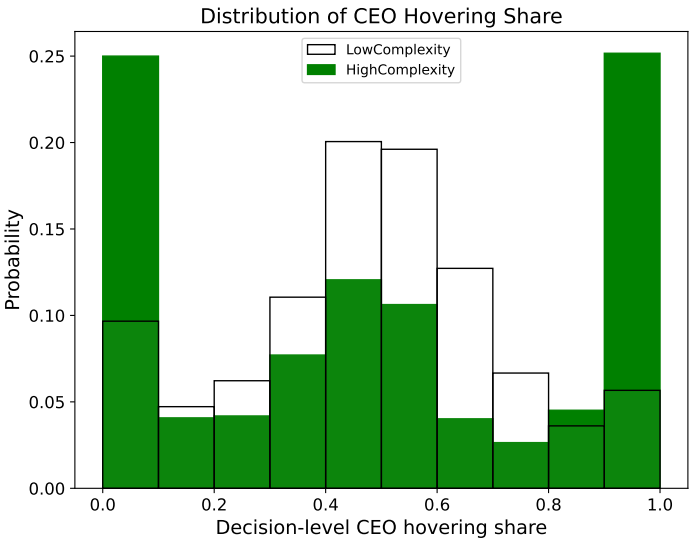

Figure 2.A.28. Distribution of decision-level CEO hovering share in the Equally Likely Models Experiment. *The figure plots the distribution of the share of time that respondents spent looking at the values of the "The CEO is key" model, using the lenient sample of the Equally Likely Models study with 452 participants.*

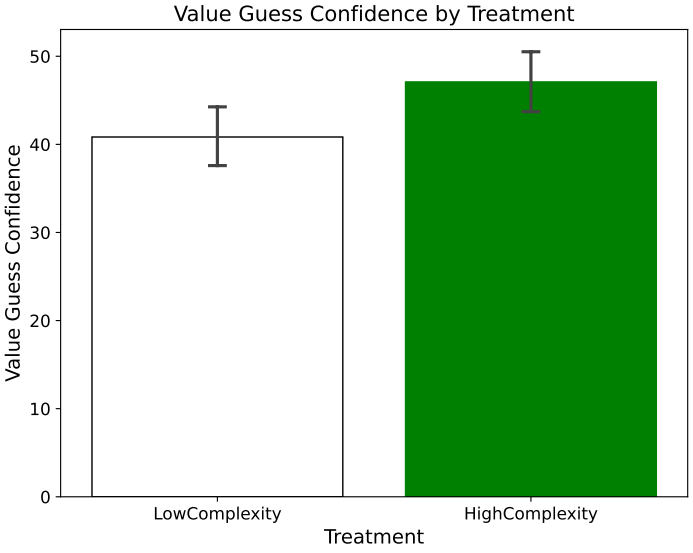

Figure 2.A.29. Average value guess confidence in the lenient sample of the Equally Likely Models Experiment. *The figure plots the average confidence that respondents had in their company value guesses, using the lenient sample of the Equally Likely Models study with 452 participants.*

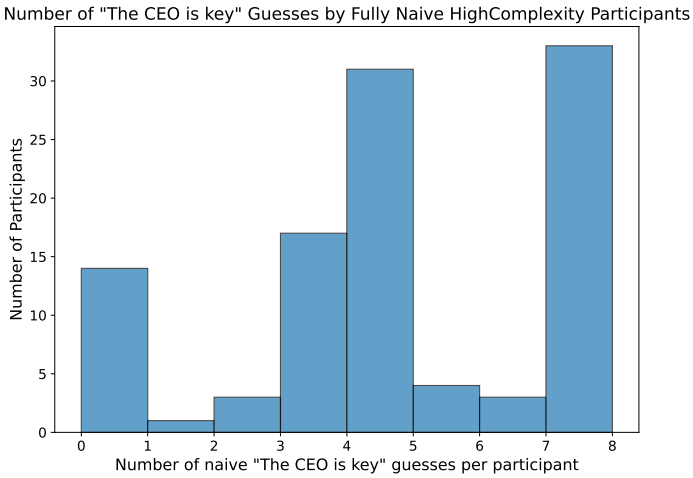

Figure 2.A.30. Distribution of naive decision-making in the lenient sample of the Equally Likely Models Experiment. *The figure plots the distribution of the number of times participants selected the value corresponding to the "The CEO is key" model, using the lenient sample of the HighComplexity treatment of the Equally Likely Models study. The sample is limited to the 106 participants in the HighComplexity condition who made only fully naive guesses, meaning they always selected a value corresponding to either the "The CEO is key" or "Products are crucial" model.*

2.A.2.3 Investment Experiment 1: Results in Lenient Sample

Here, we present the results for the lenient sample of the Investment Experiment 1, featuring 265 participants who solved at least one of the example screens.

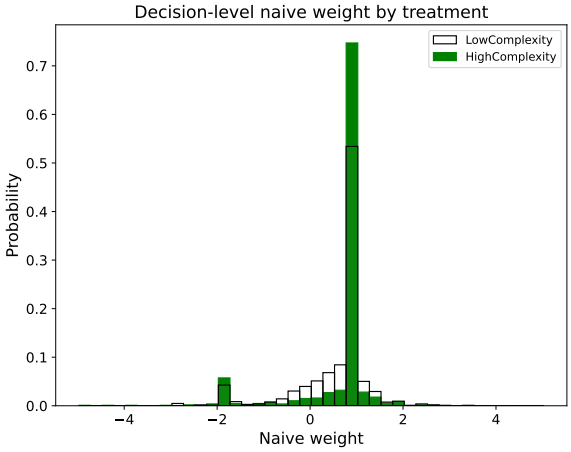

(a) Distribution of decision-level naive weights

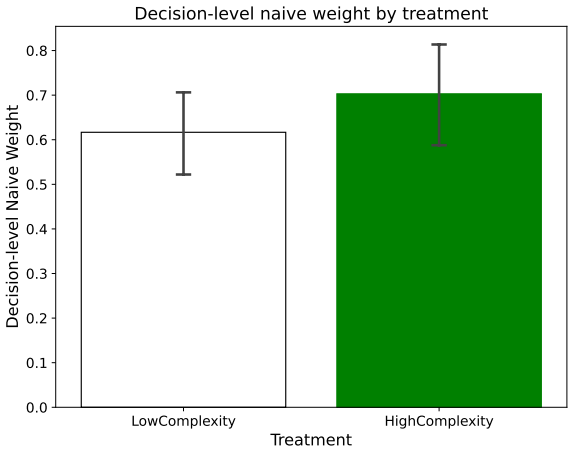

(b) Mean decision-level naive weights

Figure 2.A.31. Decision-level naive weights in the lenient sample of the Investment Experiment 1. *Panel (a) plots the distribution of naive weights λ calculated as specified in Equation 2.2, using the lenient sample of the Investment Experiment 1 with 265 participants. Panel (b) plots average naive weights.*

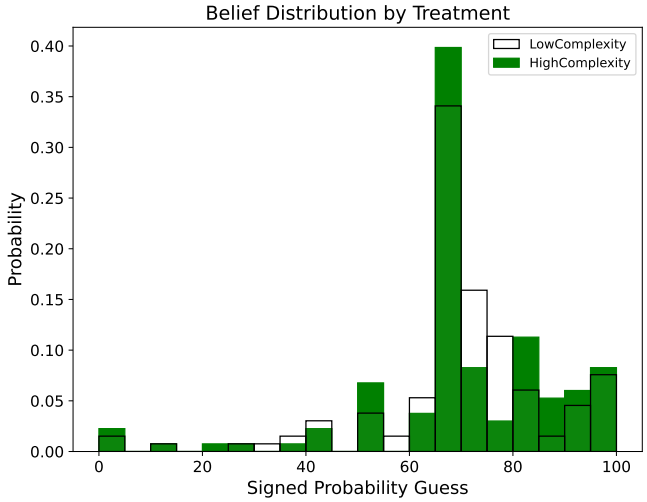

(a) Histogram of signed probability guesses.

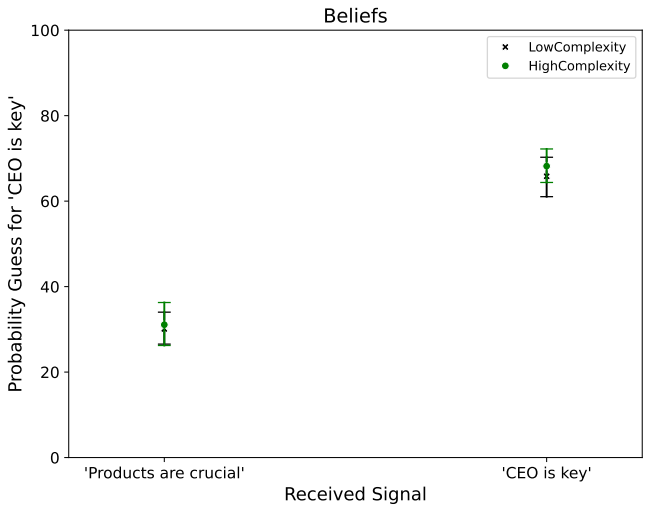

(b) Mean probability guesses for "The CEO is key".

Figure 2.A.32. Beliefs in the Lenient Sample of the Investment Experiment 1. *In Panel (a) beliefs are converted into the direction of the more likely model, so that 65 corresponds to the Bayesian probability. Panel (b) plots mean probability guesses for the model "The CEO is key". The figure is based on the lenient sample of the Investment Experiment 1 with 265 participants.*

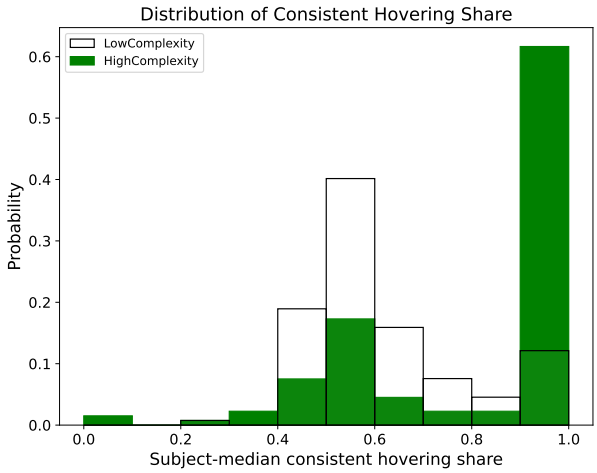

Figure 2.A.33. Distribution of subject-medians of the consistent hovering shares in the lenient sample of the Investment Experiment 1. *The figure plots the distribution of the share of time that respondents spent looking at the values of the signal-consistent model, using the lenient sample of the Investment Experiment 1 with 265 participants. Only the median consistent share for each participant is plotted.*

Table 2.A.13. Company Bids in the Lenient Sample of the Investment Experiment 1

Dependent variable:	Company Bids		
Sample:	LowComplexity (1)	HighComplexity (2)	Pooled (3)
Rational Benchmark	0.432***	0.363***	0.432***
	(0.052)	(0.065)	(0.052)
Naive Benchmark	0.596***	0.674***	0.596***
	(0.043)	(0.056)	(0.043)
Rational B. × HighComplexity			-0.070
			(0.083)
Naive B. × HighComplexity			0.078
			(0.071)
R^2	0.915	0.905	0.910
Observations	1056	1064	2120

The table presents OLS regressions of respondents' company bids on the rational and naive benchmarks as detailed in Section 2.3.1. The table is based on the lenient sample of the Investment Experiment 1. Column (1) uses observations from the LowComplexity treatment, column (2) from the HighComplexity treatment, and column (3) from both. Stars highlight significant differences from 0 with * for $p < 0.10$, ** for $p < 0.05$, *** for $p < 0.01$. Standard errors are clustered on the subject level.

Table 2.A.14. Beliefs and Recall in the Lenient Sample of the Investment Experiment 1

Dependent variable:	Probability Guess	Correct Recall
Sample:	Pooled	Pooled
	(1)	(2)
Constant	67.915***	0.955***
	(1.515)	(0.018)
HighComplexity	0.629	0.023
	(2.237)	(0.022)
R^2	0.000	0.004
Observations	265	265

The table presents OLS regressions using the lenient sample of the Investment Experiment 1. In column (1), the dependent variable is the probability guess for the likely state. Column (2) uses a dummy for whether respondents correctly recall the more likely model. All columns use observations from both the HighComplexity and LowComplexity conditions. Stars highlight significant differences from 0 with * for $p < 0.10$, ** for $p < 0.05$, *** for $p < 0.01$. Standard errors are clustered on the subject level.

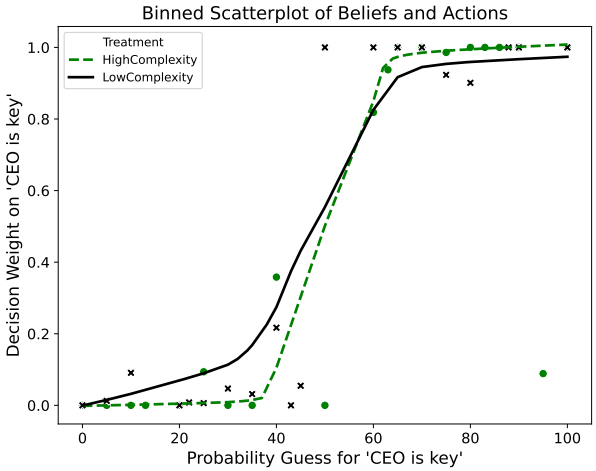

Figure 2.A.34. The relationship between decisions and immediate beliefs in the lenient sample of the Investment Experiment 1. *The figure shows a binned scatterplot using the lenient sample of the Investment Experiment 1 with 265 participants. It has the probability guess that "The CEO is key" is the more likely model on the horizontal axis, and the decision weight γ as defined in Equation 2.3 on the vertical axis. The lines show LOWESS regressions based on all datapoints.*

2.A.2.4 Investment Experiment 1: Results in Full Sample

Here, we present the results for the full sample of the Investment Experiment 1, featuring all 400 participants.

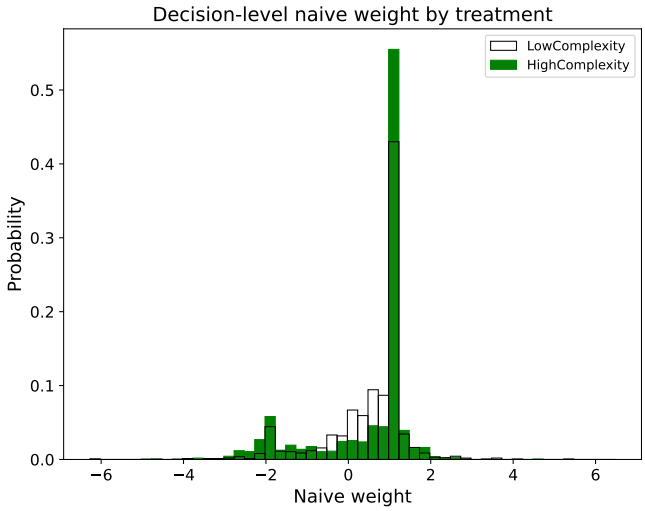

(a) Distribution of decision-level naive weights

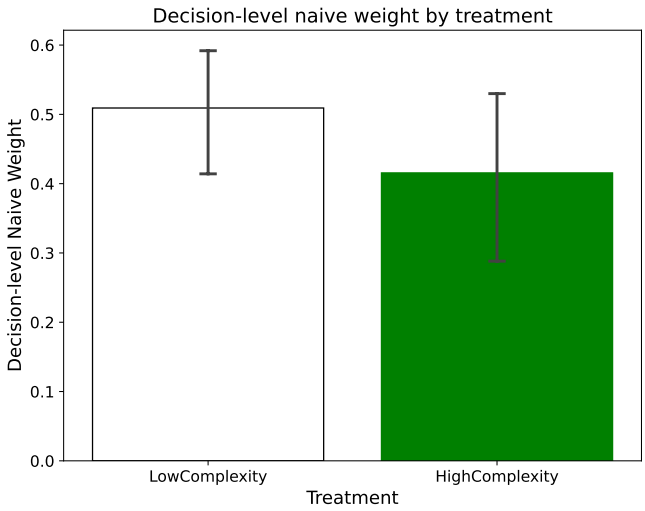

(b) Mean decision-level naive weights

Figure 2.A.35. Decision-level naive weights in the full sample of the Investment Experiment 1. *Panel (a) plots the distribution of naive weights λ calculated as specified in Equation 2.2, using the full sample of the Investment Experiment 1 with 400 participants. Panel (b) plots average naive weights.*

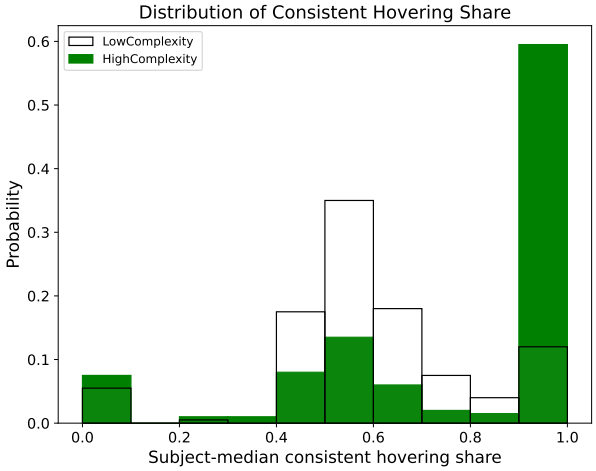

Figure 2.A.36. Distribution of subject-medians of the consistent hovering shares in the full sample of the Investment Experiment 1. *The figure plots the distribution of the share of time that respondents spent looking at the values of the signal-consistent model, using the full sample of the Investment Experiment 1 with 400 participants. Only the median consistent share for each participant is plotted.*

Table 2.A.15. Company Bids in the Full Sample of the Investment Experiment 1

Dependent variable:	Company Bids		
Sample:	LowComplexity (1)	HighComplexity (2)	Pooled (3)
Rational Benchmark	0.598***	0.551***	0.598***
	(0.053)	(0.057)	(0.053)
Naive Benchmark	0.463***	0.423***	0.463***
	(0.043)	(0.055)	(0.043)
Rational B. × HighComplexity			-0.047
			(0.078)
Naive B. × HighComplexity			-0.040
			(0.069)
R^2	0.888	0.816	0.854
Observations	1600	1600	3200

The table presents OLS regressions of respondents' company bids on the rational and naive benchmarks as detailed in Section 2.3.1. The table is based on the full sample of the Investment Experiment 1. Column (1) uses observations from the LowComplexity treatment, column (2) from the HighComplexity treatment, and column (3) from both. Stars highlight significant differences from 0 with * for $p < 0.10$, ** for $p < 0.05$, *** for $p < 0.01$. Standard errors are clustered on the subject level.

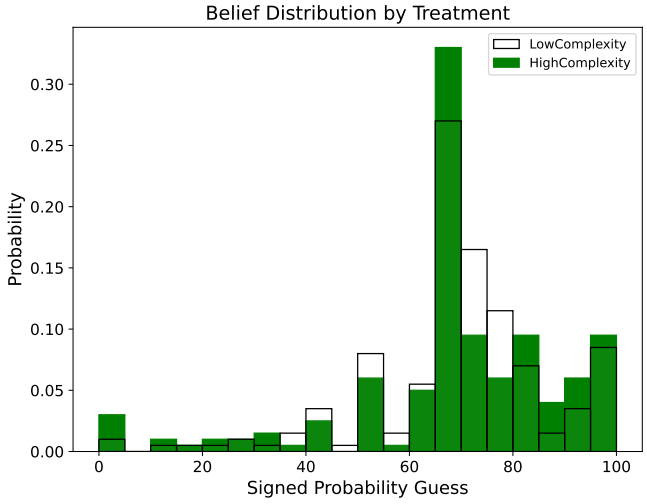

(a) Histogram of signed probability guesses.

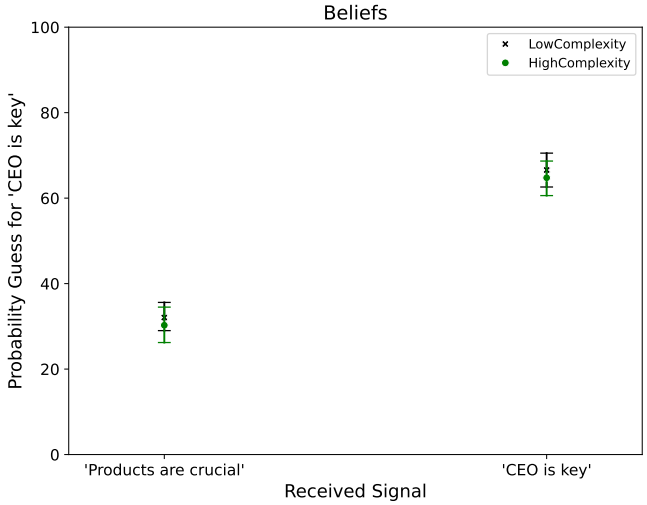

(b) Mean probability guesses for "The CEO is key".

Figure 2.A.37. Beliefs in the full sample of the Investment Experiment 1. *In Panel (a) beliefs are converted into the direction of the more likely model, so that 65 corresponds to the Bayesian probability. Panel (b) plots mean probability guesses for the model "The CEO is key". The figure is based on the full sample of the Investment Experiment 1 with 400 participants.*

Table 2.A.16. Beliefs and Recall in the Full Sample of the Investment Experiment 1

Dependent variable:	Probability Guess	Correct Recall
Sample:	Pooled (1)	Pooled (2)
Constant	67.319***	0.945***
	(1.259)	(0.016)
HighComplexity	-0.092	0.030
	(1.958)	(0.020)
R^2	0.000	0.006
Observations	400	400

The table presents OLS regressions using the full sample of the Investment Experiment 1. In column (1), the dependent variable is the probability guess for the likely state. Column (2) uses a dummy for whether respondents correctly recall the more likely model. All columns use observations from both the HighComplexity and LowComplexity conditions. Stars highlight significant differences from 0 with * for $p < 0.10$, ** for $p < 0.05$, *** for $p < 0.01$. Standard errors are clustered on the subject level.

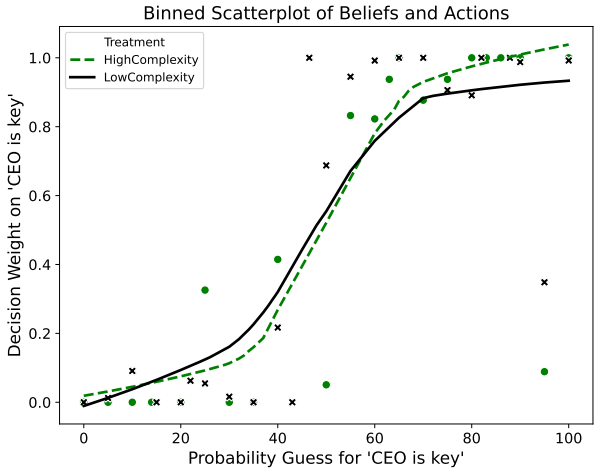

Figure 2.A.38. The relationship between decisions and beliefs in the full sample of the Investment Experiment 1. *The figure shows a binned scatterplot using the full sample of the Investment Experiment 1 with 400 participants. It has the probability guess that "The CEO is key" is the more likely model on the horizontal axis, and the decision weight γ as defined in Equation 2.3 on the vertical axis. The lines show LOWESS regressions based on all datapoints.*

2.A.2.5 Investment Experiment 2: Results in Lenient Sample

Here, we present the results for the lenient sample of the Investment Experiment 2, featuring 430 participants who solved at least one of the example screens.

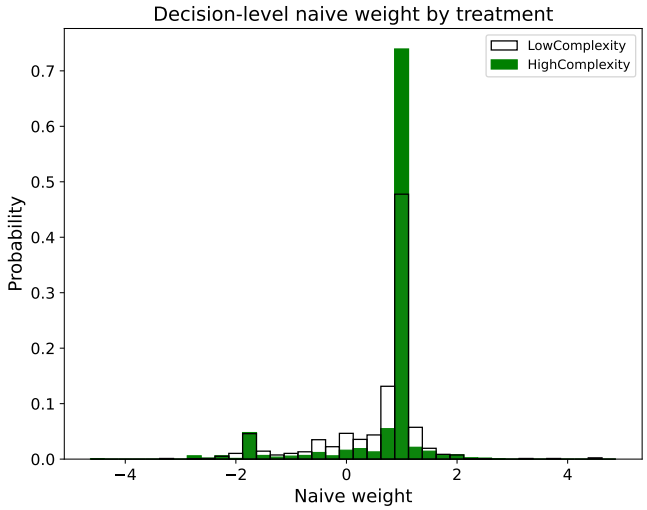

(a) Distribution of decision-level naive weights

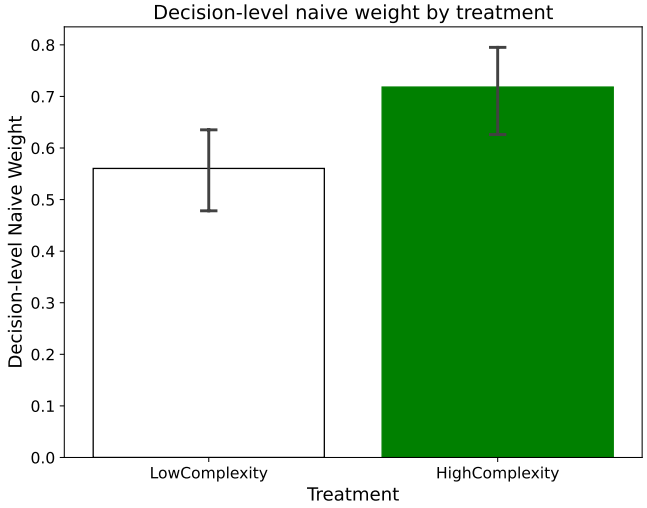

(b) Mean decision-level naive weights

Figure 2.A.39. Decision-level naive weights in the lenient sample of the Investment Experiment 2. *Panel (a) plots the distribution of naive weights λ calculated as specified in Equation 2.2, using the lenient sample of the Investment Experiment 2 with 430 participants. Panel (b) plots average naive weights.*

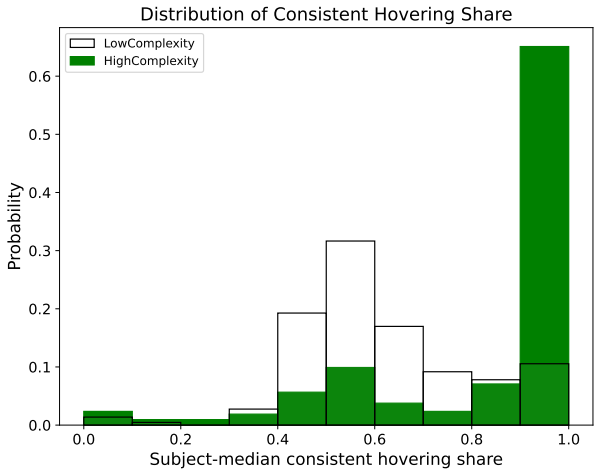

Figure 2.A.40. Distribution of subject-medians of the consistent hovering shares in the lenient sample of the Investment Experiment 2. *The figure plots the distribution of the share of time that respondents spent looking at the values of the signal-consistent model, using the lenient sample of the Investment Experiment 2 with 430 participants. Only the median consistent share for each participant is plotted.*

Table 2.A.17. Company Bids in the Lenient Sample of the Investment Experiment 2

Dependent variable:	Company Bids		
Sample:	LowComplexity (1)	HighComplexity (2)	Pooled (3)
Rational Benchmark	0.500***	0.283***	0.500***
	(0.045)	(0.043)	(0.045)
Naive Benchmark	0.532***	0.715***	0.532***
	(0.038)	(0.039)	(0.038)
Rational B. × HighComplexity			-0.218***
			(0.062)
Naive B. × HighComplexity			0.183***
			(0.054)
R^2	0.891	0.904	0.898
Observations	1744	1696	3440

The table presents OLS regressions of respondents' company bids on the rational and naive benchmarks as detailed in Section 2.3.1. The table is based on the lenient sample of the Investment Experiment 2. Column (1) uses observations from the LowComplexity treatment, column (2) from the HighComplexity treatment, and column (3) from both. Stars highlight significant differences from 0 with * for $p < 0.10$, ** for $p < 0.05$, *** for $p < 0.01$. Standard errors are clustered on the subject level.

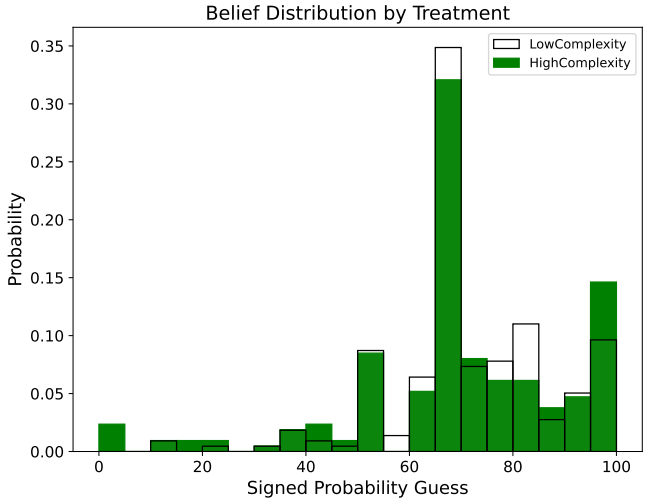

(a) Histogram of signed probability guesses.

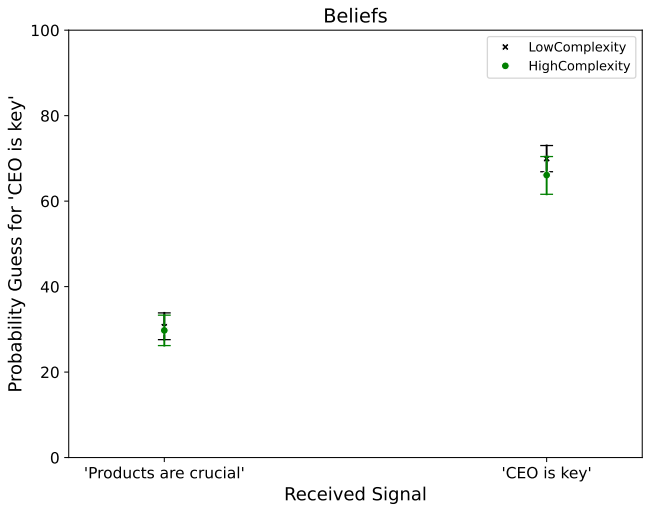

(b) Mean probability guesses for "The CEO is key".

Figure 2.A.41. Beliefs in the lenient sample of the Investment Experiment 2. *In Panel (a) beliefs are converted into the direction of the more likely model, so that 65 corresponds to the Bayesian probability. Panel (b) plots mean probability guesses for the model "The CEO is key". The figure is based on the lenient sample of the Investment Experiment 2 with 430 participants.*

Table 2.A.18. Beliefs and Recall in the Lenient Sample of the Investment Experiment 2

Dependent variable:	Probability Guess	Correct Recall
Sample:	Pooled (1)	Pooled (2)
Constant	69.656***	0.973***
	(1.096)	(0.011)
HighComplexity	-1.464	-0.020
	(1.843)	(0.018)
R^2	0.002	0.003
Observations	430	430

The table presents OLS regressions using the lenient sample of the Investment Experiment 2. In column (1), the dependent variable is the probability guess for the likely state. Column (2) uses a dummy for whether respondents correctly recall the more likely model. All columns use observations from both the HighComplexity and LowComplexity conditions. Stars highlight significant differences from 0 with * for $p < 0.10$, ** for $p < 0.05$, *** for $p < 0.01$. Standard errors are clustered on the subject level.

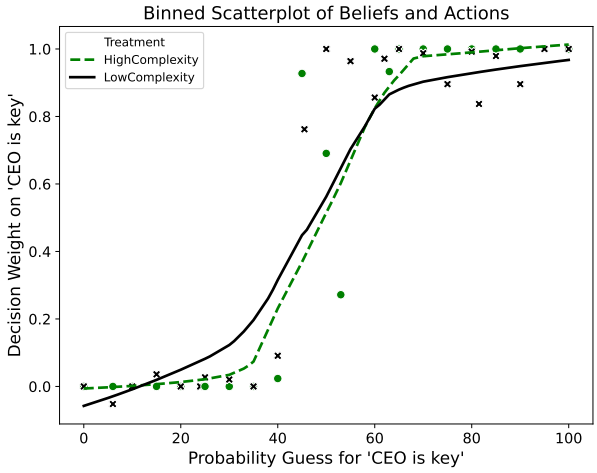

Figure 2.A.42. The relationship between decisions and beliefs in the lenient sample of the Investment Experiment 2. *The figure shows a binned scatterplot using the lenient sample of the Investment Experiment 2 with 430 participants. It has the probability guess that "The CEO is key" is the more likely model on the horizontal axis, and the decision weight γ as defined in Equation 2.3 on the vertical axis. The lines show LOWESS regressions based on all datapoints.*

2.A.2.6 Baseline Confidence: Results in Lenient Sample

Here, we present the results for the lenient sample of the Baseline Confidence Experiment, featuring 445 participants who solved at least one of the example screens.

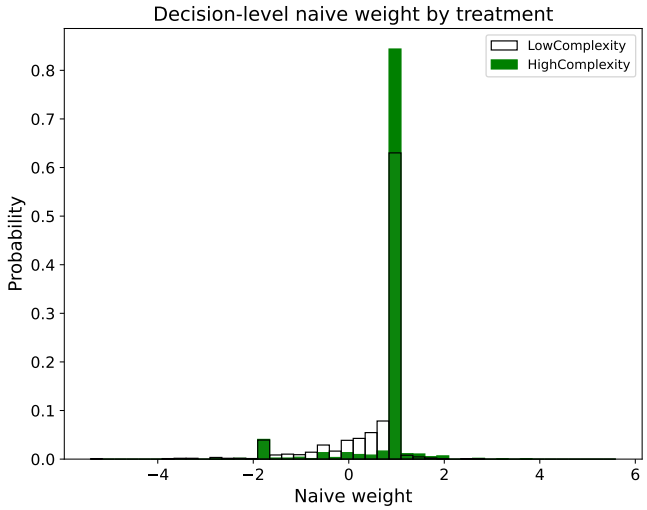

(a) Distribution of decision-level naive weights

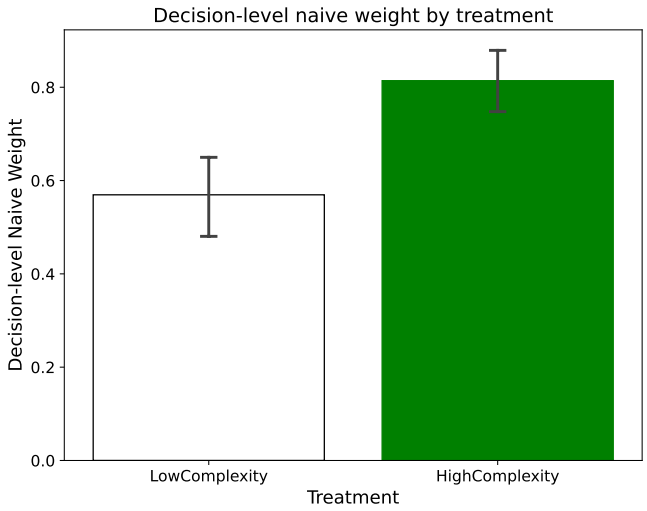

(b) Mean decision-level naive weights

Figure 2.A.43. Decision-level naive weights in the lenient sample of the Baseline Confidence Experiment. *Panel (a) plots the distribution of naive weights λ calculated as specified in Equation 2.2, using the lenient sample of the Baseline Confidence Experiment with 445 participants. Panel (b) plots average naive weights.*

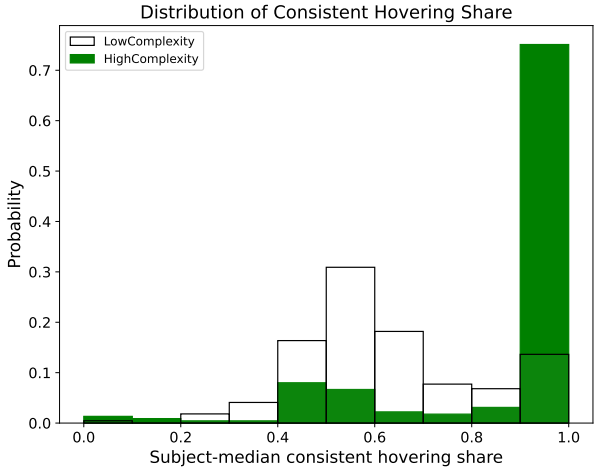

Figure 2.A.44. Distribution of subject-medians of the consistent hovering shares in the lenient sample of the Baseline Confidence Experiment. *The figure plots the distribution of the share of time that respondents spent looking at the values of the signal-consistent model, using the lenient sample of the Baseline Confidence Experiment with 445 participants. Only the median consistent share for each participant is plotted.*

Table 2.A.19. Company Value Guesses in the Lenient Sample of Baseline Confidence Experiment

Dependent variable:	Company Value Guess		
Sample:	LowComplexity (1)	HighComplexity (2)	Pooled (3)
Rational Benchmark	0.490***	0.227***	0.490***
	(0.048)	(0.037)	(0.048)
Naive Benchmark	0.560***	0.793***	0.560***
	(0.040)	(0.033)	(0.040)
Rational B. × HighComplexity			-0.263***
			(0.060)
Naive B. × HighComplexity			0.233***
			(0.052)
R^2	0.912	0.932	0.922
Observations	1760	1800	3560

The table presents OLS regressions of respondents' company value guesses on the rational and naive benchmarks as detailed in Section 2.3.1. The table is based on the lenient sample of the Baseline Confidence Experiment. Column (1) uses observations from the LowComplexity treatment, column (2) from the HighComplexity treatment, and column (3) from both. Stars highlight significant differences from 0 with * for $p < 0.10$, ** for $p < 0.05$, *** for $p < 0.01$. Standard errors are clustered on the subject level.

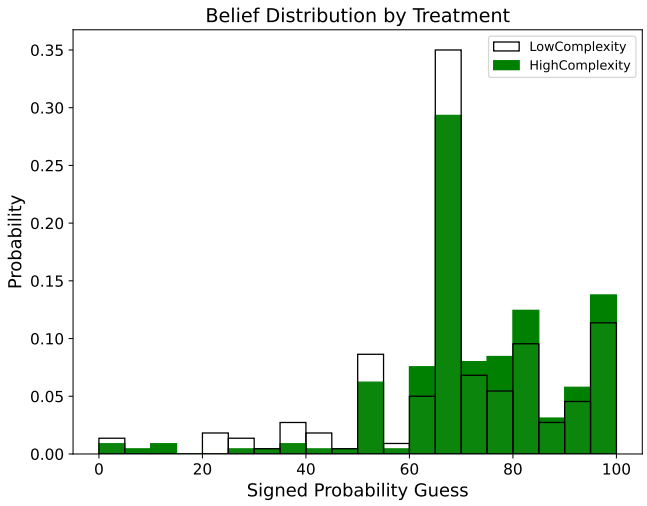

(a) Histogram of signed probability guesses.

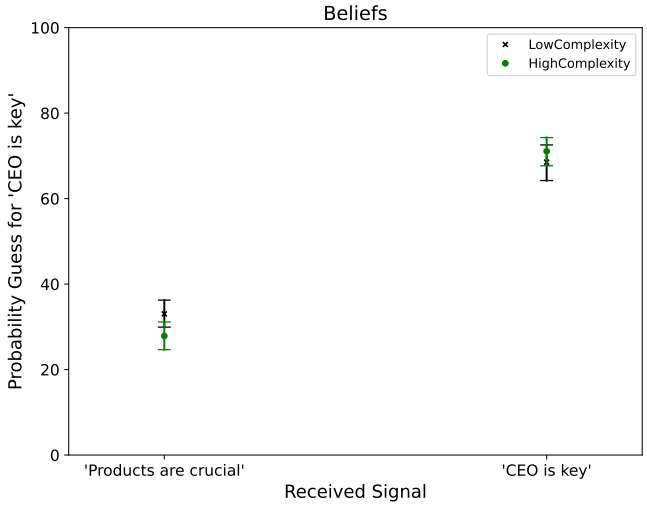

(b) Mean probability guesses for "The CEO is key".

Figure 2.A.45. Beliefs in the lenient sample of the Baseline Confidence Experiment. *In Panel (a) beliefs are converted into the direction of the more likely model, so that 65 corresponds to the Bayesian probability. Panel (b) plots mean probability guesses for the model "The CEO is key". The figure is based on the lenient sample of the Baseline Confidence Experiment with 445 participants.*

Table 2.A.20. Beliefs and Recall in the Lenient Sample of the Baseline Confidence Experiment

Dependent variable:	Probability Guess	Correct Recall
Sample:	Pooled (1)	Pooled (2)
Constant	67.759***	0.936***
	(1.315)	(0.017)
HighComplexity	3.814**	0.033
	(1.787)	(0.020)
R^2	0.010	0.006
Observations	445	445

The table presents OLS regressions using the lenient sample of the Baseline Confidence Experiment. In column (1), the dependent variable is the probability guess for the likely state. Column (2) uses a dummy for whether respondents correctly recall the more likely model. All columns use observations from both the HighComplexity and LowComplexity conditions. Stars highlight significant differences from 0 with * for $p < 0.10$, ** for $p < 0.05$, *** for $p < 0.01$. Standard errors are clustered on the subject level.

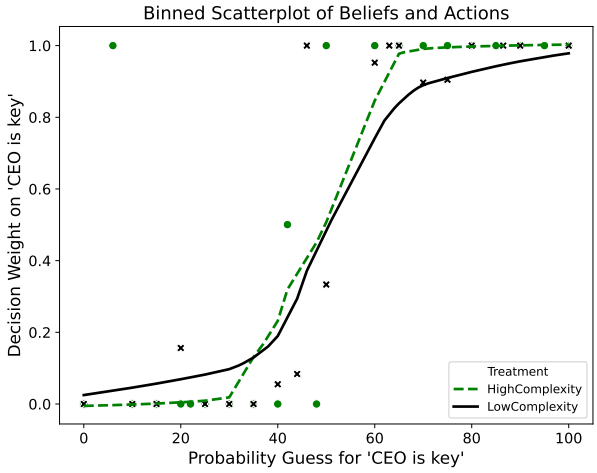

Figure 2.A.46. The relationship between decisions and beliefs in the lenient sample of the Baseline Confidence Experiment. *The figure shows a binned scatterplot using the lenient sample of the Baseline Confidence Experiment with 336 participants. It has the probability guess that "The CEO is key" is the more likely model on the horizontal axis, and the decision weight γ as defined in Equation 2.3 on the vertical axis. The lines show LOWESS regressions based on all datapoints.*

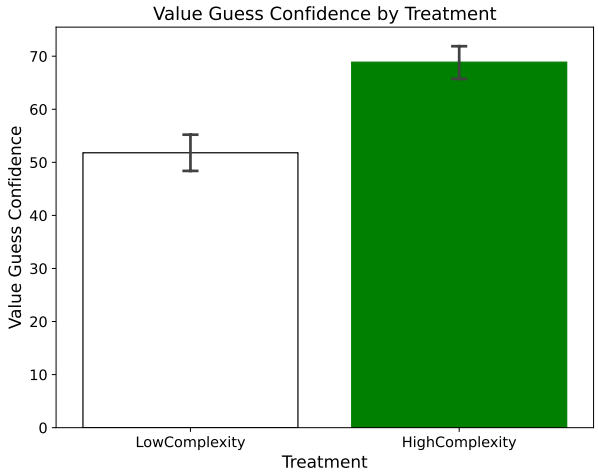

Figure 2.A.47. Average value guess confidence in the lenient sample of the Baseline Confidence Experiment. *The figure plots the average confidence that respondents had in their company value guesses, using the lenient sample of the Baseline Confidence study with 445 participants.*

2.A.2.7 Incentivized Confidence: Results in Lenient Sample

Here, we present the results for the lenient sample of the Incentivized Confidence Experiment, featuring 281 participants who solved at least one of the example screens.

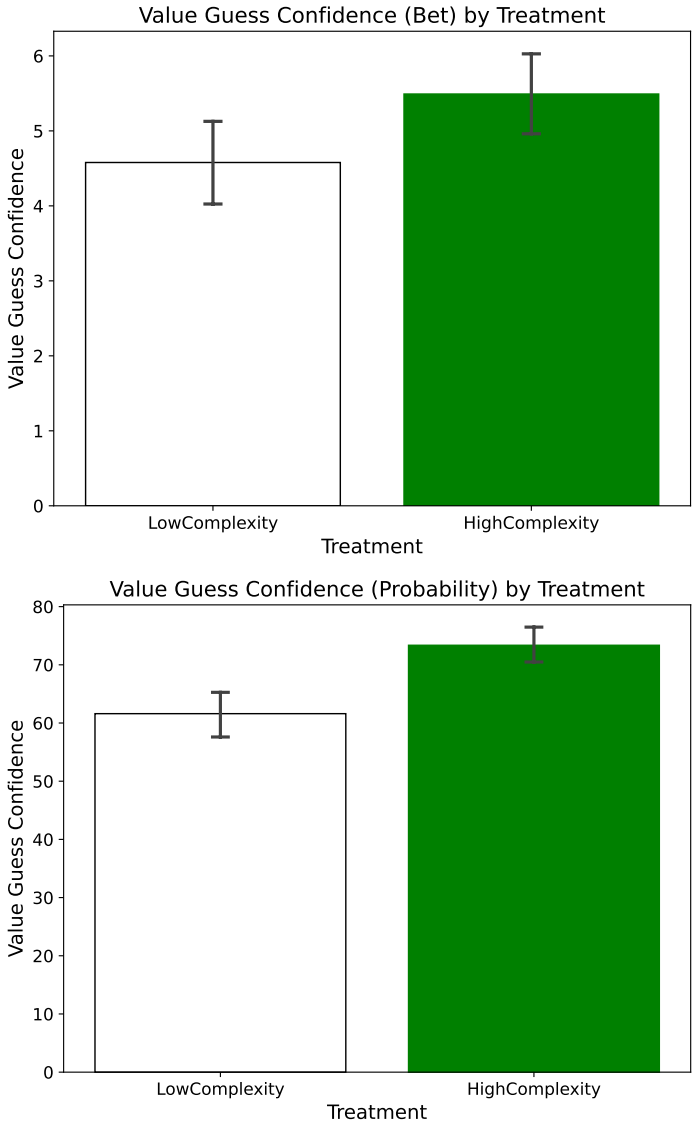

Figure 2.A.48. Average value guess confidence in the lenient sample of the Incentivized Confidence Experiment. *The top figure plots the average incentivized confidence measure, while the bottom figure plots the non-incentivized measure, both using the lenient sample of the Incentivized Confidence study with 281 participants.*

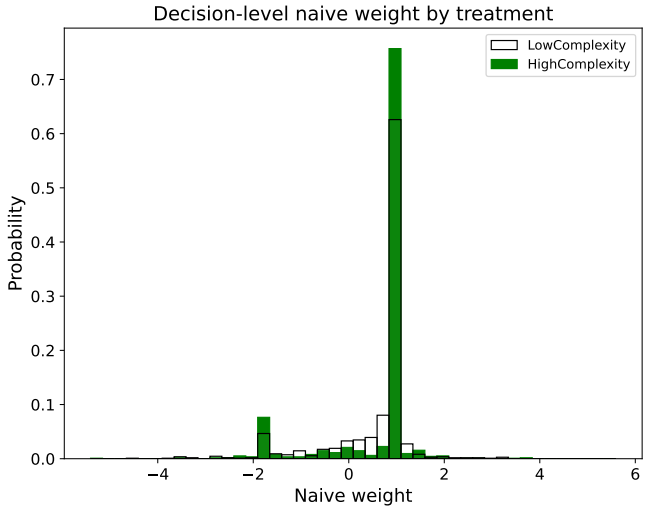

(a) Distribution of decision-level naive weights

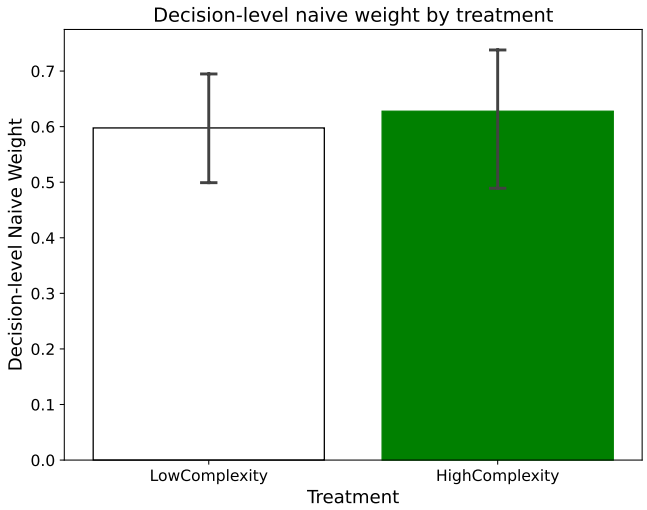

(b) Mean decision-level naive weights

Figure 2.A.49. Decision-level naive weights in the lenient sample of the Incentivized Confidence Experiment. *Panel (a) plots the distribution of naive weights λ calculated as specified in Equation 2.2, using the lenient sample of the Incentivized Confidence Experiment with 281 participants. Panel (b) plots average naive weights.*

2.A.3 Model Framework

We present a simple model that can generate the key results of our experiments. The model is an augmented and simplified version of Bordalo et al. (2025) and formalizes the construction of a mental representation of the decision problem. This process is shaped both by bottom-up and top-down attention. Upon presentation of a decision-problem, a bottom-up process of cue-dependent memory determines

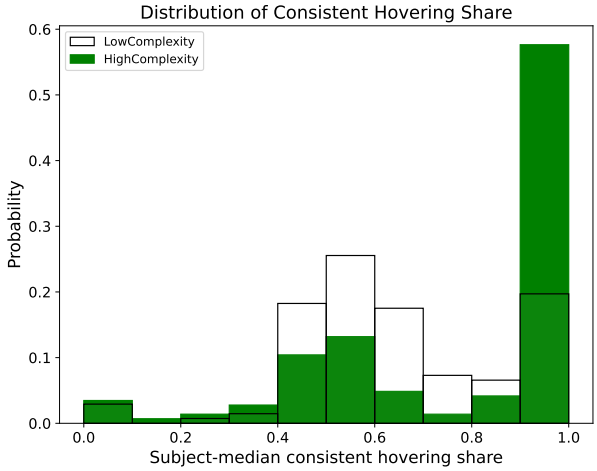

Figure 2.A.50. Distribution of subject-medians of the consistent hovering shares in the lenient sample of the Incentivized Confidence Experiment. *The figure plots the distribution of the share of time that respondents spent looking at the values of the signal-consistent model, using the lenient sample of the Incentivized Confidence Experiment with 281 participants. Only the median consistent share for each participant is plotted.*

Table 2.A.21. Company Value Guesses in the Lenient Sample of the Incentivized Confidence Experiment

Dependent variable:	Company Value Guess		
Sample:	LowComplexity (1)	HighComplexity (2)	Pooled (3)
Rational Benchmark	0.518***	0.451***	0.518***
	(0.063)	(0.071)	(0.062)
Naive Benchmark	0.564***	0.591***	0.564***
	(0.050)	(0.062)	(0.050)
Rational B. × HighComplexity			-0.067
			(0.094)
Naive B. × HighComplexity			0.027
			(0.079)
R^2	0.909	0.885	0.897
Observations	1096	1152	2248

The table presents OLS regressions of respondents' company value guesses on the rational and naive benchmarks as detailed in Section 2.3.1. The table is based on the lenient sample of the Incentivized Confidence Experiment. Column (1) uses observations from the LowComplexity treatment, column (2) from the HighComplexity treatment, and column (3) from both. Stars highlight significant differences from 0 with * for $p < 0.10$, ** for $p < 0.05$, *** for $p < 0.01$. Standard errors are clustered on the subject level.

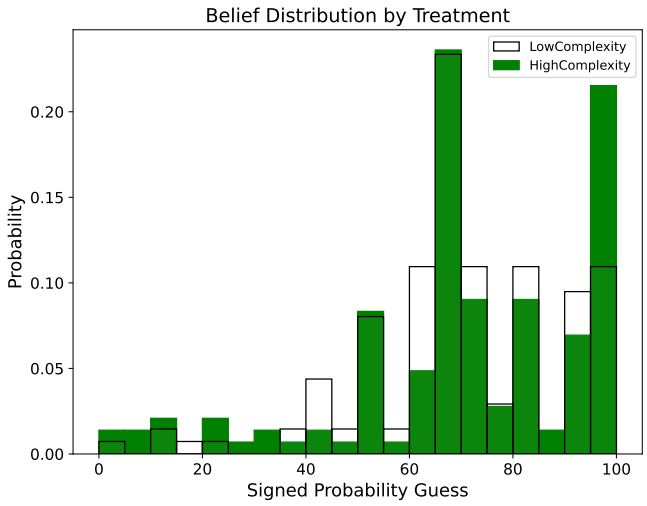

(a) Histogram of signed probability guesses.

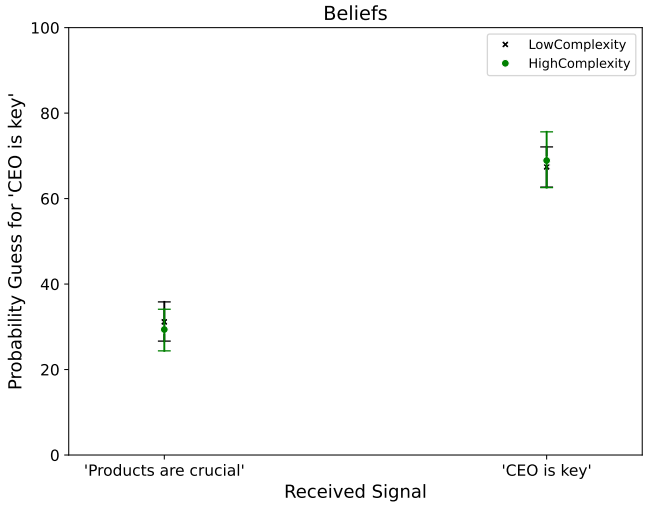

(b) Mean probability guesses for "The CEO is key".

Figure 2.A.51. Beliefs in the lenient sample of the Incentivized Confidence Experiment. *In Panel (a) beliefs are converted into the direction of the more likely model, so that 65 corresponds to the Bayesian probability. Panel (b) plots mean probability guesses for the model "The CEO is key". The figure is based on the lenient sample of the Incentivized Confidence Experiment with 281 participants.*

Table 2.A.22. Beliefs and Recall in the Lenient Sample of the Incentivized Confidence Experiment

Dependent variable:	Probability Guess	Correct Recall
Sample:	Pooled	Pooled
	(1)	(2)
Constant	68.139***	0.964***
	(1.676)	(0.016)
HighComplexity	1.714	0.002
	(2.626)	(0.022)
R^2	0.002	0.000
Observations	281	281

The table presents OLS regressions using the lenient sample of the Incentivized Confidence Experiment. In column (1), the dependent variable is the probability guess for the likely state. Column (2) uses a dummy for whether respondents correctly recall the more likely model. All columns use observations from both the HighComplexity and LowComplexity conditions. Stars highlight significant differences from 0 with * for $p < 0.10$, ** for $p < 0.05$, *** for $p < 0.01$. Standard errors are clustered on the subject level.

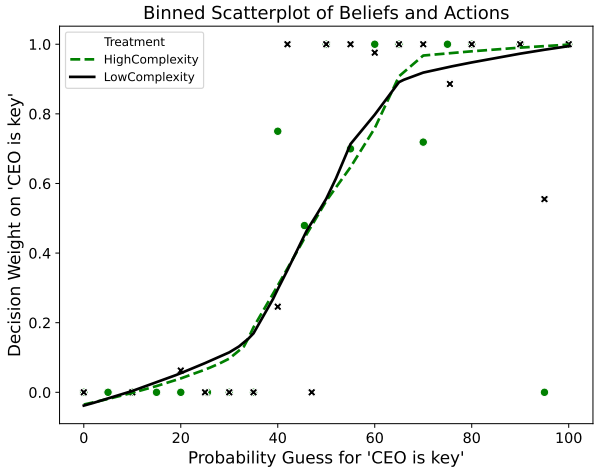

Figure 2.A.52. The relationship between decisions and beliefs in the lenient sample of the Incentivized Confidence Experiment. *The figure shows a binned scatterplot using the lenient sample of the Incentivized Confidence Experiment with 281 participants. It has the probability guess that "The CEO is key" is the more likely model on the horizontal axis, and the decision weight γ as defined in Equation 2.3 on the vertical axis. The lines show LOWESS regressions based on all datapoints.*

which of the currently stored mental representation is top of mind. Then, in a top-down process, the agent decides whether they want to further simplify this representation.

2.A.3.1 Setup

The model closely follows our experimental environment. There are two models of the world that can explain how company values are determined, i.e. $m \in \{A, B\}$. While only one model is correct, ex-ante there is a 50-50 chance of each model being correct. After observing a noisy but informative signal, the Bayesian posterior is given by $\pi = \Pr(m = A) = 0.65$, i.e. we assume without loss of generality that A is the model that is more likely to be correct.

2.A.3.2 Representations

When faced with a decision (action, confidence or belief elicitation), the agent forms a mental representation of decision problem. We assume that at any moment t the agent holds a database (a finite set) of representations

$$\boldsymbol{R}_t \;=\; \{\, r_i \;=\; (\hat{\pi}(r_i),\, V_{r_i}^{\text{context}}) \,\}$$

The scalar $\hat{\pi}(r_i)$ stores the probability attached to model A, while $V_{r_i}^{\text{context}}$ collects the contextual features of the environment in which r was formed, including the type of decision (action, confidence or belief elicitation).

Given a decision problem, the agent forms a mental representation, which we model as a two stage process. The first stage is based on bottom-up attention and similarity based recall (Bordalo et al. (2025)), while the second stage is a top-down attention decision of whether or not to simplify the representation that was recalled in the first stage.

First Stage: Bottom-Up Similarity Based Recall. When a decision screen with cue vector ξ is displayed, the agent retrieves the stored representations based on the similarity of the contextual features V_r^{context} and ξ. In particular, the recall probability $p(r)$ for a given representation is given by the relative similarity of its contextual features to the cue:

$$p(r) \;=\; \frac{S\big(\xi,\; V_r^{\text{context}}\big)}{\displaystyle\sum_{r' \in \boldsymbol{R}_t} S\big(\xi,\; V_{r'}^{\text{context}}\big)}, \tag{2.A.1}$$

where $S(\cdot,\cdot)$ is the similarity kernel and $\boldsymbol{R}_t$ is the current set of stored representations.

Second Stage: Top-Down Simplification Decision. After retrieving a non-trivial representation r with $0.5 \le \hat{\pi}(r) < 1$ the agent may keep it as it is or collapse it into a simplified version that sets $\hat{\pi} = 1$.

Optimal Action Given a Posterior π. For the case of actions (value estimates) let us first characterize the optimal estimate given some fixed representation. Under the binarised scoring rule with prize P the agent's utility is

$$u_m(a) \;=\; (1 \;-\; (\tfrac{a_m}{100} - \tfrac{a}{100})^2) \times u(P),$$

when the true model is $m \in \{A, B\}$ and a_m is the company value guess provided by model m. Hence the expected utility from guess a is

$$\mathbb{E}_{m\sim\pi}[u_m(a)] \;=\; (1 - \frac{1}{10000}[\,\pi\,(a_A - a)^2 + (1-\pi)(a_B - a)^2]) \times u(P).$$

Because P and the constant are irrelevant for the maximisation, the agent chooses the a that minimises $\pi(a_A - a)^2 + (1-\pi)(a_B - a)^2$, yielding

$$a^*(\pi) \;=\; \pi\, a_A + (1-\pi)\, a_B.$$

Complexity and Cognitive Cost. Decisions differ in their level of computational complexity $c \in \{\emptyset, \text{low}, \text{high}\}$. Here, $\emptyset$ means negligible complexity. We set $c = \text{high}$ for company-value guesses in our high-complexity treatment, $c = \text{low}$ for company-value guesses in the low-complexity treatment, and $c = \emptyset$ for the confidence and belief elicitations. Recall that we label the models so that the relevant posterior always satisfies $\pi \geq 0.5$ (i.e. model A is the more likely one), so that simplifying means collapsing to full certainty $\pi = 1$. Accordingly, we specify the cognitive cost of acting on belief π under complexity level c as

$$K(c, \pi) \;=\; (1-\pi)\kappa_c$$

where we assume

$$\kappa_\emptyset \;=\; 0 \;<\; \kappa_{\text{low}} \;<\; \kappa_{\text{high}} \quad \text{and} \quad \kappa_{\text{low}} \;<\; \frac{\Delta U(0.65)}{1 - 0.65} \;<\; \kappa_{\text{high}},$$

where

$$\Delta U(\pi) \;=\; \mathbb{E}_{m\sim\pi}[u_m(a^*(\pi))] - \mathbb{E}_{m\sim\pi}[u_m(a^*(1))]$$

is the utility cost of simplifying the representation due to loss of precision in the company value guess if the belief for model A is π. Thus, for company-value guesses, a broad representation is more costly than a simplified one, with a larger gap under high than low complexity, while for confidence and belief screens with $c = \emptyset$ no cost difference arises.

Utility Comparison for Simplification. Given a stored representation r with implied belief $\hat{\pi}(r)$ and complexity level c, maintaining the broad representation yields

$$\mathbb{E}_{m\sim\hat{\pi}(r)}[u_m\big(a^*(\hat{\pi}(r))\big)] \;-\; K\big(c,\hat{\pi}(r)\big),$$

whereas collapsing to certainty ($\pi = 1$) yields

$$\mathbb{E}_{m\sim\hat{\pi}(r)}[u_m\big(a^*(1)\big)] \;-\; K(c,1).$$

Hence, the agent simplifies if and only if

$$K(c,\hat{\pi}(r)) - K(c,1) \;>\; \mathbb{E}_{m\sim\hat{\pi}(r)}[u_m\big(a^*(\hat{\pi}(r))\big)] \;-\; \mathbb{E}_{m\sim\hat{\pi}(r)}[u_m\big(a^*(1)\big)] \;=\; \Delta U\big(\hat{\pi}(r)\big),$$

i.e. the cognitive-cost saving $K(c,\hat{\pi}(r)) - K(c,1)$ exceeds the expected-utility loss $\Delta U\big(\hat{\pi}(r)\big)$. Because $K(\emptyset,\hat{\pi}(r)) - K(\emptyset,1) = 0$, simplifying on the confidence and belief screens ($c = \emptyset$) would reduce expected utility without any cost-saving. Hence the agent never simplifies under $c = \emptyset$.

Using and Storing Representations. Once a representation has been formed and used to guide a decision, it is added to the database of representations. $V_r^{context}$ of such a representation contains the contextual features of the decision problem in which it was used.

2.A.3.3 Timeline of the Experiment

Before stating our predictions, let us briefly recap the timeline of the experiment. Respondents first read the instructions and observe the signal about which model is more likely to be correct. We assume that this induces the default representation

$$r^{\text{initial}} \;=\; (0.65,\, V^{\text{context}}_{\text{initial}}), \qquad \boldsymbol{R}_0 \;=\; \{r^{\text{initial}}\}.$$

We assume that $V^{\text{context}}_{\text{initial}}$ contains features relating to the general setup of the environment, in particular the two models, signal structure and received signal.

After going through the instructions, respondents provide the company value guesses, state their decision confidence and then finally their beliefs about which model is correct.

2.A.3.4 Estimation of Company Values

Predictions. In the high-complexity condition, respondents will simplify the initial representation to $\hat{\pi} = 1$ and choose $a^* = a_A$. In the low-complexity condition, they will keep $\hat{\pi} = 0.65$ and choose $a^* = 0.65\,a_A + 0.35\,a_B$.

Proof. Since $R_0 = \{r^{\text{initial}}\}$, participants must decide whether or not to collapse that single representation. Hence they form

$$r^{\text{guess}} = \begin{cases} (1, V^{\text{context}}_{\text{guess}}) & \text{if } K(c, 0.65) - K(c, 1) > \Delta U(0.65), \\ (0.65, V^{\text{context}}_{\text{guess}}) & \text{otherwise.} \end{cases}$$

Substituting $K(c, \pi) = (1 - \pi)\kappa_c$ shows that simplification occurs exactly when $(1 - 0.65)\kappa_c > \Delta U(0.65)$. By assumption $(1 - 0.65)\kappa_{\text{low}} < \Delta U(0.65) < (1 - 0.65)\kappa_{\text{high}}$, hence only high-complexity agents collapse to $\pi = 1$, yielding the stated actions.

2.A.3.5 Confidence

Predictions. Assuming that the reported confidence for $\hat{\pi} = 1$ exceeds that for $\hat{\pi} = 0.65$, on average, respondents in the high-complexity condition will report higher confidence than those in the low-complexity condition.

Proof. Confidence is assessed within the formed mental representation. On the confidence screen, the set of available representations is $R_1 = \{r^{\text{initial}}, r^{\text{guess}}\}$. We assume the prompt 'How certain are you that your answer is within 10 percentage points of the best possible guess?' makes the guess context more salient, so

$$S(\xi_{\text{confidence}}, V^{\text{context}}_{\text{guess}}) > S(\xi_{\text{confidence}}, V^{\text{context}}_{\text{initial}}).$$

Therefore the most likely case is that r^{guess} is retrieved. Because $c = \emptyset$ on this screen, no further simplification occurs, so high-complexity subjects retain $\hat{\pi} = 1$ and low-complexity subjects retain $\hat{\pi} = 0.65$. By the assumption that higher stored $\hat{\pi}$ yields higher reported confidence, high-complexity subjects report greater confidence.

If r^{initial} is retrieved instead, once again no simplification occurs because of $c = \emptyset$ and participants in both conditions report a low confidence associated with $\hat{\pi} = 0.65$.

2.A.3.6 Belief

Predictions. Simplification in the action space might not carry over to the belief space. The most likely outcome is that respondents in both complexity conditions will state their belief as $\hat{\pi} = 0.65$. Respondents in the high-complexity condition may exhibit a greater tendency to state the simplified belief of $\hat{\pi} = 1$.

Proof. On the belief screen, $R_2 = \{r^{\text{initial}}, r^{\text{guess}}, r^{\text{confidence}}\}$. We assume the prompt 'Which of the two rules generated correct company value estimates?' explicitly echoes the original signal description, so

$$S(\xi_{\text{belief}}, V^{\text{context}}_{\text{initial}}) > \max\{S(\xi_{\text{belief}}, V^{\text{context}}_{\text{guess}}), S(\xi_{\text{belief}}, V^{\text{context}}_{\text{confidence}})\}.$$

Thus the most likely case is that r^{initial} is retrieved. As $c = \emptyset$ here too, no collapse occurs and the stated belief is $\hat{\pi} = 0.65$ in both complexity conditions.

If r^{guess} or $r^{\text{confidence}}$ are retrieved instead, once again no simplification occurs because of $c = \emptyset$. For respondents in the high-complexity condition this implies $\hat{\pi}(r^{\text{guess}}) = \hat{\pi}(r^{\text{confidence}}) = 1$, while under low complexity $\hat{\pi}(r^{\text{guess}}) = \hat{\pi}(r^{\text{confidence}}) = 0.65$.

2.A.4 Instructions

2.A.4.1 Welcome Screen and Attention Check

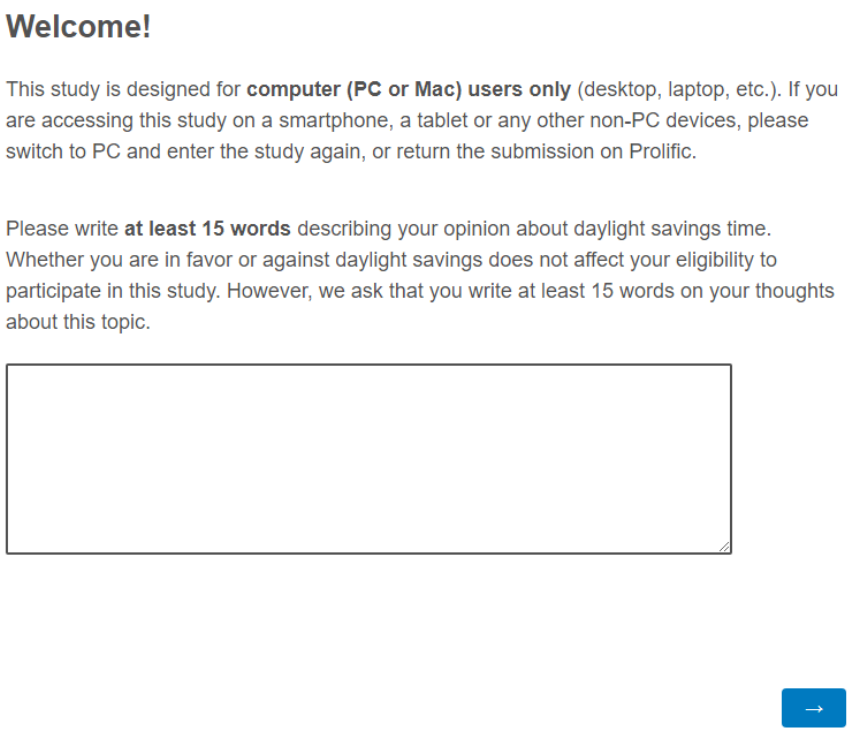

2.A.4.2 General Instructions

General Instructions

Thank you for participating in this study.

Today's survey will take approximately 20 minutes. You will earn a reward of $4.00 for participating (implying an hourly wage of $12/hr).

To earn your reward, you have to **read all instructions carefully and correctly answer the comprehension questions**.

To receive your payment, it is crucial that you pay attention throughout the whole study!

Every tenth participant has the chance to get an **additional bonus of $10.**

Feel free to use a piece of paper and pen during the study.

→

2.A.4.3 Bonus Information

Bonus

Today's survey consists of **two parts.**

If you are selected for the bonus, **an answer in one of the two parts will be randomly selected to determine your bonus**.

Therefore, you should answer all questions carefully.

→

2.A.4.4 Instructions

Part 1 - Instructions

Basic Setting

In this study, you will need to guess the value of 8 different hypothetical companies. All companies have a value between 10 and 100.

There are two rules that generate estimates for the value of the companies. One of these rules will generate the correct estimate, the other will generate an incorrect and uninformative estimate. Both rules take different variables as inputs from which the estimates can be calculated.

These two rules are:

1. **"The CEO is key"**
   Under the "The CEO is key" rule, the CEO competence **C** and supporting staff **S** are the only variables that matter.
   The estimate under this rule is given by $C \cdot S - C - S + 10$.

2. **"Products are crucial"**
   Under the "Products are crucial" rule, the number of products **P** and the research cost **R** are the only variables that matter.
   The estimate under this rule is given by $P \cdot (10 - R) + R - P$.

The realizations of the variables are different for each company and are always between 0 and 10. The company value estimates under the two rules will also differ between companies and are always between 10 and 100.

Indication of the Correct Rule

At the beginning of the study, the computer will randomly determine with a fair coin flip which of the two rules generates correct estimates. Therefore, only either the "The CEO is key" or "Products are crucial" rule will generate correct estimates for all 8 companies.

Importantly, you will not learn the outcome of the coin flip directly. Instead, we will provide you with an indication for which rule is the correct one. This indication will give you an idea for which rule is more likely to produce correct estimates, but you will not learn the correct rule with certainty.

Versions of the Study

There are two versions of the study. After you have read the instructions, the computer will randomly assign you to one of these versions.

If you are assigned to **study version 1**, you need to calculate the company value estimates under both rules yourself. For each company, you can learn about the realizations of the different variables by hovering over the variable names, which can be plugged into the formulas to obtain the company value estimates under each rule. Then you need to provide a guess about the value of that company.

If you are assigned to **study version 2,** the company value estimates under both rules will be calculated for you. For each company, you can learn about the company value estimates under both decision rules by hovering over the name of the respective rule. Then you need to provide a guess about the value of that company.

Example (Study Version 1)

The following is an example for the information you will see on the decision screen if you are assigned to study version 1.

Hover over the elements below to display the formulas:

- **"The CEO is key"**:
- **"Products are crucial"**:

Hover over the elements below to display the variables which can be used to calculate the company value estimates under both rules:

- CEO competence **C**:
- Supporting staff **S**:
- Number of products **P**:
- Research cost **R**:

Given these variable realizations, the company value estimates can be calculated under each decision rule:

- **Estimate for "The CEO is key":** $C \cdot S - C - S + 10 = 5 \cdot 4 - 5 - 4 + 10 = 21$
- **Estimate for "Products are crucial":** $P \cdot (10 - R) + R - P = 7 \cdot (10 - 2) + 2 - 7 = 51$

Example (Study Version 2)

The following is an example for the information you will see on the decision screen if you are assigned to study version 2.

Hover over the elements below to display the company value estimates under each decision rule:

- **Estimate for "The CEO is key"**:
- **Estimate for "Products are crucial"**:

Bonus Payment

We will randomly select one of your 8 answers. **The closer your guess is to the correct value of the company, the higher the likelihood that you receive the bonus of $10.** If you click on the triangle below, the precise formula will be displayed. While this formula might seem complicated, the underlying principle is very simple: the smaller the difference between your estimate and the truth, the higher the likelihood that you win $10. It is hence in your best interest to simply state your best guess.
Importantly, it does not directly matter for your earnings whether the company value is high or low. All that matters is that you guess it correctly. Your chance of receiving the bonus only depends on how close your guess is to the actual company value, which is not necessarily the higher one.

$$\text{Likelihood of winning the bonus (in \%)} = 100 - 100 \cdot (\text{Value} \div 100 - \text{Guess} \div 100)^2$$

2.A.4.5 Comprehension Questions

Comprehension Questions

Please answer the following comprehension questions. You need to answer them correctly to continue the study.

For how many companies will you need to provide a guess?

- 8
- 12
- 10

Which one of the following statements is correct?

- The rule that generates correct company value estimates might differ between the companies.
- A fair coin flip will determine which rule generates correct company value estimates. This rule will then be relevant for all companies.

Which one of the following statements is correct?

- The guesses I make in this study can affect my payoff. The study involves real stakes.
- The guesses I make in this study will not affect my payoff. The study is purely hypothetical.

The computer will secretly determine which rule generates correct company value estimates. Which of the following is true?

- The computer is more likely to select "The CEO is key".
- The computer is more likely to select "Products are crucial".
- Both rules are equally likely.

In the example of study version 1, what is the variable realization for Number of products P?

What is true regarding your chance of winning the 10$ bonus?

The chance of getting the bonus is highest when I guess the larger of the two estimates proposed by the two rules.

The chance of getting the bonus is highest when I state the guess that is closest to the true company value, which is not necessarily the higher of the two estimates.

The chance of getting the bonus is highest when I guess the sum of the two estimates proposed by the two rules.

→

2.A.4.6 Example Decision Screens for Selecting Restricted Sample

You passed the comprehension questions!

Before the actual tasks start, we give you the opportunity to familiarize yourself with the calculations of study version 1 on the next two pages.

This is not an additional comprehension check, but you can use the opportunity to get familiar with the calculations.

→

Example for Study Version 1

Hover over the elements below to display the formulas:

- **"The CEO is key"**: $C \cdot S - C - S + 10$
- **"Products are crucial"**:

Hover over the elements below to display the variables which can be used to calculate the company value estimates under both rules:

- CEO competence **C**:
- Supporting staff **S**:
- Number of products **P**:
- Research cost **R**:

What is the company value estimate under the rule **"The CEO is key"**?

(Your response needs to lie between 10 and 100.)

→

Example for Study Version 1

Hover over the elements below to display the formulas:

- **"The CEO is key"**:
- **"Products are crucial"**:

Hover over the elements below to display the variables which can be used to calculate the company value estimates under both rules:

- CEO competence **C**:
- Supporting staff **S**:
- Number of products **P**:
- Research cost **R**: 5

What is the company value estimate under the rule **"Products are crucial"**?

(Your response needs to lie between 10 and 100.)

→

2.A.4.7 Treatment Assignment - HighComplexity

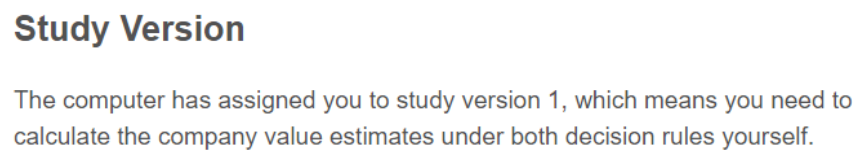

2.A.4.8 Treatment Assignment - LowComplexity

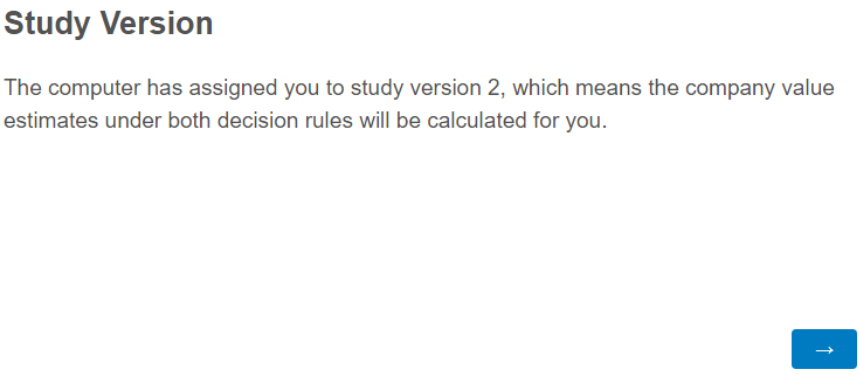

2.A.4.9 Determining the Correct Rule

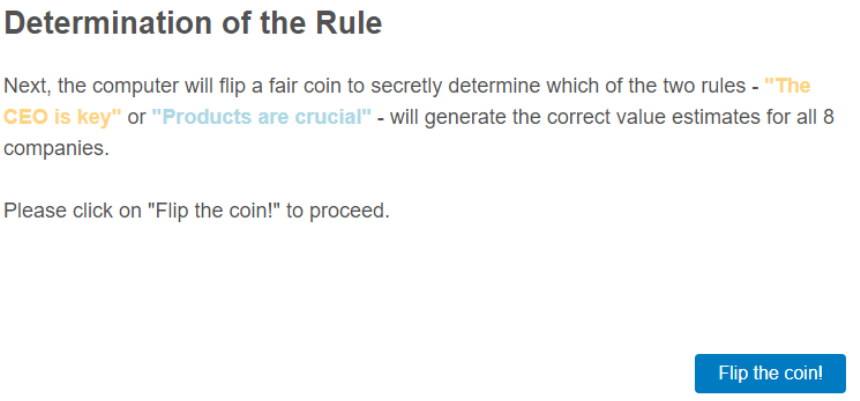

2.A.4.10 Timeline - HighComplexity

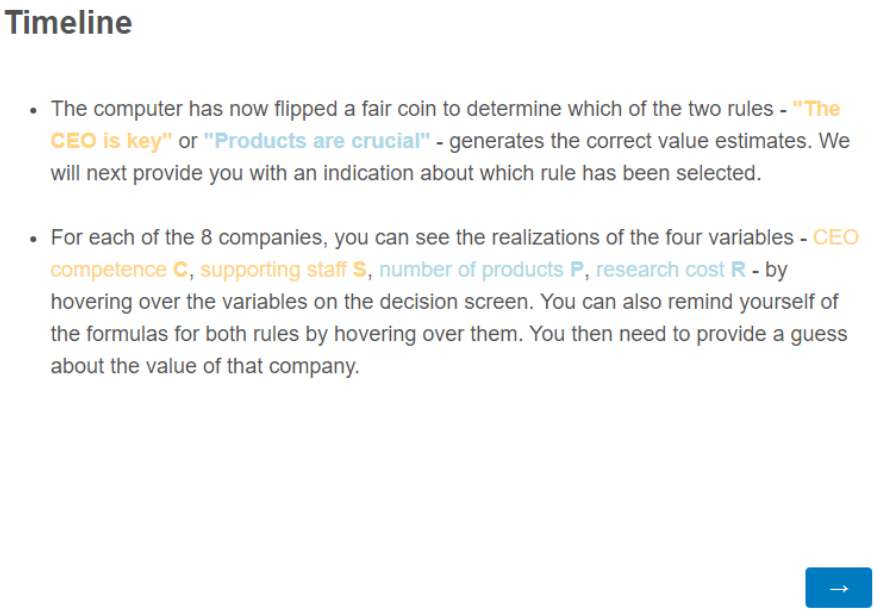

2.A.4.11 Timeline - LowComplexity

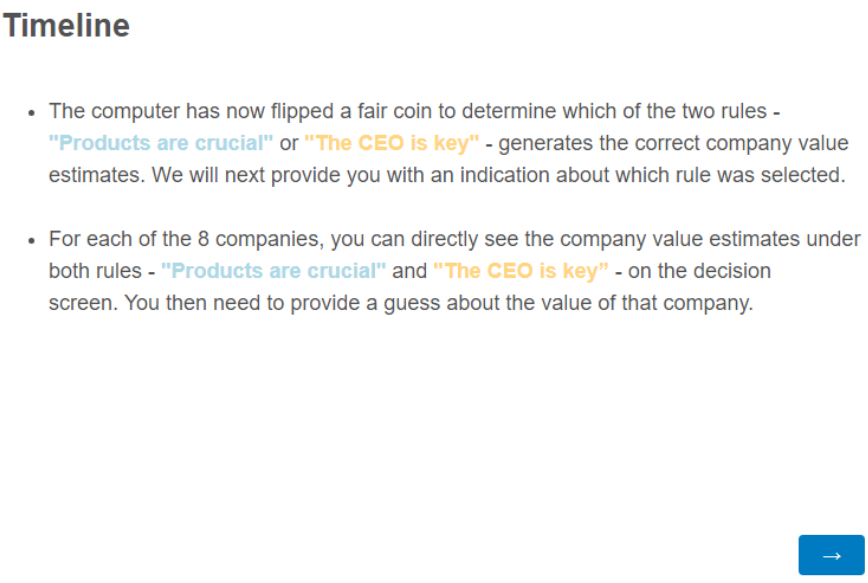

2.A.4.12 Drawing the Signal

Indication about the Rule

The computer has secretly flipped a coin to determine which of the two rules - **"The CEO is key"** or **"Products are crucial"** - generates the correct company value estimates.

On the next page, you will observe an indication for which of the two rules was randomly chosen.

You will get this information in the form of a ball that is drawn from one of the two urns displayed below.

- If **"The CEO is key"** generates correct estimates, the computer will draw a single ball from **Urn A, which contains 65 orange balls and 35 blue balls**.
- If **"Products are crucial"** generates correct estimates, the computer will draw a single ball from **Urn B, which contains 35 orange balls and 65 blue balls**.

This means that a **orange ball** is rather indicative of **"The CEO is key"**, and a **blue ball** is rather indicative of **"Products are crucial"**. Notice again that these are only indications**, you will not know with certainty which rule has been selected** by the computer.

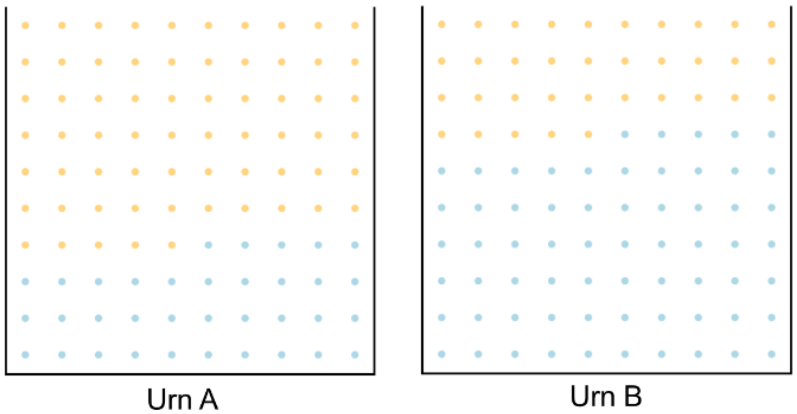

Draw the ball!

2.A.4.13 Signal Screen - CEO is key

The drawn ball is orange, which means **"The CEO is key"** is more likely to generate correct estimates.

You can proceed to the next screen in 15 seconds.

2.A.4.14 Signal Screen - Products are crucial

The drawn ball is blue, which means **"Products are crucial"** is more likely to generate correct estimates.

You can proceed to the next screen in 15 seconds.

2.A.4.15 Signal Check

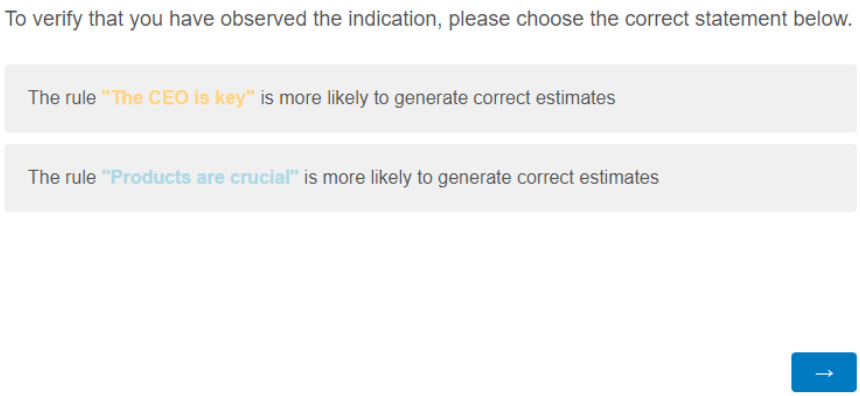

2.A.4.16 Decision Screen - HighComplexity

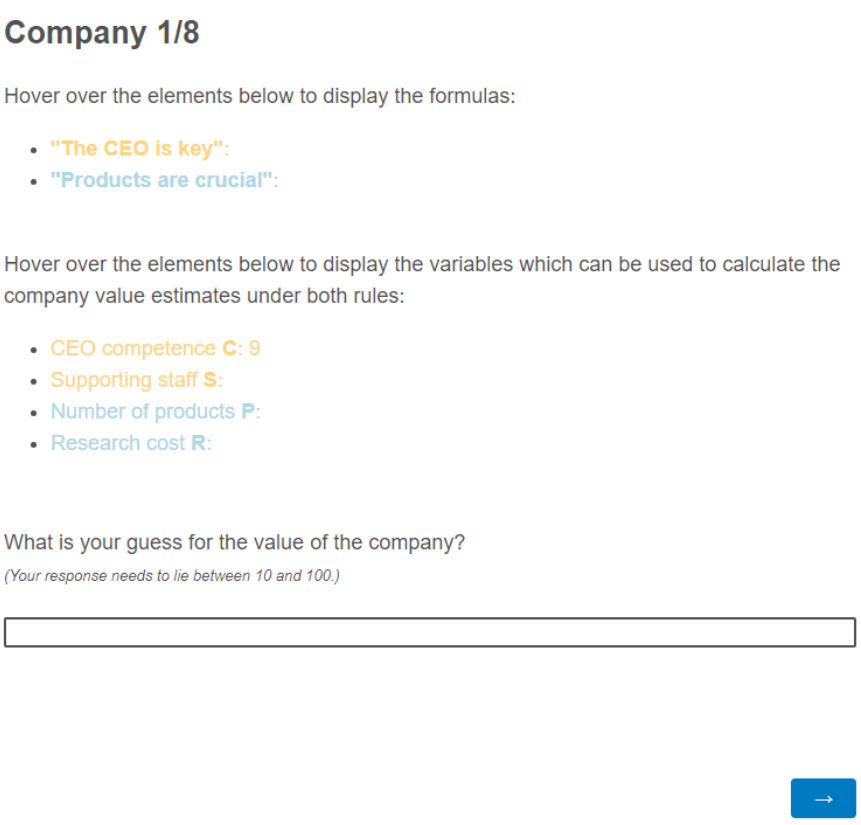

The formulas and variable realizations were displayed upon hovering over the respective text.

2.A.4.17 Decision Screen - LowComplexity

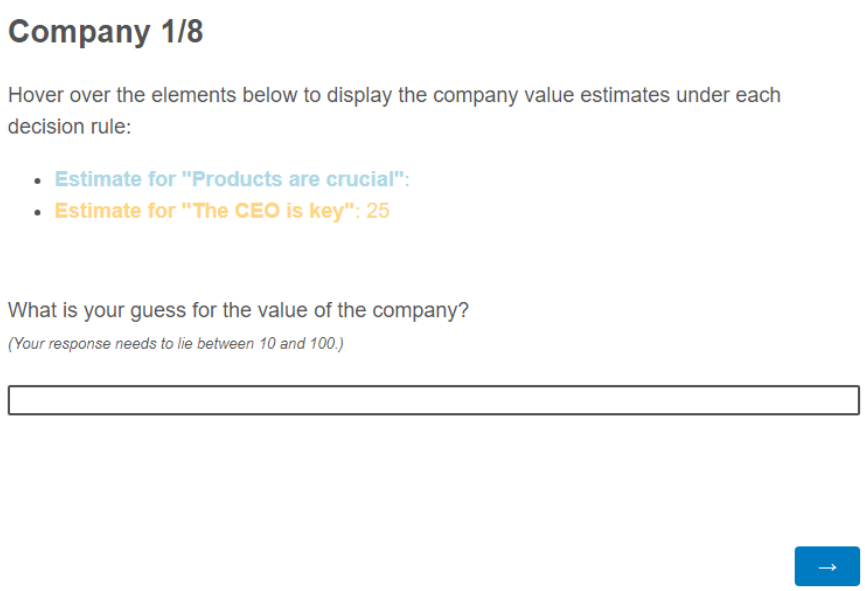

The pre-calculated company value estimates were displayed upon hovering over the respective text.

2.A.4.18 Instructions

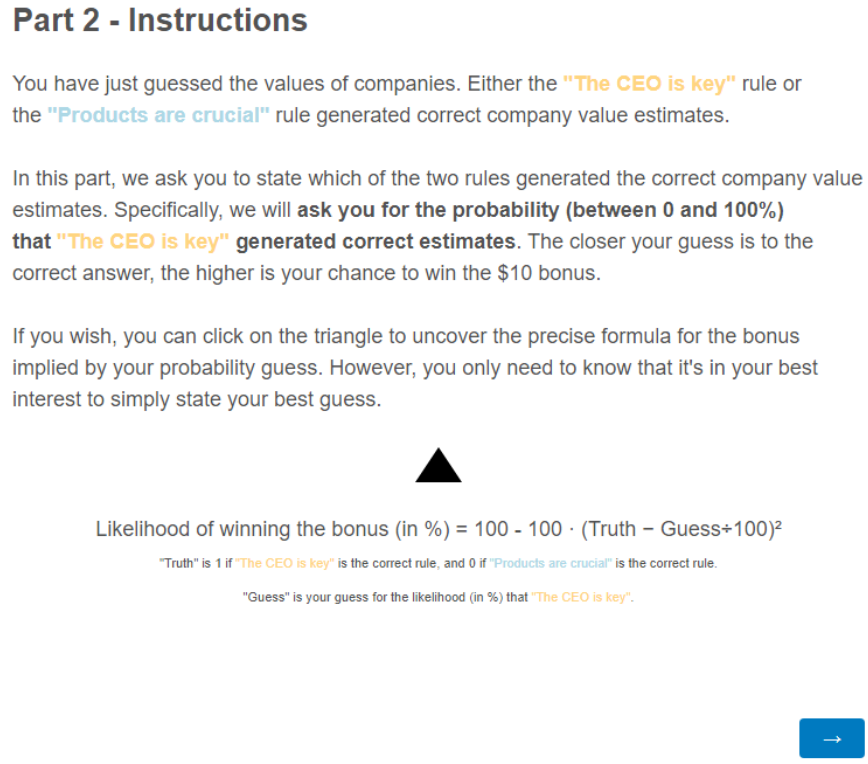

2.A.4.19 Decision Screen

In the first part, which of the two rules - "The CEO is key" or "Products are crucial" - generated correct company value estimates?

"The CEO is key": 34 %.

"Products are crucial": 66%.

→

2.A.4.20 Confidence Elicitation

You guessed that the likelihood for "The CEO is key" is 34%.

How certain are you that this answer is within 10 percentage points of the best possible guess given the information you received in the first part?

(This question does not affect your bonus.)

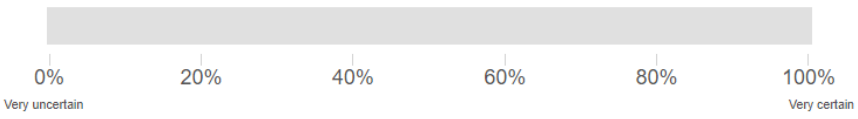

2.A.4.21 Direct Recall

In the first part, before providing guesses for the company values, you received an indication for which rule was selected by the computer to generate correct company value estimates.

Do you recall which rule was indicated to be more likely to generate correct value estimates?

(This question does not affect your bonus.)

"The CEO is key"

"Products are crucial"

I don't remember

→

References

Aina, Chiara. 2024. "Tailored Stories." *Working Paper.* [64]

Aina, Chiara, and Florian Schneider. 2025. "Weighting Competing Models." *Working Paper.* [64]

Ameriks, John, et al. 2020. "Heterogeneity in expectations, risk tolerance, and household stock shares: The attenuation puzzle." *Journal of Business & Economic Statistics* 38 (3): 633–46. [63, 79]

Andre, Peter, et al. 2024. "Narratives about the Macroeconomy." *Working Paper.* [64]

Arrieta, Gonzalo, and Kirby Nielsen. 2024. "Procedural Decision-Making In The Face Of Complexity." *Working Paper.* [65]

Augenblick, Ned, Eben Lazarus, and Michael Thaler. Forthcoming. "Overinference from Weak Signals and Underinference from Strong Signals." *Quarterly Journal of Economics.* [65]

Ba, Cuimin, J Aislinn Bohren, and Alex Imas. 2022. "Over-and underreaction to information." *Available at SSRN 4274617.* [63, 65, 67]

Barron, Kai, and Tilman Fries. 2024. "Narrative Persuasion." *Working Paper.* [62, 64]

Benjamin, Daniel J. 2019. "Errors in probabilistic reasoning and judgment biases." *Handbook of Behavioral Economics: Applications and Foundations 1* 2: 69–186. [67]

Beutel, Johannes, and Michael Weber. 2023. "Beliefs and portfolios:Causal evidence." *Working Paper.* [63, 79]

Bordalo, Pedro, Nicola Gennaioli, and Andrei Shleifer. 2012. "Salience theory of choice under risk." *Quarterly journal of economics* 127 (3): 1243–85. [63, 64]

Bordalo, Pedro, Nicola Gennaioli, and Andrei Shleifer. 2020. "Memory, Attention, and Choice." *Quarterly Journal of Economics* 135 (3): 1399–442. https://doi.org/10.1093/qje/qjaa012. [84]

Bordalo, Pedro, et al. 2023a. "How People Use Statistics." *Working Paper.* [63, 64]

Bordalo, Pedro, et al. 2023b. "Memory and Probability." *Quarterly Journal of Economics* 138 (1): 265–311. [63, 64, 84]

Bordalo, Pedro, et al. 2025. "A Cognitive Theory of Reasoning and Choice." *Working Paper.* [64, 84, 137, 141]

Bushong, Benjamin, Matthew Rabin, and Joshua Schwartzstein. 2021. "A model of relative thinking." *Review of Economic Studies* 88 (1): 162–91. [64]

Caplin, Andrew, Mark Dean, and John Leahy. 2019. "Rational inattention, optimal consideration sets, and stochastic choice." *Review of Economic Studies* 86 (3): 1061–94. [64]

Charles, Constantin, Cary Frydman, and Mete Kilic. 2024. "Insensitive investors." *Journal of Finance* 79 (4): 2473–503. [80]

Danz, David, Lise Vesterlund, and Alistair Wilson. 2022. "Belief Elicitation and Behavioral Incentive Compatibility." *American Economic Review* 112 (9): 2851–83. [66]

Dertwinkel-Kalt, Markus, et al. 2022. "Concentration bias in intertemporal choice." *Review of Economic Studies* 89 (3): 1314–34. [64]

Eliaz, Kfir, and Rani Spiegler. 2022. "A Model of Competing Narratives." *American Economic Review* 110 (12): 3786–816. [64]

Enke, Benjamin. 2020. "What you see is all there is." *Quarterly Journal of Economics* 135 (3): 1363–98. [65]

Enke, Benjamin. 2024. "The Cognitive Turn in Behavioral Economics." *Working Paper.* [65]

Enke, Benjamin, and Thomas Graeber. 2023. "Cognitive Uncertainty." *Quarterly Journal of Economics* 138 (4): 2021–67. [64, 67, 81]

Enke, Benjamin, Thomas Graeber, and Ryan Oprea. 2023. "Confidence, Self-selection and Bias in the Aggregate." *American Economic Review* 113 (7): 1933–66. [64, 81]

Enke, Benjamin, Thomas Graeber, and Ryan Oprea. Forthcoming. "Complexity and Time." *Journal of the European Economic Association.* [64, 81]

Enke, Benjamin, Frederik Schwerter, and Florian Zimmermann. 2024. "Associative Memory, Beliefs and Market Interactions." *Journal of Financial Economics* 157: 103853. https://doi.org/10.1016/j.jfineco.2024.103853. [84]

Enke, Benjamin, and Cassidy Shubatt. 2024. "Quantifying Lottery Choice Complexity." *Working Paper.* [65]

Enke, Benjamin, and Florian Zimmermann. 2019. "Correlation neglect in belief formation." *review of economic studies* 86 (1): 313–32. [65]

Enke, Benjamin, et al. 2024. "Behavioral Attenuation." Working paper. National Bureau of Economic Research. [80]

Esponda, Ignacio, Ryan Oprea, and Sevgi Yuksel. 2023. "Seeing What is Representative." *Quarterly Journal of Economics* 138 (4): 2607–57. [65]

Fan, Tony, Yucheng Liang, and Cameorn Peng. 2024. "The Inference-Forecast Gap in Belief Updating." *Working Paper.* [65]

Frick, Mira, Ryota Iijima, and Yuhta Ishii. 2022. "Dispersed behavior and perceptions in assortative societies." *American Economic Review* 112 (9): 3063–105. [62, 64]

Gabaix, Xavier. 2014. "A sparsity-based model of bounded rationality." *Quarterly Journal of Economics* 129 (4): 1661–710. [64]

Gagnon-Bartsch, Tristan, Matthew Rabin, and Joshua Schwartzstein. 2023. "Channeled attention and stable errors." *Working Paper.* [64]

Giglio, Stefano, et al. 2021. "Five facts about beliefs and portfolios." *American Economic Review* 111 (5): 1481–522. [63, 79]

Graeber, Thomas. 2023. "Inattentive inference." *Journal of the European Economic Association* 21 (2): 560–92. [65]

Graeber, Thomas, Christopher Roth, and Constantin Schesch. 2024. "Explanations." *Working Paper.* [64]

Graeber, Thomas, Christopher Roth, and Florian Zimmermann. 2024. "Stories, statistics, and memory." *Quarterly Journal of Economics.* [64, 84]

Hartzmark, Samuel, Samuel Hirshman, and Alex Imas. 2021. "Ownership, Learning, and Beliefs." *Quarterly Journal of Economics* 136 (3): 1665–717. [65]

Heidhues, Paul, Botond Kőszegi, and Philipp Strack. 2018. "Unrealistic Expectations and Misguided Learning." *Econometrica* 86 (4): 1159–214. [62, 64]

Heidhues, Paul, Botond Kőszegi, and Philipp Strack. 2023. "Misinterpreting Yourself." *Working Paper.* [64]

Jiang, Zhengyang, et al. Forthcoming. "Investor Memory and Biased Beliefs: Evidence from the Field." *Quarterly Journal of Economics.* [84]

Kendall, Chad, and Ryan Oprea. 2024. "On the Complexity of Forming Mental Models." *Quantitative Economics* 15: 175–211. [65]

Kőszegi, Botond, and Filip Matějka. 2020. "Choice simplification: A theory of mental budgeting and naive diversification." *Quarterly Journal of Economics* 135 (2): 1153–207. [64]

Kőszegi, Botond, and Adam Szeidl. 2013. "A model of focusing in economic choice." *Quarterly journal of economics* 128 (1): 53–104. [64]

Laudenbach, Christina, et al. Forthcoming. "Beliefs About the Stock Market and Investment Choices: Evidence from a Field Experiment." *Review of Financial Studies.* [63, 80]

Mailath, George J, and Larry Samuelson. 2019. "The Wisdom of a Confused Crowd: Model-Based Inference." *Working Paper*. [64]

Montiel Olea, José Luis, et al. 2022. "Competing models." *Quarterly Journal of Economics* 137 (4): 2419–57. [64]

Oprea, Ryan. 2020. "What Makes a Rule Complex?" *American Economic Review* 110 (12): 3913–51. [65]

Schwartzstein, Joshua. 2014. "Selective attention and learning." *Journal of the European Economic Association* 12 (6): 1423–52. [64]

Schwartzstein, Joshua, and Adi Sunderam. 2021. "Using models to persuade." *American Economic Review* 111 (1): 276–323. [62, 64]

Shiller, Robert. 2017. "Narrative Economics." *American Economic Review* 107 (4): 967–1004. [64]

Shubatt, Cassidy, and Jeffrey Yang. 2024. "Tradeoffs and Comparison Complexity." Working paper. Working Paper. [65]

Somerville, Jason. 2022. "Range-Dependent Attribute Weighting in Consumer Choice: An Experimental Test." *Econometrica* 90 (2): 799–830. [64]

Yang, Jeffrey. 2023. "On the Decision-Relevance of Subjective Beliefs." *Available at SSRN 4425080*. [63, 80]

Chapter 3

Misperceptions and Politically Motivated Reasoning*

Joint with Guy Yanay

3.1 Introduction

Mass political polarization is not a uniquely American phenomenon, but its rise in the United States since the 1990s has highlighted its prevalence and the threat it poses to democracies around the world. Beyond a substantial affective component, partisan divisions in factual beliefs are widely perceived, with a 2025 Pew Research Center poll reporting that 80% of Americans believe Republican and Democratic voters disagree on basic facts.[1] A common explanation for these gaps is politically motivated reasoning—a bias in information processing that leads individuals across the political spectrum to draw different conclusions from identical evidence by interpreting its strength and relevance through the lens of their political identity and prior beliefs.

This chapter presents results from a series of experiments designed to examine politically motivated reasoning, measure its scope, and assess whether both the general public and the academic community hold realistic expectations about its prevalence. The first experiment tests whether individuals engage in politically motivated reasoning when updating their own beliefs. Participants complete incentivized estimation tasks on politically charged topics, make initial estimates, receive directional feedback, and then revise their beliefs. The second and third experiments investi-

* Funding by the Deutsche Forschungsgemeinschaft (DFG) through CRC TR 224 (Project A01) and under Germany's Excellence Strategy - EXC 2126/1 - 390838866 is gratefully acknowledged. Proofreading assistance for this chapter was provided by the LLM ChatGPT; the research design, analyses, and all arguments remain solely the work of the authors.

1. https://www.pewresearch.org/short-reads/2025/07/30/most-americans-say-republican-and-democratic-voters-cannot-agree-on-basic-facts/

gate people's expectations regarding the beliefs and belief updating of others. In these studies, participants are asked to guess both the initial and revised estimates made by participants in Study 1, allowing us to assess expectations of politically motivated reasoning in others. Studies 1 and 2 recruited participants through the online platform Prolific, whereas Study 3 involved an expert sample surveyed via the Social Science Prediction Platform.

When examining actual belief updating, we find no evidence of politically motivated reasoning on average. Participants respond similarly to both pro-party and anti-party signals when updating their beliefs. However, belief updating is generally more pronounced for political topics than for non-political ones. Participants who are more affectively polarized, or who prefer interacting with political in-group members, show greater responsiveness to pro-party information. Additionally, we document partisan asymmetries: Democrats tend to respond more strongly to pro-party signals, whereas Republicans show greater responsiveness to anti-party information.

Our main results on expected belief updating show that both Prolific participants and experts from the Social Science Prediction Platform expect that others do engage in politically motivated reasoning, i.e. update their beliefs more strongly in response to signals aligned with their political affiliation. This pattern persists when controlling for the target's actual prior belief but reverses when accounting for the expected prior that participants attribute to the target. This may suggest that the expected asymmetry reflects exaggerated assumptions about the extremity of others' political beliefs, rather than accurate predictions about belief updating itself. Moreover, experts anticipate a significantly higher degree of motivated reasoning than Prolific participants.

Our work contributes to several strands of related literature. Research on affective polarization documents increasing partisan hostility in the U.S. and its consequences for perceptions and behavior (see Iyengar et al. (2019) for a review). Lelkes (2016) provides evidence of increasing ideological polarization—i.e., divergence in political opinions—between partisans in the United States. Beyond partisan preferences and affect, individuals also differ in their factual beliefs (Nyhan (2020), Rekker and Harteveld (2022), Lee et al. (2021), Flynn, Nyhan, and Reifler (2017)). These forms of polarization are often interconnected and have important manifestations in domains such as immigration (Sides and Citrin (2007), Gorodzeisky and Semyonov (2020)) and inequality (Kuziemko et al. (2015)). We contribute to this literature by studying factual belief polarization in a controlled, incentivized setting, before and after participants receive truthful and informative information.

Several explanations have been proposed for the rise in political polarization. On one hand, numerous studies have documented that partisans select into news sources that align with and reinforce their pre-existing political views (Gentzkow and Shapiro (2011), Iyengar and Hahn (2009), Peterson, Goel, and Iyengar (2021), Garrett (2009), Chopra, Haaland, and Roth (2024)). On the other hand, politically

motivated reasoning refers to the tendency of partisans to interpret information through the lens of their prior beliefs, leading to divergent conclusions even when presented with the same evidence (Kahan, Jenkins-Smith, and Braman (2011), Kahan (2013), Taber and Lodge (2006), Lodge and Taber (2013)). In this study, we isolate politically motivated reasoning by holding information exposure constant and using incentivized belief elicitation.

Another strand of literature documents systematic misperceptions about others' political beliefs and characteristics. Individuals frequently overestimate both the extremity of their political opponents' views and the magnitude of the ideological divide (Ahler (2014), Bursztyn and Yang (2022), Westfall et al. (2015)), and misperceive the demographic composition of political parties, which in turn shapes attitudes and stereotypes (Ahler and Sood (2018)). Relatedly, Agranov and Detkova (2024) examine expectations about belief updating in non-political contexts, showing that individuals systematically misperceive how others incorporate new evidence. We contribute to this literature by examining beliefs about political information processing—specifically, individuals' expectations of how others respond to information that aligns with or challenges their political identity. Importantly, we find systematic overestimation of the extent to which others interpret political information in a biased manner.

Finally, our research contributes to a large economic literature on motivated reasoning and belief updating. To date, this literature has primarily examined ego-relevant domains such as intelligence (Eil and Rao (2011), Möbius et al. (2022), Zimmermann (2020)), and financial decision-making (Gotthard-Real (2017), Coutts (2019), Barron (2021), Gödker, Jiao, and Smeets (2025)). Recently, Thaler (2024) extended this work to political contexts, focusing on how individuals evaluate the trustworthiness of political information, whereas our study directly examines belief updating.

The rest of the chapter is structured as follows. Section 3.2 presents our research design. Section 3.3 presents our main results, followed by Section 3.4 which delves into heterogeneity analyses. Section 3.5 then concludes.

3.2 Research Design

3.2.1 Experimental Design

We conducted three studies. Study 1 examines whether individuals engage in motivated reasoning when updating their own beliefs. Studies 2 and 3 explore individuals' expectations about others' belief updating behavior. Studies 1 and 2 were conducted with participants recruited via the survey platform Prolific, while Study 3 was an expert survey administered through the Social Science Prediction Platform.

All studies were preregistered, including details on the design, sample size, exclusion criteria, and main hypotheses.[2] The questionnaires for all studies can be found in Appendix 3.A.4.

3.2.1.1 Study 1: Motivated Reasoning

We recruited U.S.-based participants via Prolific who identified with either the Republican or Democratic Party. Participants answered a series of probabilistic questions covering both politically charged and neutral topics. All responses were expressed as percentages from 0% to 100%. For example, a politically charged question asked: "What share of those convicted of fentanyl trafficking between 2018 and 2021 were illegal immigrants?" A politically neutral example was: "What share of U.S. presidents were left-handed?" The political questions were directionally framed such that higher estimates aligned with the typical views of one party (i.e., the Republican or Democratic Party), while lower estimates aligned with the opposing party. We refer to questions aligned with Republican views as *Republican-coded*, and those aligned with Democratic views as *Democratic-coded*. All 12 questions are listed in Appendix 3.A.3.

In the first stage, participants provided initial guesses for five of the twelve questions. Each participant was randomly assigned five questions: two Democratic-coded, two Republican-coded, and one neutral. After each estimate, participants reported their confidence in the accuracy of their guess on a scale from 1 to 7. They also responded to follow-up questions regarding the importance of the issue and their beliefs about the average guesses of Republicans and Democrats.

In the second stage, participants received feedback on each question they had answered in the first stage. They were truthfully informed whether the correct answer was higher or lower than their initial guess and were then asked to revise their response.

Both initial and revised guesses were incentivized. For each participant, one response was randomly selected to determine their probability of earning a $1.50 bonus. The bonus probability decreased with the squared distance between the participant's estimate and the true value, ensuring incentive compatibility regardless of risk preferences (similar to Hossain and Okui (2013)).[3] Comprehension checks ensured participants understood the incentive mechanism and the truthfulness of the feedback.

At the end of the study, we collected political demographic measures, including party preference and strength of identification (measured on a 1–7 Likert scale). We

2. The preregistration can be found at https://www.socialscienceregistry.org/trials/15498/history/255622.

3. Specifically, the probability of winning the bonus (in percent) was given by $100 - 100 \times (\text{Guess}/100 - \text{Truth}/100)^2$.

also administered a feeling thermometer, where participants rated their warmth toward Republicans and Democrats on a scale from 0 (very cold/unfavorable) to 10 (very warm/favorable). Finally, we elicited participants' willingness to take part in various potential follow-up studies, both political and non-political, featuring interactions with political peers or members of the opposite party.

3.2.1.2 Study 2: Expectations of Motivated Reasoning

As in Study 1, we recruited U.S.-based participants via Prolific who identified as either Democrats or Republicans. Participants in Study 2 (henceforth *evaluators*) were asked to infer the responses of participants from Study 1 (henceforth *targets*).

The first two stages of Study 2 followed a shortened version of the Study 1 protocol. Evaluators answered only two questions: one Democratic-coded and one Republican-coded.

In Stage 3, evaluators completed five rounds in which they guessed the initial responses provided by targets in Study 1. In each round, evaluators were shown a question and the political affiliation of a target participant and were asked to guess the target's initial response. Each evaluator completed this task for two Democratic-coded questions, two Republican-coded questions, and one neutral question. For the partisan questions, evaluators guessed the responses of both a Democrat and a Republican. For the neutral item, the target's political affiliation was randomly assigned.

In Stage 4, evaluators guessed the updated (Stage 2) responses of the targets for the same issues that were featured in Stage 3. Each round mirrored the structure of Stage 3, with evaluators shown the question text, target affiliation, the target's initial estimate, and the directional feedback signal (i.e., whether the true value was higher or lower than the initial guess). Evaluators then guessed the target's revised estimate.

In Stages 3 and 4, the political affiliations, initial responses, and feedback signals associated with each target were randomly drawn from a pre-selected pool of Study 1 participants, selected to ensure variation in prior beliefs, received signals, and political affiliation.

Evaluators' own responses to the belief questions, as well as their guesses about the targets' estimates, were incentivized. For each evaluator, one response was randomly selected to determine their probability of earning a $1.50 bonus. As in Study 1, the bonus probability decreased with the squared distance between the guess and the true value. Comprehension checks ensured that participants understood both the incentive structure and the feedback mechanism.

At the end of the study, we collected the same demographic and attitudinal measures as in Study 1.

3.2.1.3 Study 3: Expert Survey

Study 3 was an expert survey administered via the Social Science Prediction Platform. The study focused on expert expectations regarding the extent to which Study 1 participants engaged in politically motivated reasoning.

In Part 1, experts were asked to guess the updated (Stage 2) responses of target participants from Study 1. This task consisted of five rounds. In each round, experts were shown the question text, the political affiliation of a target participant, the participant's initial estimate, and a feedback signal indicating whether the true value was higher or lower than the initial estimate. Each expert made predictions for two Democratic-coded questions, two Republican-coded questions, and one neutral question. For the partisan items, predictions were made for both Democratic and Republican targets. For the neutral item, the target's political affiliation was randomly assigned.

In Part 2, we elicited general expectations about the belief-updating behavior of Study 1 participants.

At the end of the study, we collected information on each expert's preferred political party—specifically, which party they would vote for if eligible to vote in the United States—and the strength of their partisan identification, measured on a 1–7 Likert scale.

3.2.2 Procedures

Study 1 was conducted via Prolific on March 19, 2025. The mean completion time was 17 minutes, and participants received an average total compensation of $3.36, corresponding to an average hourly rate of $15.54. The study was completed by 400 participants. 15 participants were excluded for failing to indicate a party affiliation, failing an attention check near the end of the survey, or reporting the use of external assistance (e.g., Google) to inform their responses. The final sample consists of 385 participants: 196 Republicans and 185 Democrats.

Study 2 was conducted via Prolific on March 27, 2025. The mean completion time was 24 minutes, and participants received an average total compensation of $5.28, corresponding to an average hourly rate of $17.96. The study was completed by 400 participants. 16 participants were excluded using the same criteria as in Study 1. The final sample consists of 384 evaluators: 201 Republicans and 183 Democrats.

Study 3 was conducted via the Social Science Prediction Platform between March 18 and June 30, 2025. The sample consists of 78 registered members of the platform.

3.2.3 Analysis Plan

3.2.3.1 Estimation Strategy

Our main measure of belief updating is the difference between a participant's initial and revised estimate for a given question Δ_{ij}, coded in the direction of the feedback received. For example, if the true value was higher than the initial guess, $\Delta_{ij} = Revised_{ij} - Initial_{ij}$. We standardize these measures across participants and questions.

We estimate the following regression:

$$\Delta_{ij} = \alpha_0 + \alpha_1 \cdot ProParty_{ij} + \alpha_2 \cdot AntiParty_{ij} + \alpha_3 \cdot X_{ij} + \gamma_i + \epsilon \tag{3.1}$$

where $ProParty_{ij}$ and $AntiParty_{ij}$ indicate whether the signal aligned or misaligned with the participant's political affiliation, X_{ij} includes controls such as the initial estimate and self-reported confidence, and γ_i are individual fixed effects. The omitted category is non-political signals. Our key parameter of interest is $\alpha_\Delta = \alpha_1 - \alpha_2$, which captures the difference in updating between politically aligned and misaligned signals.

In addition to this standardized measure, we also pre-registered alternative measures of belief updating (percentage-point difference, log odds ratio, and share of possible movement). These measures are formally defined in Appendix 3.A.2.1, and results using them are reported in the Results section (see Tables 3.A.3 and 3.A.15).

For Study 1, signal direction is mechanically linked to the initial estimate, meaning that more extreme priors are more likely to generate anti-party signals. As a result, treatment assignment is not randomized, and we interpret the results descriptively rather than causally.

For Studies 2 and 3 (expectations of updating), we apply the same specification with minor adjustments (e.g., replacing confidence with the expected prior of the target and omitting individual fixed effects). Because evaluators were randomly assigned to a balanced set of signals and target affiliations, expectation estimates are not subject to the same selection issue. Expected revisions are standardized across studies for comparability.

3.2.3.2 Bayesian Benchmark

In binary updating tasks, Bayesian posteriors provide a clear benchmark of symmetric updating across signal types. In our continuous setting, we approximate priors using each participant's mean estimate and reported confidence, which proxy for the distribution's mean and variance. The benchmark predicts no difference between pro- and anti-party signals ($\alpha_\Delta = 0$). Details on the benchmark logic, alternative prior distributions, and robustness checks are reported in Appendix 3.A.2.2.

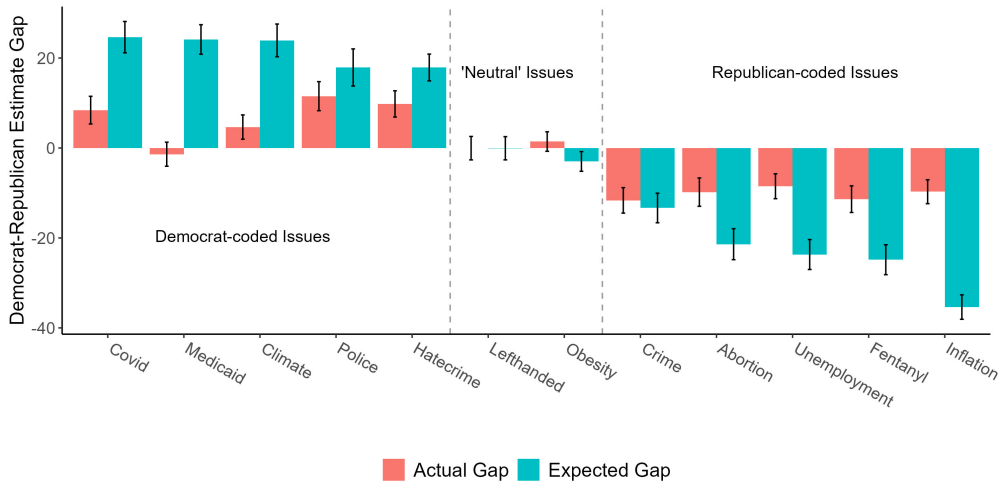

Figure 3.1. Actual and Expected Party Differences in Prior Beliefs. *The red (left) bars represent the actual difference in prior beliefs between Democrats and Republicans, i.e., the average prior belief of Democrats minus that of Republicans. The green (right) bars represent the expected difference, i.e., the average belief expected to be held by a Democrat minus that expected to be held by a Republican.*

3.2.3.3 Heterogeneity Analysis

We also examine whether updating patterns vary across individual traits (party affiliation, strength of identification, affective polarization, homophily) and issue characteristics (perceived importance, partisan gaps in priors). To do so, we interact these variables with the signal indicators in the regression:

$$\begin{aligned}\Delta_{ij} = \beta_0 + \beta_1 \cdot ProParty_{ij} + \beta_2 \cdot NonPolitical_{ij} + \beta_3 \cdot Z_{ij} \\ + \beta_4 \cdot (ProParty_{ij} \times Z_{ij}) + \beta_5 \cdot (NonPolitical_{ij} \times Z_{ij}) + \epsilon_{ij}\end{aligned} \tag{3.2}$$

where Z_{ij} denotes the heterogeneity variable of interest.

In this specification, anti-party signals serve as the omitted category. This choice allows the coefficient on the interaction term β_4 to be read directly as the effect of the moderating variable on the difference between pro- and anti-party updating.

3.3 Main Results

3.3.1 Prior Beliefs

In this section, we provide a descriptive analysis of prior beliefs—specifically, the mean guesses reported by participants in Studies 1 and 2 before any feedback was provided. We focus on the observed belief differences between Republicans and Democrats and compare them to the differences expected by Study 2 participants.

Figure 3.1 displays the actual and expected differences in prior beliefs between Democrats and Republicans for each of the 12 questions listed in Appendix 3.A.3. Democratic-coded questions appear on the left, non-political questions in the center, and Republican-coded questions on the right. The red bars indicate actual differences in prior beliefs. For the Democratic-coded questions, Democrats generally report significantly higher beliefs than Republicans, with the exception of the question on "Medicaid for all." For the Republican-coded questions, the pattern is reversed, with Republicans reporting higher prior beliefs. These patterns indicate meaningful partisan differences in prior beliefs. Table 3.A.1 shows significant partisan differences in average prior beliefs for 8 out of the 10 political questions. For the non-political questions, no sizable belief differences are observed.

Turning to expectations, the green bars in Figure 3.1 show that participants in Study 2 systematically overestimate partisan belief gaps. For nearly all political questions, the expected difference exceeds the actual difference—often by a factor of two or more. Table 3.A.1 confirms significant expected partisan differences for all 10 political questions, as well as a significant difference between expected and actual gaps for 7 of the 10 political items. For the non-political questions, actual and expected differences are closely aligned, with no notable discrepancies.

Result 3.1. Partisan differences in political prior beliefs are evident. However, we document substantial misperceptions: expected gaps between Democrats and Republicans are considerably larger than the actual belief differences.

3.3.2 Actual Belief Updating

The top panel of Figure 3.2 displays the distribution of belief updating in response to pro-party and anti-party signals. Belief updating is measured as the difference between a participant's posterior and prior belief, signed in the direction of the received signal. The figure shows that belief updating in response to pro-party and anti-party signals is broadly similar. The distributions largely overlap, and we observe no substantial difference in responsiveness based on signal alignment.

Column (1) of Table 3.1 confirms this pattern by regressing standardized belief movement on indicators for pro-party and anti-party news, with non-political items as the omitted category. We find no statistically significant difference in belief updating between pro-party and anti-party news. However, belief movement is generally larger for political than for non-political questions.[4]

4. Hartzmark, Hirshman, and Imas (2021) find that participants respond more strongly to both positive and negative news about goods they own relative to those they do not. Their preferred interpretation is that ownership increases the salience of the information. Our finding that participants respond more strongly to political than non-political signals—regardless of whether the signal is pro- or anti-party—may reflect a similar salience mechanism.

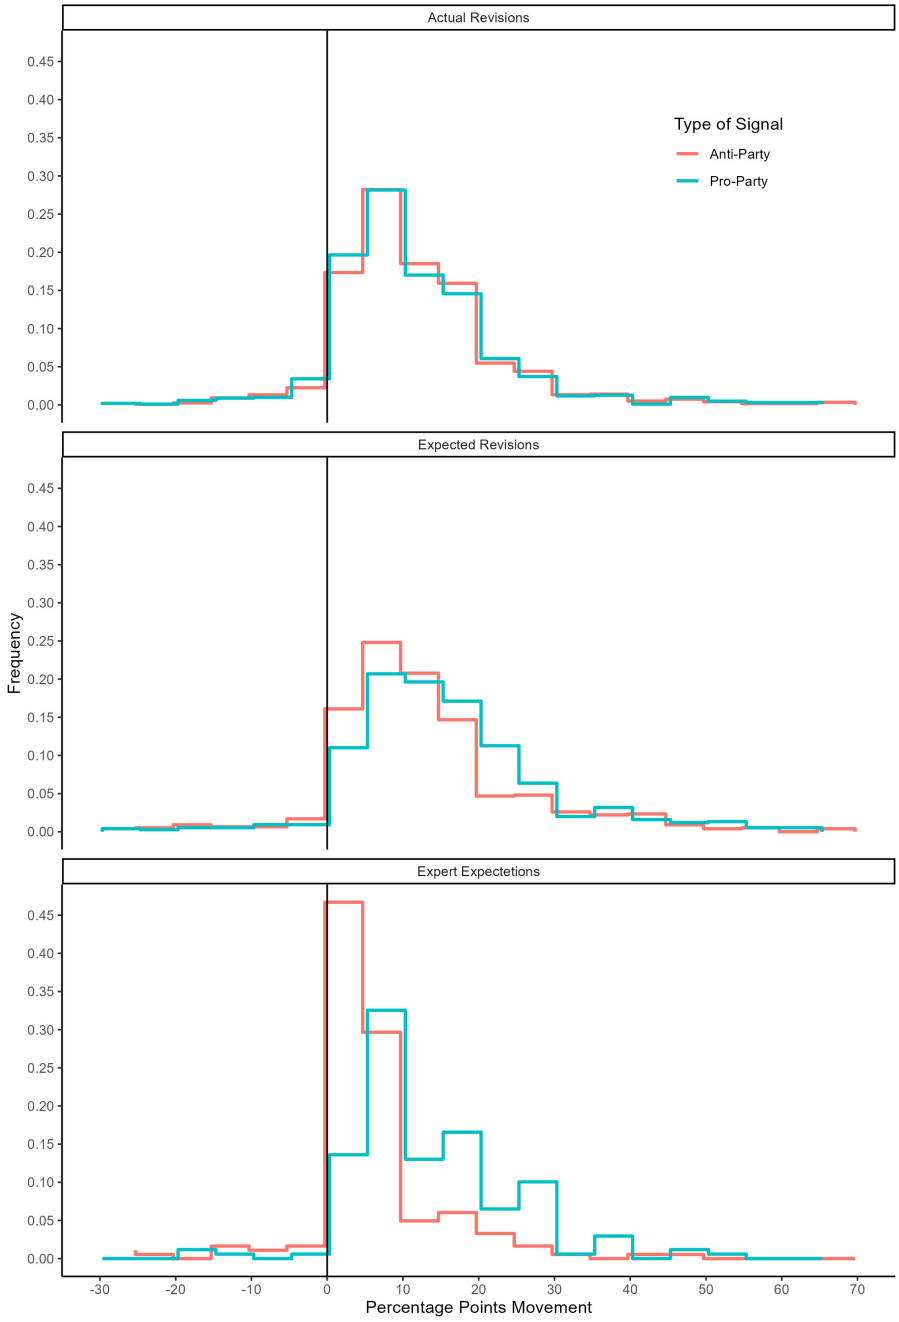

Figure 3.2. Belief Updating by Received Signal. *This figure shows the distribution of belief updating (in percentage points), defined as the difference between posterior and prior beliefs, signed in the direction of the received signal. The top panel shows actual updates (Study 1), the middle panel shows expectations of partisan participants (Study 2), and the bottom panel shows expert expectations (Study 3).*

Table 3.1. Regressions for Standardized Actual Belief Movement

Dependent variable:	Standardized Belief Movement		
	(1)	(2)	(3)
Constant	-0.094**	0.510***	
	(0.045)	(0.078)	
Anti-Party	0.132***	0.179***	0.158***
	(0.048)	(0.049)	(0.050)
Pro-Party	0.082*	0.205***	0.214***
	(0.047)	(0.048)	(0.050)
Prior		-0.010***	-0.009***
		(0.001)	(0.001)
Confidence		-0.065***	-0.087***
		(0.015)	(0.020)
Pro/Anti-Party Difference	-0.050	0.026	0.055
	(0.045)	(0.044)	(0.045)
Fixed Effects	None	None	Individual
Observations	2,612	2,612	2,612
R^2	0.002	0.065	0.533

The table reports regressions of standardized belief movement on various predictors. The sample includes belief updates in response to pro-party, anti-party, and neutral news. *Pro-Party* and *Anti-Party* are dummy variables equal to 1 if the belief update followed pro-party or anti-party news, respectively. *Confidence* measures participants' self-reported confidence in their initial estimate on a Likert scale from 1 to 7. *Prior* refers to participants' initial guess before receiving information. Standard errors are clustered at the individual level. Stars indicate significance levels: * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Column (2) adds controls for participants' prior beliefs and their reported confidence in those beliefs. Both controls are associated with lower belief updating, as expected: participants with more extreme priors and higher confidence in their initial estimate update less in response to new information. With these controls included, the difference in belief movement between political and non-political signals becomes more pronounced, while the difference between pro-party and anti-party signals remains statistically insignificant.

Because Study 1 does not involve random assignment of signal direction, the results should not be interpreted causally. In particular, participants who receive pro-party versus anti-party signals may differ systematically in the intensity of their political beliefs. For instance, a politically extreme respondent may report a high prior on a party-aligned question, increasing the likelihood of receiving an anti-party signal (i.e., learning that the true value is lower than their estimate).

To partially account for this selection in signal direction, column (3) of Table 3.1 includes individual fixed effects. Even with these controls, we continue to find no statistically significant difference in belief updating between pro-party and anti-party signals. Table 3.A.2 presents additional specifications with question fixed effects, which confirm the same pattern.

Table 3.A.3 reports alternative versions of Table 3.1, using different preregistered measures of belief movement. Across all specifications—except the one based on the share of possible belief movement—we continue to find no significant difference in belief updating between pro-party and anti-party signals.

Result 3.2. We find no evidence of motivated reasoning on average: belief updating in response to pro-party signals is not significantly greater than updating in response to anti-party signals. However, participants exhibit significantly larger belief movement in response to political signals compared to non-political ones.

3.3.3 Expected Belief Updating

3.3.3.1 Expectations by Prolific Participants

The middle panel of Figure 3.2 displays the distribution of expected belief movement as reported by evaluating participants in Study 2 (Prolific), conditional on whether the target participant received pro-party or anti-party news. The distributions indicate that evaluators expect targets to update their beliefs more in response to pro-party signals than to anti-party signals.

The first two columns of Table 3.2 provide regression evidence on evaluators' expectations about belief updating. Column (1) regresses standardized expected belief movement on indicators for pro-party and anti-party signals, with non-political signals as the omitted category. The results confirm a significant perceived asymmetry: evaluators expect stronger belief updating in response to pro-party signals than to anti-party signals. Column (2) adds a control for the actual prior belief reported

Table 3.2. Regressions for Standardized Expected Belief Movement

Dependent variable:	Standardized Expected Belief Movement			
Sample:	Prolific	Prolific	Experts	Experts
	(1)	(2)	(3)	(4)
Pro-Party	0.079	0.241***	0.133	0.314***
	(0.055)	(0.054)	(0.109)	(0.104)
Anti-Party	-0.148***	0.016	-0.411***	-0.251***
	(0.053)	(0.051)	(0.084)	(0.087)
Provided Prior		-0.019***		-0.010***
		(0.001)		(0.002)
Pro/Anti-Party Difference	0.227***	0.225***	0.544***	0.565***
	(0.052)	(0.049)	(0.091)	(0.087)
Fixed Effects	None	None	None	None
Observations	1,920	1,920	440	440
R^2	0.011	0.156	0.086	0.149

The table presents regressions of standardized expected belief movement on various variables. The first two columns use evaluating participants from Prolific, and the last two columns use experts. The sample includes expected belief updates in response to pro-party and anti-party news. Expected Belief Movement was standardized across studies to enable comparisons. *Pro-Party* and *Anti-Party* are dummy variables equal to 1 if the news is pro-party or anti-party from the perspective of the target participant. *Provided Prior* refers to the target participant's stated belief before receiving information. Standard errors are clustered at the evaluating participant level. Stars indicate significance levels: * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table 3.3. Standardized Expected Belief Movement With Expected Prior Control

Dependent variable:	Standardized Expected Belief Movement
Sample:	Prolific
	(1)
Pro-Party	0.039
	(0.051)
Anti-Party	0.135***
	(0.048)
Provided Prior	-0.022***
	(0.001)
Expected Prior	0.014***
	(0.001)
Pro/Anti-Party Difference	-0.096*
	(0.055)
Fixed Effects	None
Observations	1,920
R^2	0.242

The table presents regressions of standardized expected belief movement on various variables, using evaluating participants from Prolific. The sample includes expected belief updates in response to pro-party and anti-party news. Expected Belief Movement was standardized across studies to enable comparisons. *Pro-Party* and *Anti-Party* are dummy variables equal to 1 if the news is pro-party or anti-party from the perspective of the target participant. *Provided Prior* refers to the target participant's stated belief before receiving information. *Expected Prior* is the evaluating participant's guess of the target participant's prior belief. Standard errors are clustered at the evaluating participant level. Stars indicate significance levels: * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

by the target participant. The estimated difference in expected belief updating between pro-party and anti-party signals remains robust.

Table 3.3 adds a control for the expected prior belief—that is, the belief the evaluating participant assumes the target held prior to receiving the signal—alongside the actual prior. This adjustment accounts for systematic misperceptions about others' baseline beliefs, which we previously showed to be polarized. Once the expected prior is included, the earlier pattern reverses: participants no longer expect stronger updating in response to pro-party signals. Instead, they predict somewhat greater belief movement in response to anti-party signals.

We interpret this reversal as evidence that individuals tend to hold exaggerated assumptions about the political extremity of others. When expected belief updating

is evaluated relative to moderate actual priors, it appears biased toward pro-party signals. However, when the same updating is assessed relative to already extreme expected priors, the expected response to pro-party signals no longer appears disproportionately strong.

Table 3.A.14 reports specifications that include individual and question fixed effects. These controls do not materially affect the results. In Table 3.A.15, we present alternative versions of Table 3.2 using different pre-registered measures of expected belief updating. Across all specifications that do not control for expected priors, we consistently observe the perceived asymmetry favoring pro-party signals. Once expected priors are included, this asymmetry disappears or reverses, confirming the central role of biased assumptions about others' prior beliefs.

Result 3.3. Controlling for actual prior beliefs, evaluating participants on average expect targets to engage in motivated reasoning, as reflected in stronger expected belief updating in response to pro-party signals than to anti-party signals. However, when we additionally control for expected prior beliefs, this pattern reverses—consistent with the interpretation that evaluators hold exaggerated assumptions about the political extremity of others' beliefs.

We can also evaluate signals from the perspective of the evaluating participant by classifying news as pro-party or anti-party based on whether it aligns with the evaluator's own political affiliation. Table 3.A.16 presents this analysis. We find no evidence that evaluators are motivated in this sense: they do not systematically expect others to update more strongly in response to news that aligns with their own political views.

3.3.3.2 Expectations by Experts

The bottom panel of Figure 3.2 displays the distribution of expected belief updating by experts, conditional on whether the target received pro-party or anti-party news. The figure shows a clear asymmetry: experts expect substantially stronger updating in response to pro-party signals than to anti-party signals. Notably, the mode of the distribution for anti-party signals is close to zero, suggesting that many experts expect such information to be disregarded entirely.

Columns (3) and (4) of Table 3.2 report regressions of expected belief updating by experts. The results confirm the visual pattern observed in Figure 3.2: experts anticipate significant motivated reasoning, expecting stronger belief updating in response to pro-party signals than to anti-party signals. Table 3.A.14 shows that these conclusions are robust to the inclusion of individual and question fixed effects.

Because the standardization of expected belief updating was performed jointly for experts and Prolific participants, the coefficients are directly comparable across groups. A comparison of the estimates reveals that experts anticipate a substantially greater degree of motivated reasoning than lay participants.

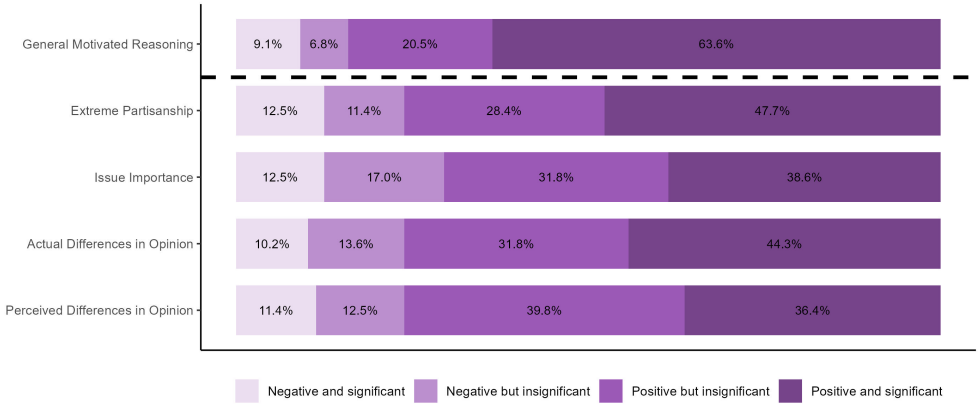

Figure 3.3. General Predictions in the Expert Study. *This figure summarizes qualitative predictions provided by participants in the expert study. "General Motivated Reasoning" refers to whether experts expected stronger belief updating in response to pro-party signals compared to anti-party signals. "Extreme Partisanship" captures whether this gap was predicted to be larger among individuals with stronger partisan identities. "Issue Importance" concerns whether the gap was expected to be larger for issues rated as more important by the target's co-partisans. "Actual Differences in Opinion" refers to whether experts anticipated larger gaps on issues with greater observed disagreement in prior beliefs. "Expected Differences in Opinion" captures whether the gap was predicted to widen with greater expected disagreement in prior beliefs.*

Figure 3.3 summarizes the additional general predictions we elicited from expert participants. Overall, 64% of experts anticipated a significantly stronger reaction to pro-party than to anti-party news. In addition, many experts predicted a larger pro–anti-party gap for people with more extreme partisan identities, for issues considered more important, and for questions where initial beliefs showed greater actual or expected differences. While the data provide (borderline) significant evidence in line with expert predictions for issue importance and partisan extremity, we find no significant heterogeneity along the dimensions of actual or expected differences in priors (see Section 3.4).

Taken together, the results suggest that experts systematically overestimated the extent of politically motivated reasoning observed in our data.

Result 3.4. Experts expect belief movement to reflect strong motivated reasoning, anticipating significantly larger updates in response to pro-party than anti-party signals. These expectations are stronger than those of Prolific participants.

3.4 Heterogeneity

3.4.1 Actual Belief Updating

3.4.1.1 Heterogeneity by Individual Characteristics

Figure 3.4 examines how the difference in belief updating between pro- and anti-party signals varies with individual-level characteristics. All specifications include individual fixed effects.

We first split the sample at the median of participants' self-reported political identification strength, measured by their response to the question: "On a scale from 1 (not important) to 7 (very important), how important is your political identity to you?" We find no statistically significant difference in the pro- versus anti-party updating gap across the two groups.

Next, we examine heterogeneity by affective polarization, using feeling thermometer ratings of Republicans and Democrats. Among participants with high affective polarization, we observe significantly greater belief updating in response to pro-party signals than to anti-party signals. Moreover, the pro/anti-party gap is significantly larger in this group than among participants with low affective polarization.

In the third analysis, we consider preferences for within-party social interaction. Participants were asked to rank potential follow-up surveys involving interactions with either co-partisans or members of the opposing party. We classify participants as *homophilic* if they expressed a preference for interacting with co-partisans in both political and non-political follow-up scenarios. Within this group, we find a significantly larger response to pro-party than anti-party signals. The size of this gap is also significantly greater than in the non-homophilic group.

Finally, we examine heterogeneity by political affiliation. This analysis reveals that the aggregate null finding masks substantial subgroup differences. Among Democratic participants, we observe significantly stronger belief updating in response to pro-party signals compared to anti-party signals. In contrast, Republican participants exhibit the opposite pattern, updating more in response to anti-party signals.

Figure 3.A.1 presents the same heterogeneity analyses without individual fixed effects. The results are broadly consistent, with one exception: the significant pro/anti-party difference for the homophilic group is no longer observed. Figure 3.A.2 further disaggregates the first three heterogeneity dimensions by party affiliation, highlighting the central role of Democratic participants in driving the overall patterns.

Result 3.5. The aggregate null finding on motivated reasoning masks meaningful heterogeneity. Controlling for prior beliefs and confidence levels, we find that Democratic participants update more in response to pro-party signals, while Republican

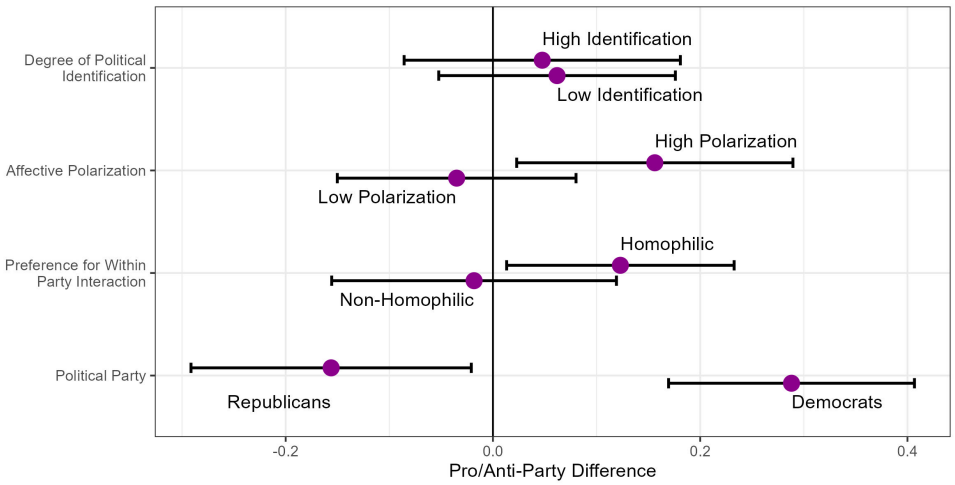

Figure 3.4. Heterogeneity of Actual Belief Movement by Target Participant Characteristics. *This figure displays the coefficient difference in belief responses to pro- and anti-party news for different subgroups of the sample. Degree of Political Identification performs a median split based on participants' indicated importance of their political identitiy. Affective Polarization performs a median split based on participants' feelings thermometer difference for their preferred vs the opposite party. Preference for Within Party Interaction is Homophilic if participants indicate they prefer to discuss both political and neutral issues with a political peer rather than a supporter of the opposite party. Political Party refers to the political affiliation of the participant. The figure is based on regression results including individual fixed effects displayed in Table 3.A.11.*

participants update more in response to anti-party signals. There is also suggestive evidence that individuals with stronger partisan identities may respond more to pro-party signals, though this pattern is not consistently significant.

3.4.1.2 Heterogeneity by Issue Characteristics

Figure 3.5 examines how question-level characteristics relate to the difference in belief updating between pro- and anti-party signals.

First, we find that a one standard deviation increase in issue importance is associated with a significantly larger pro/anti-party difference in belief updating.[5] This effect is statistically significant when individual fixed effects are excluded and remains marginally significant when fixed effects are included.

Next, we explore whether questions associated with greater polarization predict stronger motivated reasoning. We consider three measures: the actual gap in prior beliefs between Democrats and Republicans, the expected gap as perceived by Study 2 participants, and the difference between the two. None of these measures significantly predict the observed difference in belief updating between pro- and anti-party signals.

Result 3.6. More important issues are associated with a larger pro-/anti-party difference in belief updating. In contrast, neither actual nor expected polarization in prior beliefs significantly moderates this differential response.

3.4.2 Expected Belief Updating

Recall that in the analysis of actual belief updating, we found that Democrats respond more strongly to pro-party signals than to anti-party signals, while Republicans exhibit the opposite pattern. The first analysis in Figure 3.6 examines whether participants in Study 2 correctly anticipate this partisan asymmetry.

When controlling only for the actual prior belief held by the target, evaluators expect a significant pro-party bias in belief updating for Democrats, but no such asymmetry for Republicans. Once we additionally control for the expected prior—that is, the prior belief evaluators assume the target held—these asymmetries disappear (see Table 3.A.18).

This pattern mirrors the main analysis and supports the interpretation that perceived motivated reasoning among Democrats arises when evaluated against moderate actual priors, but not when assessed relative to evaluators' already polarized expectations.

In the second analysis, we examine heterogeneity by the political affiliation of the evaluating participant. While both Democratic and Republican evaluators expect

5. Figure 3.A.3 plots normalized importance ratings of political issues by party. Republican-coded issues tend to be rated as more important by Republicans, and Democratic-coded issues by Democrats, validating our classification.

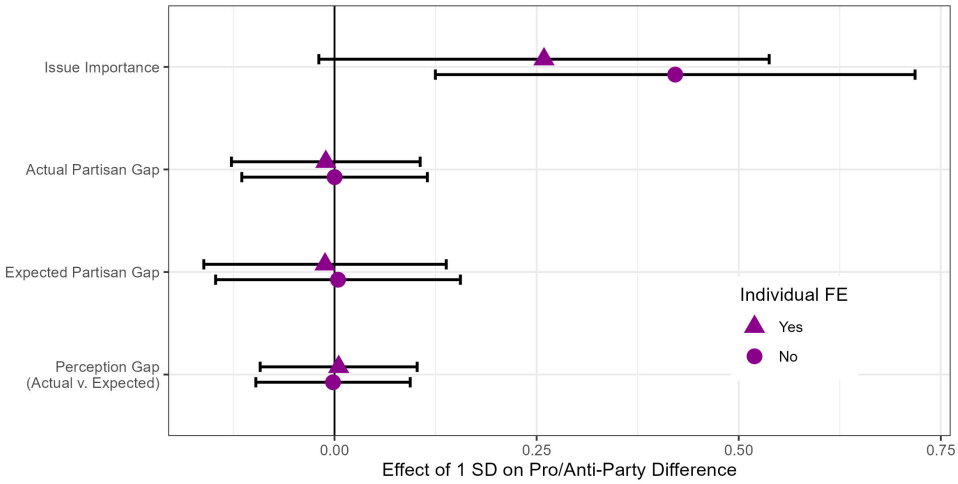

Figure 3.5. Heterogeneity of Actual Belief Movement by Issue Characteristics. *This figure displays how a one standard deviation increase in different variables affects the coefficient difference in belief responses to pro- and anti-party news. Issue importance is the standardized average importance assigned to an issue by study 1 participants who share the same political affiliation as the current participant. Actual Partisan Gap is the standardized absolute difference in prior beliefs between Democrats and Republicans for the current question, based on participants' self-reported priors. Expected Partisan Gap is the standardized absolute difference in prior beliefs between Democrats and Republicans as anticipated by the participant for the current question. Perception Gap (Actual v. Expected) is the standardized absolute difference between the actual and expected partisan prior differences for the current question. The figure is based on regression results presented in Tables 3.A.4 and 3.A.5.*

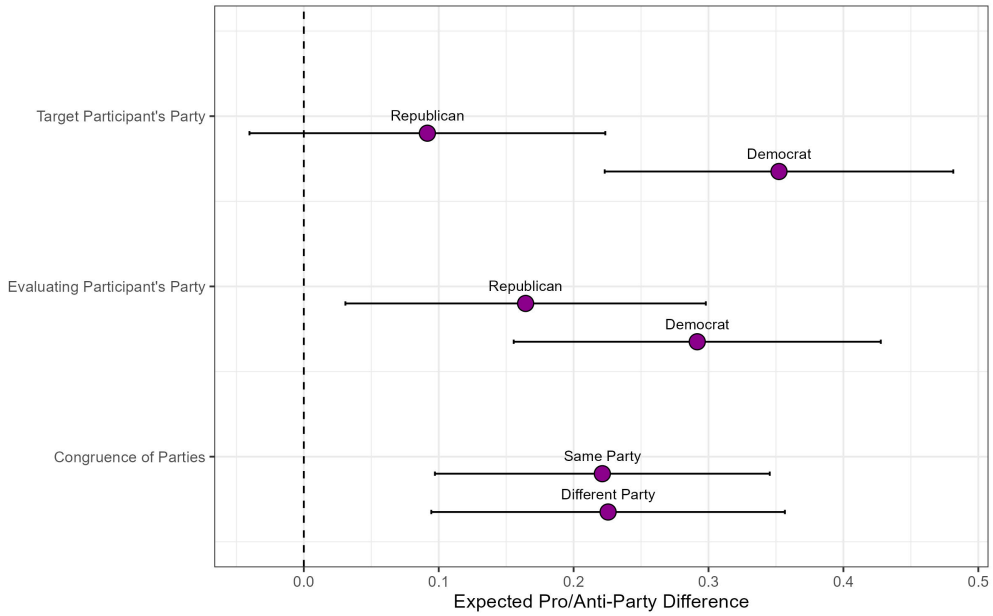

Figure 3.6. Heterogeneity of Expected Belief Updating by Political Affiliation. *This figure displays the coefficient difference in expected belief responses to pro- and anti-party news for different subgroups of the sample. Target Participant's Party splits the sample by whether the target individual affiliates with the Democratic or Republican party. Evaluating Participant's Party instead splits the sample by the political affiliation of the evaluating individual. Congruence of Parties splits the sample by whether the evaluating and target individual share the same political affiliation or not. The figure is based on regressions presented in Table 3.A.17, which feature a control for the actual prior belief of targets but not the prior belief that was expected by evaluators. Table 3.A.18 presents regressions that additionally control for the expected prior belief.*

targets to engage in motivated reasoning, the expected asymmetry is directionally stronger among Democratic evaluators. However, once we control for expected priors, these differences are no longer statistically significant.

We also investigate whether evaluators' expectations depend on whether they share the same political affiliation as the target. In the final analysis in Figure 3.6, we split the sample based on partisan alignment between evaluator and target. We find no evidence that shared political affiliation influences expected belief updating.

Tables 3.A.19 and 3.A.20 explore heterogeneity along other dimensions analyzed for actual belief updating. None of these yield additional statistically significant effects.

Result 3.7. Participants expect a greater degree of motivated reasoning among Democrats than among Republicans, consistent with observed patterns in actual belief updating. The political affiliation of the evaluating participant, as well as whether the evaluator and target share the same affiliation, does not significantly moderate these expectations.

3.5 Conclusion

This chapter investigates the extent to which politically motivated reasoning shapes belief updating in response to factual information. Using a series of incentivized experiments, we assess both actual belief updating behavior and expectations about how others process political information.

Our main finding is that, on average, participants do not update their beliefs more strongly in response to pro-party signals than to anti-party signals, suggesting limited evidence of politically motivated reasoning in belief revision. However, this aggregate pattern conceals important heterogeneity. We find that Democrats tend to update more in response to pro-party signals, while Republicans update more in response to anti-party signals. In addition, participants with higher levels of affective polarization or a preference for within-group interaction exhibit stronger partisan asymmetries in updating.

In contrast to actual behavior, both Prolific participants and experts systematically overestimate the extent of politically motivated reasoning in others. Participants expect others—especially Democrats—to engage in belief updating biased toward their own party. Experts anticipate even greater asymmetries than laypeople. However, once we control for participants' biased expectations about others' prior beliefs, these perceived asymmetries largely disappear.

Together, our findings highlight a divergence between actual and perceived political bias in information processing. While politically motivated reasoning may be more limited than commonly assumed, people tend to overestimate its prevalence in others. This misperception may itself have important social and political consequences, such as reinforcing distrust or fueling polarization.

Future research could explore the implications of such distorted second-order beliefs. In particular, overestimating the extent of motivated reasoning in others may contribute to reduced willingness to engage in political discussion, especially with those holding opposing views. If individuals perceive others as systematically biased or unresponsive to evidence, they may view political discourse as unproductive or even harmful. Understanding how these perceptions shape political behavior and social interaction is an important direction for future work.

Appendix 3.A Appendix

3.A.1 Additional Tables and Figures

3.A.1.1 Prior Beliefs

Table 3.A.1. Summary of Actual and Expected Priors by Party

	Actual Priors			Expected Priors			Actual - Expected
Question	Democrat	Republican	Difference	Democrat	Republican	Difference	Diff-in-Diff
Covid	62.76	54.35	8.42***	70.12	45.47	24.65***	-16.23***
	(2.03)	(2.31)	(3.07)	(1.94)	(2.9)	(3.48)	(4.65)
Medicaid	48.69	50.08	-1.38	59.82	35.7	24.13***	-25.51***
	(2.01)	(1.76)	(2.67)	(2.41)	(2.23)	(3.28)	(4.23)
Climate	80.21	75.56	4.65*	78.16	54.24	23.92***	-19.27***
	(1.74)	(2.07)	(2.7)	(2.16)	(2.92)	(3.64)	(4.53)
Police	41.92	30.39	11.53***	49.8	31.89	17.91***	-6.38
	(2.49)	(2.06)	(3.23)	(3.12)	(2.7)	(4.12)	(5.24)
Hatecrime	60.67	50.87	9.8***	63.49	45.59	17.9***	-8.1*
	(2.25)	(1.86)	(2.92)	(2.1)	(2.11)	(2.98)	(4.17)
Lefthanded	28.19	28.23	-0.04	29.95	30.01	-0.06	0.02
	(1.87)	(1.81)	(2.61)	(1.64)	(1.99)	(2.58)	(3.67)
Obesity	45.49	44.06	1.44	44.43	47.42	-2.99	4.43
	(1.67)	(1.39)	(2.17)	(1.34)	(1.73)	(2.19)	(3.08)
Crime	25.79	37.44	-11.65***	32.34	45.68	-13.34***	1.69
	(2.03)	(1.97)	(2.83)	(2.15)	(2.49)	(3.28)	(4.34)
Abortion	30.62	40.43	-9.81***	30.65	52.04	-21.39***	11.58**
	(2.11)	(2.34)	(3.15)	(2.41)	(2.46)	(3.44)	(4.67)
Unemployment	40.22	48.73	-8.51***	34.29	57.97	-23.69***	15.18***
	(2.05)	(1.86)	(2.76)	(2.23)	(2.47)	(3.32)	(4.32)
Fentanyl	26.13	37.51	-11.38***	31.45	56.29	-24.84***	13.46***
	(2.02)	(2.16)	(2.96)	(2.22)	(2.48)	(3.33)	(4.45)
Inflation	37.45	47.17	-9.72***	33.19	68.56	-35.36***	25.64***
	(1.86)	(1.9)	(2.66)	(1.99)	(1.85)	(2.72)	(3.8)

This table reports question-level averages of participants' priors (in percent) for both Democrats and Republicans. "Actual Prior" refers to participants' own prior beliefs based on their political affiliation. "Expected Prior" refers to what participants expected members of the party to believe. "Difference" is the difference in means between Democrats and Republicans. "Diff-in-Diff" is the difference between actual and expected differences. Standard errors are shown in parentheses. *, **, and *** denote significance at the 10%, 5%, and 1% level, respectively.

3.A.1.2 Actual Belief Updating

Table 3.A.2. Question Fixed Effects Specifications for Standardized Actual Belief Movement

Dependent variable:	Standardized Belief Movement		
	(1)	(2)	(3)
Pro-Party	-0.049	0.019	0.041
	(0.045)	(0.044)	(0.045)
Prior		-0.010***	-0.008***
		(0.001)	(0.001)
Confidence		-0.063***	-0.081***
		(0.015)	(0.022)
Fixed Effects	Question	Question	Individual + Q
Observations	2,612	2,612	2,612
R^2	0.032	0.077	0.542

The table reports regressions of standardized belief movement on various predictors. The sample includes belief updates in response to pro-party, anti-party, and neutral news. *Pro-Party* and *Non-Political* are dummy variables equal to 1 if the belief update followed pro-party or neutral news, respectively. *Confidence* measures participants' self-reported confidence in their initial estimate on a Likert scale from 1 to 7. *Prior* refers to participants' initial guess before receiving information. Standard errors are clustered at the individual level. Stars indicate significance levels: * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table 3.A.3. Regressions for Alternative Measures of Actual Belief Movement

Dependent variable:	Percentage Point Belief Movement			Belief Movement Likelihood Ratio			Share of Possible Belief Movement		
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
Constant	12.117***	18.856***		0.599***	0.920***		0.181***	0.164***	
	(0.498)	(0.871)		(0.026)	(0.046)		(0.008)	(0.017)	
Anti-Party	1.480***	1.999***	1.767***	0.165***	0.201***	0.190***	0.058***	0.042***	0.034***
	(0.540)	(0.548)	(0.554)	(0.028)	(0.027)	(0.028)	(0.010)	(0.009)	(0.009)
Pro-Party	0.919*	2.287***	2.387***	0.123***	0.209***	0.220***	0.089***	0.056***	0.061***
	(0.521)	(0.535)	(0.561)	(0.028)	(0.028)	(0.030)	(0.008)	(0.008)	(0.009)
Prior		-0.111***	-0.102***		-0.007***	-0.006***		0.003***	0.003***
		(0.010)	(0.011)		(0.001)	(0.001)		(0.000)	(0.000)
Confidence		-0.727***	-0.973***		-0.022**	0.006		-0.016***	-0.009**
		(0.166)	(0.228)		(0.009)	(0.014)		(0.003)	(0.004)
Pro/Anti-Party Difference	-0.560	0.288	0.620	-0.042	0.008	0.030	0.031***	0.015	0.027***
	(0.499)	(0.494)	(0.500)	(0.027)	(0.027)	(0.027)	(0.010)	(0.010)	(0.010)
Fixed Effects	None	None	Individual	None	None	Individual	None	None	Individual
Observations	2,612	2,612	2,612	2,600	2,600	2,600	2,612	2,612	2,612
R^2	0.002	0.065	0.533	0.009	0.084	0.547	0.018	0.114	0.536

The table presents regressions for three dependent variables: (1) percentage point belief movement, (2) the belief movement likelihood ratio (computed as the log likelihood ratio of posterior vs. prior), and (3) the actual share of possible belief movement. The sample includes belief updates in response to pro-party, anti-party, and neutral news. *Pro-Party* and *Anti-Party* are dummy variables equal to 1 if the belief update followed pro-party or anti-party news, respectively. *Prior* refers to participants' stated belief before receiving information. *Confidence* measures participants' self-reported confidence in their initial estimate on a Likert scale from 1 to 7. Standard errors are clustered at the individual level. Stars indicate significance levels: * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

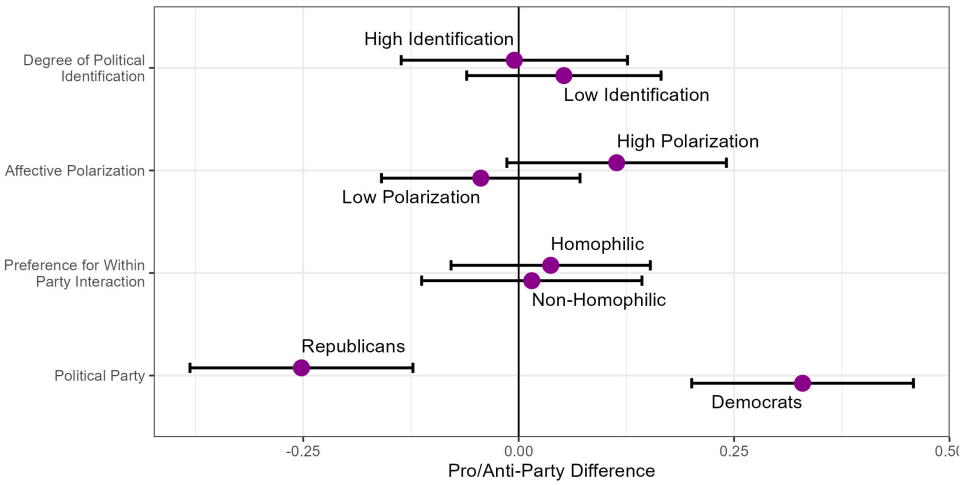

Figure 3.A.1. Heterogeneity of Actual Belief Movement by Individual Characteristics. *This figure displays the coefficient difference in belief responses to pro- and anti-party news for different subgroups of the sample. Degree of Political Identification performs a median split based on participants' indicated importance of their political identitiy. Affective Polarization performs a median split based on participants' feelings thermometer difference for their preferred vs the opposite party. Preference for Within Party Interaction is Homophilic if participants indicate they prefer to discuss both political and neutral issues with a political peer rather than a supporter of the opposite party. Political Party refers to the political affiliation of the participant. The figure is based on regression results with no fixed effects displayed in Table 3.A.10.*

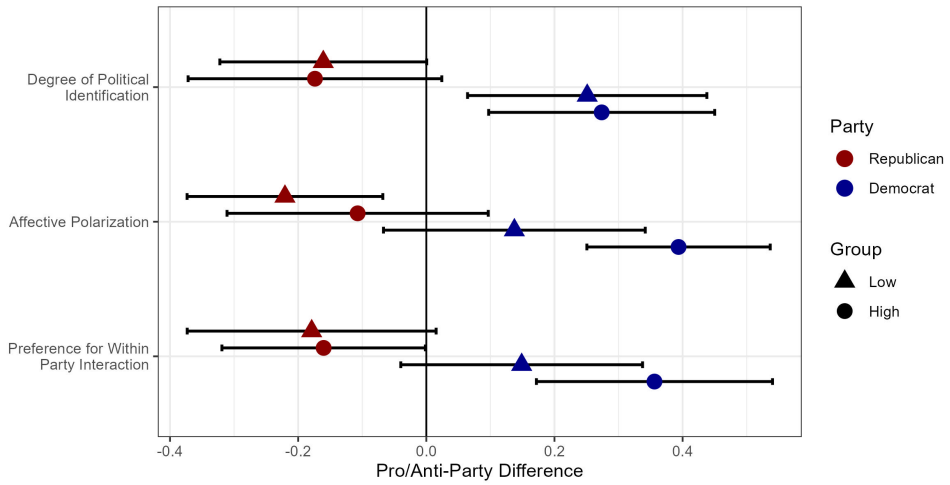

Figure 3.A.2. Heterogeneity of Actual Belief Movement by Individual Characteristics and Political Affiliation. *This figure displays the coefficient difference in belief responses to pro- and anti-party news for different subgroups of the sample. Degree of Political Identification performs a median split based on participants' indicated importance of their political identitiy. Affective Polarization performs a median split based on participants' feelings thermometer difference for their preferred vs the opposite party. Preference for Within Party Interaction is Homophilic if participants indicate they prefer to discuss both political and neutral issues with a political peer rather than a supporter of the opposite party. Party refers to the political affiliation of the participant. The figure is based on regression results with individual fixed effects displayed in Tables 3.A.12 and 3.A.13.*

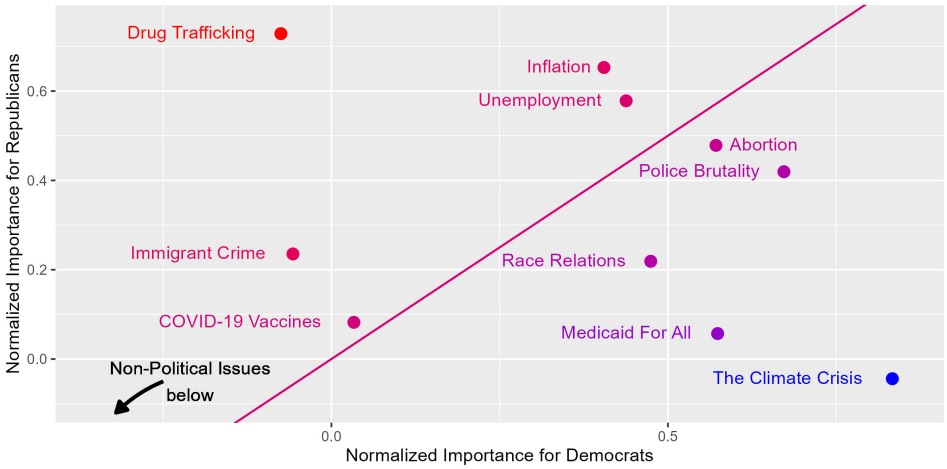

Figure 3.A.3. Normalized Importance of Political Issues for Republicans and Democrats. *Normalized importance ratings are plotted for Democrats on the x-axis and for Republicans on the y-axis. The line denotes the 45 degree line, meaning that issues on these line would be regarded as equally important by both Democrats and Republicans. Regarding the non-political issues, Left-handedness has a normalized importance of* −2.68 *for Democrats and* −3.00 *for Republicans. Obesity has a normalized importance of* −1.19 *for Democrats and* −0.40 *for Republicans.*

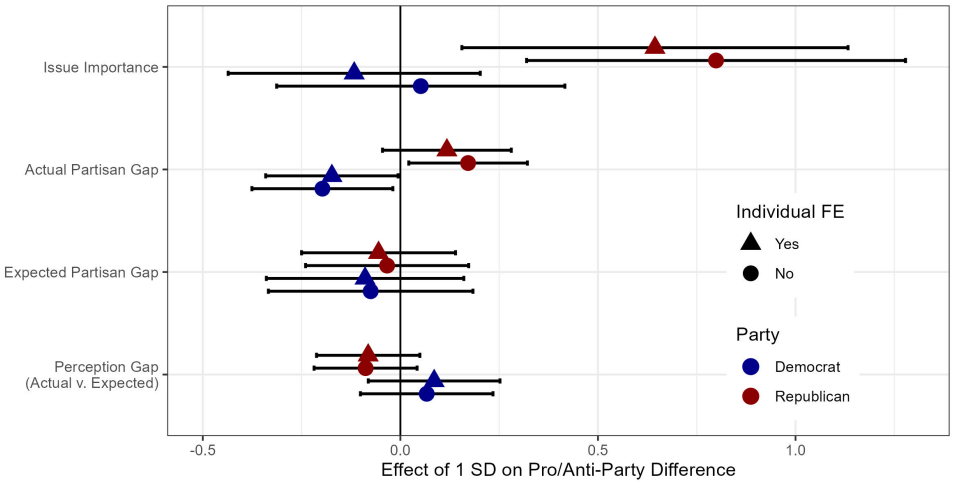

Figure 3.A.4. Heterogeneity of Actual Belief Movement by Issue Characteristics and Political Affiliation. *This figure displays how a one standard deviation increase in different variables affects the coefficient difference in belief responses to pro- and anti-party news. Issue importance is the standardized average importance assigned to an issue by study 1 participants who share the same political affiliation as the current participant. Actual Partisan Gap is the standardized absolute difference in prior beliefs between Democrats and Republicans for the current question, based on participants' self-reported priors. Expected Partisan Gap is the standardized absolute difference in prior beliefs between Democrats and Republicans as expected by the participant for the current question. Perception Gap (Actual v. Expected) is the standardized absolute difference between the actual and expected partisan prior differences for the current question. The figure is based on regression results presented in Tables 3.A.6, 3.A.7, 3.A.8 and 3.A.9.*

Table 3.A.4. Regressions for Issue Heterogeneity of Actual Belief Movement

Dependent variable:	Standardized Belief Movement			
	(1)	(2)	(3)	(4)
Constant	0.719***	0.702***	0.654***	0.658***
	(0.082)	(0.078)	(0.079)	(0.076)
Pro-Party	-0.119	0.020	0.025	0.023
	(0.073)	(0.050)	(0.054)	(0.046)
Non-Political	-0.177*	-0.020	0.093	-0.073
	(0.099)	(0.418)	(0.531)	(0.200)
Confidence	-0.063***	-0.066***	-0.065***	-0.065***
	(0.015)	(0.015)	(0.015)	(0.015)
Prior	-0.010***	-0.010***	-0.010***	-0.010***
	(0.001)	(0.001)	(0.001)	(0.001)
Importance	-0.068			
	(0.103)			
Importance × Pro-Party	0.421***			
	(0.151)			
Importance × Non-Political	0.085			
	(0.109)			
Actual Prior Difference		-0.067		
		(0.041)		
Act. Prior Diff. × Pro-Party		0.000		
		(0.059)		
Act. Prior Diff. × Non-Political		0.183		
		(0.277)		
Expected Prior Difference			0.082*	
			(0.046)	
Exp. Prior Diff × Pro-Party			0.004	
			(0.077)	
Exp. Prior Diff × Non-Political			0.052	
			(0.305)	
Perception Gap (Actual v. Expected)				0.074**
				(0.032)
Perception Gap × Pro-Party				-0.002
				(0.049)
Perception Gap × Non-Political				-0.001
				(0.178)
Fixed Effects	None	None	None	None
Observations	2,612	2,612	2,612	2,612
R^2	0.069	0.067	0.067	0.068

The table presents regressions of standardized belief movement on various predictors. The sample includes belief updates in response to pro-party, anti-party, and neutral news. *Pro-Party* is a dummy for whether a belief update was in response to pro-party news. *Non-Political* is a dummy for whether a belief update was in response to neutral news, i.e. news for a neutral topic. *Confidence* measures participantsśelf-reported confidence in their initial estimate on a Likert scale from 1 to 7. *Prior* refers to participantsínitial guess before receiving information. *Importance* refers to the standardized average importance that Study 1 participants with the same political affiliation as the current participant assigned to the question. *Actual Prior Difference* is the standardized absolute actual difference in prior beliefs between Democrats and Republicans for the current question. *Expected Prior Difference* is the standardized absolute expected difference in prior beliefs for Democrats and Republicans for the current question. *Perception Gap (Actual v. Expected)* is the standardized absolute difference of the actual and expected difference in prior beliefs for Democrats and Republicans for the current question. Standard errors are clustered at the individual level. Stars indicate significance levels: * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table 3.A.5. Regressions for Issue Heterogeneity of Actual Belief Movement (with Individual Fixed Effects)

Dependent variable:	Standardized Belief Movement			
	(1)	(2)	(3)	(4)
Pro-Party	-0.023	0.046	0.060	0.047
	(0.070)	(0.052)	(0.054)	(0.045)
Non-Political	-0.109	0.010	0.144	-0.043
	(0.098)	(0.405)	(0.515)	(0.194)
Confidence	-0.091***	-0.094***	-0.087***	-0.088***
	(0.021)	(0.021)	(0.021)	(0.021)
Prior	-0.010***	-0.008***	-0.009***	-0.009***
	(0.001)	(0.001)	(0.001)	(0.001)
Importance	0.068			
	(0.102)			
Importance × Pro-Party	0.259*			
	(0.142)			
Importance × Non-Political	-0.053			
	(0.107)			
Actual Prior Difference		-0.101**		
		(0.042)		
Act. Prior Diff. × Pro-Party		-0.011		
		(0.060)		
Act. Prior Diff. × Non-Political		0.231		
		(0.269)		
Expected Prior Difference			0.086*	
			(0.049)	
Exp. Prior Diff × Pro-Party			-0.012	
			(0.077)	
Exp. Prior Diff × Non-Political			0.065	
			(0.300)	
Perception Gap (Actual v. Expected)				0.093***
				(0.033)
Perception Gap × Pro-Party				0.005
				(0.050)
Perception Gap × Non-Political				-0.013
				(0.175)
Fixed Effects	Individual	Individual	Individual	Individual
Observations	2,612	2,612	2,612	2,612
R^2	0.535	0.536	0.534	0.536

The table presents regressions of standardized belief movement on various predictors. The sample includes belief updates in response to pro-party, anti-party, and neutral news. *Pro-Party* is a dummy for whether a belief update was in response to pro-party news. *Non-Political* is a dummy for whether a belief update was in response to neutral news, i.e. news for a neutral topic. *Confidence* measures participants self-reported confidence in their initial estimate on a Likert scale from 1 to 7. *Prior* refers to participants initial guess before receiving information. *Importance* refers to the standardized average importance that Study 1 participants with the same political affiliation as the current participant assigned to the question. *Actual Prior Difference* is the standardized absolute actual difference in prior beliefs between Democrats and Republicans for the current question. *Expected Prior Difference* is the standardized absolute expected difference in prior beliefs for Democrats and Republicans for the current question. *Perception Gap (Actual v. Expected)* is the standardized absolute difference of the actual and expected difference in prior beliefs for Democrats and Republicans for the current question. Standard errors are clustered at the individual level. Stars indicate significance levels: * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table 3.A.6. Issue Heterogeneity among Democrats (No Individual FE)

Dependent variable:	Standardized Belief Movement			
	(1)	(2)	(3)	(4)
Constant	0.515***	0.554***	0.571***	0.565***
	(0.093)	(0.095)	(0.097)	(0.092)
Pro-Party	0.279***	0.324***	0.315***	0.251***
	(0.107)	(0.088)	(0.112)	(0.094)
Non-Political	0.146	0.114	0.184	0.066
	(0.166)	(0.606)	(0.772)	(0.283)
Confidence	-0.081***	-0.073***	-0.074***	-0.074***
	(0.021)	(0.021)	(0.021)	(0.021)
Prior	-0.010***	-0.009***	-0.010***	-0.010***
	(0.002)	(0.002)	(0.002)	(0.002)
Importance	0.238*			
	(0.129)			
Importance × Pro-Party	0.052			
	(0.186)			
Importance × Non-Political	-0.215			
	(0.156)			
Actual Prior Difference		0.013		
		(0.057)		
Act. Prior Diff. × Pro-Party		-0.197**		
		(0.091)		
Act. Prior Diff. × Non-Political		0.057		
		(0.406)		
Expected Prior Difference			0.020	
			(0.060)	
Exp. Prior Diff × Pro-Party			-0.075	
			(0.132)	
Exp. Prior Diff × Non-Political			0.076	
			(0.455)	
Perception Gap (Actual v. Expected)				0.017
				(0.041)
Perception Gap × Pro-Party				0.067
				(0.086)
Perception Gap × Non-Political				0.033
				(0.266)
Fixed Effects	None	None	None	None
Observations	1,282	1,282	1,282	1,282
R^2	0.067	0.068	0.062	0.063

The table presents regressions of standardized belief movement on various predictors. The sample includes belief updates by Democrats in response to pro-party, anti-party, and neutral news. *Pro-Party* is a dummy for whether a belief update was in response to pro-party news. *Non-Political* is a dummy for whether a belief update was in response to neutral news, i.e. news for a neutral topic. *Confidence* measures participants self-reported confidence in their initial estimate on a Likert scale from 1 to 7. *Prior* refers to participants initial guess before receiving information. *Importance* refers to the standardized average importance that Study 1 participants with the same political affiliation as the current participant assigned to the question. *Actual Prior Difference* is the standardized absolute actual difference in prior beliefs between Democrats and Republicans for the current question. *Expected Prior Difference* is the standardized absolute expected difference in prior beliefs for Democrats and Republicans for the current question. *Perception Gap (Actual v. Expected)* is the standardized absolute difference of the actual and expected difference in prior beliefs for Democrats and Republicans for the current question. Standard errors are clustered at the individual level. Stars indicate significance levels: * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table 3.A.7. Issue Heterogeneity among Republicans (No Individual FE)

Dependent variable:	Standardized Belief Movement			
	(1)	(2)	(3)	(4)
Constant	1.177***	1.105***	1.021***	1.021***
	(0.149)	(0.138)	(0.144)	(0.140)
Pro-Party	-0.530***	-0.311***	-0.233***	-0.226***
	(0.116)	(0.071)	(0.076)	(0.064)
Non-Political	-0.457***	0.014	0.240	-0.160
	(0.125)	(0.582)	(0.742)	(0.283)
Confidence	-0.069***	-0.072***	-0.075***	-0.073***
	(0.022)	(0.022)	(0.023)	(0.023)
Prior	-0.015***	-0.014***	-0.014***	-0.013***
	(0.001)	(0.001)	(0.001)	(0.001)
Importance	-0.231			
	(0.148)			
Importance × Pro-Party	0.799***			
	(0.245)			
Importance × Non-Political	0.274*			
	(0.151)			
Actual Prior Difference		-0.129**		
		(0.058)		
Act. Prior Diff. × Pro-Party		0.171**		
		(0.077)		
Act. Prior Diff. × Non-Political		0.444		
		(0.382)		
Expected Prior Difference			0.120	
			(0.075)	
Exp. Prior Diff × Pro-Party			-0.034	
			(0.105)	
Exp. Prior Diff × Non-Political			0.223	
			(0.418)	
Perception Gap (Actual v. Expected)				0.120**
				(0.052)
Perception Gap × Pro-Party				-0.088
				(0.067)
Perception Gap × Non-Political				0.074
				(0.245)
Fixed Effects	None	None	None	None
Observations	1,330	1,330	1,330	1,330
R^2	0.115	0.110	0.108	0.110

The table presents regressions of standardized belief movement on various predictors. The sample includes belief updates by Republicans in response to pro-party, anti-party, and neutral news. *Pro-Party* is a dummy for whether a belief update was in response to pro-party news. *Non-Political* is a dummy for whether a belief update was in response to neutral news, i.e. news for a neutral topic. *Confidence* measures participants self-reported confidence in their initial estimate on a Likert scale from 1 to 7. *Prior* refers to participants initial guess before receiving information. *Importance* refers to the standardized average importance that Study 1 participants with the same political affiliation as the current participant assigned to the question. *Actual Prior Difference* is the standardized absolute actual difference in prior beliefs between Democrats and Republicans for the current question. *Expected Prior Difference* is the standardized absolute expected difference in prior beliefs for Democrats and Republicans for the current question. *Perception Gap (Actual v. Expected)* is the standardized absolute difference of the actual and expected difference in prior beliefs for Democrats and Republicans for the current question. Standard errors are clustered at the individual level. Stars indicate significance levels: * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table 3.A.8. Issue Heterogeneity among Democrats (With Individual FE)

Dependent variable:	Standardized Belief Movement			
	(1)	(2)	(3)	(4)
Pro-Party	0.343***	0.295***	0.299***	0.214***
	(0.092)	(0.082)	(0.103)	(0.083)
Non-Political	0.291*	0.196	0.253	0.102
	(0.167)	(0.591)	(0.757)	(0.278)
Confidence	-0.102***	-0.086***	-0.079***	-0.079***
	(0.029)	(0.028)	(0.027)	(0.028)
Prior	-0.010***	-0.008***	-0.009***	-0.009***
	(0.002)	(0.001)	(0.002)	(0.002)
Importance	0.451***			
	(0.123)			
Importance × Pro-Party	-0.117			
	(0.163)			
Importance × Non-Political	-0.398***			
	(0.144)			
Actual Prior Difference		-0.052		
		(0.059)		
Act. Prior Diff. × Pro-Party		-0.173**		
		(0.085)		
Act. Prior Diff. × Non-Political		0.187		
		(0.397)		
Expected Prior Difference			0.039	
			(0.065)	
Exp. Prior Diff × Pro-Party			-0.089	
			(0.127)	
Exp. Prior Diff × Non-Political			0.092	
			(0.447)	
Perception Gap (Actual v. Expected)				0.048
				(0.045)
Perception Gap × Pro-Party				0.085
				(0.085)
Perception Gap × Non-Political				0.029
				(0.262)
Fixed Effects	Individual	Individual	Individual	Individual
Observations	1,282	1,282	1,282	1,282
R^2	0.544	0.543	0.536	0.538

The table presents regressions of standardized belief movement on various predictors. The sample includes belief updates by Democrats in response to pro-party, anti-party, and neutral news. *Pro-Party* is a dummy for whether a belief update was in response to pro-party news. *Non-Political* is a dummy for whether a belief update was in response to neutral news, i.e. news for a neutral topic. *Confidence* measures participants self-reported confidence in their initial estimate on a Likert scale from 1 to 7. *Prior* refers to participants initial guess before receiving information. *Importance* refers to the standardized average importance that Study 1 participants with the same political affiliation as the current participant assigned to the question. *Actual Prior Difference* is the standardized absolute actual difference in prior beliefs between Democrats and Republicans for the current question. *Expected Prior Difference* is the standardized absolute expected difference in prior beliefs for Democrats and Republicans for the current question. *Perception Gap (Actual v. Expected)* is the standardized absolute difference of the actual and expected difference in prior beliefs for Democrats and Republicans for the current question. Standard errors are clustered at the individual level. Stars indicate significance levels: * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table 3.A.9. Issue Heterogeneity among Republicans (With Individual FE)

Dependent variable:	Standardized Belief Movement			
	(1)	(2)	(3)	(4)
Pro-Party	-0.391***	-0.210***	-0.132*	-0.136**
	(0.117)	(0.076)	(0.075)	(0.065)
Non-Political	-0.409***	-0.053	0.152	-0.164
	(0.126)	(0.546)	(0.693)	(0.265)
Confidence	-0.095***	-0.100***	-0.102***	-0.101***
	(0.032)	(0.032)	(0.032)	(0.032)
Prior	-0.013***	-0.012***	-0.012***	-0.011***
	(0.002)	(0.001)	(0.001)	(0.001)
Importance	-0.233			
	(0.156)			
Importance × Pro-Party	0.644***			
	(0.249)			
Importance × Non-Political	0.266*			
	(0.160)			
Actual Prior Difference		-0.134**		
		(0.060)		
Act. Prior Diff. × Pro-Party		0.118		
		(0.083)		
Act. Prior Diff. × Non-Political		0.365		
		(0.361)		
Expected Prior Difference			0.125*	
			(0.075)	
Exp. Prior Diff × Pro-Party			-0.055	
			(0.099)	
Exp. Prior Diff × Non-Political			0.132	
			(0.395)	
Perception Gap (Actual v. Expected)				0.135***
				(0.052)
Perception Gap × Pro-Party				-0.081
				(0.067)
Perception Gap × Non-Political				-0.002
				(0.231)
Fixed Effects	Participant	Participant	Participant	Participant
Observations	1,330	1,330	1,330	1,330
R^2	0.549	0.548	0.547	0.549

The table presents regressions of standardized belief movement on various predictors. The sample includes belief updates by Republicans in response to pro-party, anti-party, and neutral news. *Pro-Party* is a dummy for whether a belief update was in response to pro-party news. *Non-Political* is a dummy for whether a belief update was in response to neutral news, i.e. news for a neutral topic. *Confidence* measures participants self-reported confidence in their initial estimate on a Likert scale from 1 to 7. *Prior* refers to participants initial guess before receiving information. *Importance* refers to the standardized average importance that Study 1 participants with the same political affiliation as the current participant assigned to the question. *Actual Prior Difference* is the standardized absolute actual difference in prior beliefs between Democrats and Republicans for the current question. *Expected Prior Difference* is the standardized absolute expected difference in prior beliefs for Democrats and Republicans for the current question. *Perception Gap (Actual v. Expected)* is the standardized absolute difference of the actual and expected difference in prior beliefs for Democrats and Republicans for the current question. Standard errors are clustered at the individual level. Stars indicate significance levels: * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table 3.A.10. Individual Heterogeneity for Actual Belief Movement (No Individual FE)

Dependent variable:	Standardized Belief Movement			
	(1)	(2)	(3)	(4)
Constant	0.666***	0.701***	0.673***	0.618***
	(0.076)	(0.086)	(0.089)	(0.073)
Pro-Party	0.052	-0.044	0.015	0.329***
	(0.059)	(0.058)	(0.065)	(0.066)
Non-Political	-0.111*	-0.252***	-0.167**	0.011
	(0.066)	(0.064)	(0.075)	(0.062)
Confidence	-0.066***	-0.067***	-0.064***	-0.072***
	(0.015)	(0.015)	(0.015)	(0.015)
Prior	-0.010***	-0.010***	-0.010***	-0.012***
	(0.001)	(0.001)	(0.001)	(0.001)
High Identification	0.055			
	(0.070)			
High Identification × Pro-Party	-0.057			
	(0.087)			
High Identification × Non-Political	-0.149			
	(0.097)			
High Polarization		-0.007		
		(0.067)		
High Polarization × Pro-Party		0.158*		
		(0.088)		
High Polarization × Non-Political		0.145		
		(0.097)		
Homophily			0.020	
			(0.068)	
Homophily × Pro-Party			0.022	
			(0.087)	
Homophily × Non-Political			-0.023	
			(0.098)	
Republican				0.388***
				(0.072)
Republican × Pro-Party				-0.582***
				(0.098)
Republican × Non-Political				-0.429***
				(0.099)
Fixed Effects	None	None	None	None
Observations	2,612	2,612	2,612	2,612
R^2	0.066	0.068	0.065	0.084

The table presents regressions of standardized belief movement on various predictors. The sample includes belief updates in response to pro-party, anti-party, and neutral news. *Pro-Party* is a dummy for whether a belief update was in response to pro-party news. *Non-Political* is a dummy for whether a belief update was in response to neutral news, i.e. news for a neutral topic. *Confidence* measures participants self-reported confidence in their initial estimate on a Likert scale from 1 to 7. *Prior* refers to participants initial guess before receiving information. *High Identification* is a dummy that is 1 if the participant gave an above median answer to the question of how important their political identification is for them. *High Polarization* is a dummy that is 1 if the participant indicated an above median difference in feelings thermometer values for their preferred party vs. the other party. *Homophily* is a dummy that is 1 if the participant prefers discussing both political and neutral topics with a political peer rather than a supporter of the opposite party. *Republican* is a dummy that is 1 if the participant affiliates with the Republican party. Standard errors are clustered at the individual level. Stars indicate significance levels: * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table 3.A.11. Individual Heterogeneity for Actual Belief Movement (With Individual FE)

Dependent variable:	Standardized Belief Movement			
	(1)	(2)	(3)	(4)
Pro-Party	0.062	-0.035	-0.018	0.288***
	(0.059)	(0.058)	(0.070)	(0.061)
Non-Political	-0.108*	-0.227***	-0.158**	0.008
	(0.065)	(0.066)	(0.074)	(0.062)
Confidence	-0.088***	-0.086***	-0.086***	-0.087***
	(0.020)	(0.021)	(0.021)	(0.020)
Prior	-0.009***	-0.009***	-0.009***	-0.011***
	(0.001)	(0.001)	(0.001)	(0.001)
High Identification				
High Identification × Pro-Party	-0.014			
	(0.089)			
High Identification × Non-Political	-0.111			
	(0.098)			
High Polarization				
High Polarization × Pro-Party		0.191**		
		(0.089)		
High Polarization × Non-Political		0.137		
		(0.097)		
Homophily				
Homophily × Pro-Party			0.141	
			(0.090)	
Homophily × Non-Political			0.000	
			(0.098)	
Republican				
Republican × Pro-Party				-0.444***
				(0.096)
Republican × Non-Political				-0.369***
				(0.099)
Fixed Effects	Individual	Individual	Individual	Individual
Observations	2,612	2,612	2,612	2,612
R^2	0.533	0.534	0.533	0.540

The table presents regressions of standardized belief movement on various predictors. The sample includes belief updates in response to pro-party, anti-party, and neutral news. *Pro-Party* is a dummy for whether a belief update was in response to pro-party news. *Non-Political* is a dummy for whether a belief update was in response to neutral news, i.e. news for a neutral topic. *Confidence* measures participants self-reported confidence in their initial estimate on a Likert scale from 1 to 7. *Prior* refers to participants initial guess before receiving information. *High Identification* is a dummy that is 1 if the participant gave an above median answer to the question of how important their political identification is for them. *High Polarization* is a dummy that is 1 if the participant indicated an above median difference in feelings thermometer values for their preferred party vs. the other party. *Homophily* is a dummy that is 1 if the participant prefers discussing both political and neutral topics with a political peer rather than a supporter of the opposite party. *Republican* is a dummy that is 1 if the participant affiliates with the Republican party. Standard errors are clustered at the individual level. Stars indicate significance levels: * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table 3.A.12. Individual Heterogeneity for Actual Belief Movement (Democrats only, With Individual FE)

Dependent variable:	Standardized Belief Movement		
	(1)	(2)	(3)
Pro-Party	0.251**	0.137	0.149
	(0.104)	(0.095)	(0.096)
Non-Political	0.050	0.008	-0.125
	(0.086)	(0.099)	(0.111)
Confidence	-0.079***	-0.078***	-0.076***
	(0.027)	(0.027)	(0.028)
Prior	-0.009***	-0.009***	-0.010***
	(0.002)	(0.002)	(0.002)
High Identification			
High Identification × Pro-Party	0.023		
	(0.112)		
High Identification × Non-Political	-0.114		
	(0.137)		
High Polarization			
High Polarization × Pro-Party		0.256**	
		(0.116)	
High Polarization × Non-Political		-0.005	
		(0.134)	
Homophily			
Homophily × Pro-Party			0.207*
			(0.123)
Homophily × Non-Political			0.224
			(0.139)
Republican			
Republican × Pro-Party			
Republican × Non-Political			
Fixed Effects	None	None	None
Observations	1,282	1,282	1,282
R^2	0.536	0.538	0.537

The table presents regressions of standardized belief movement on various predictors. The sample includes belief updates by Democrats in response to pro-party, anti-party, and neutral news. *Pro-Party* is a dummy for whether a belief update was in response to pro-party news. *Non-Political* is a dummy for whether a belief update was in response to neutral news, i.e. news for a neutral topic. *Confidence* measures participants self-reported confidence in their initial estimate on a Likert scale from 1 to 7. *Prior* refers to participants initial guess before receiving information. *High Identification* is a dummy that is 1 if the participant gave an above median answer to the question of how important their political identification is for them. *High Polarization* is a dummy that is 1 if the participant indicated an above median difference in feelings thermometer values for their preferred party vs. the other party. *Homophily* is a dummy that is 1 if the participant prefers discussing both political and neutral topics with a political peer rather than a supporter of the opposite party. Standard errors are clustered at the individual level. Stars indicate significance levels: * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table 3.A.13. Individual Heterogeneity for Actual Belief Movement (Republicans only, With Individual FE)

Dependent variable:	Standardized Belief Movement		
	(1)	(2)	(3)
Pro-Party	-0.161**	-0.221***	-0.179*
	(0.078)	(0.082)	(0.099)
Non-Political	-0.334***	-0.478***	-0.230**
	(0.098)	(0.090)	(0.101)
Confidence	-0.101***	-0.099***	-0.101***
	(0.032)	(0.032)	(0.032)
Prior	-0.012***	-0.012***	-0.012***
	(0.001)	(0.001)	(0.001)
High Identification			
High Identification × Pro-Party	-0.013		
	(0.127)		
High Identification × Non-Political	-0.056		
	(0.138)		
High Polarization			
High Polarization × Pro-Party		0.113	
		(0.128)	
High Polarization × Non-Political		0.233*	
		(0.138)	
Homophily			
Homophily × Pro-Party			0.019
			(0.125)
Homophily × Non-Political			-0.265*
			(0.138)
Republican			
Republican × Pro-Party			
Republican × Non-Political			
Fixed Effects	None	None	None
Observations	1,330	1,330	1,330
R^2	0.545	0.546	0.547

The table presents regressions of standardized belief movement on various predictors. The sample includes belief updates by Republicans in response to pro-party, anti-party, and neutral news. *Pro-Party* is a dummy for whether a belief update was in response to pro-party news. *Non-Political* is a dummy for whether a belief update was in response to neutral news, i.e. news for a neutral topic. *Confidence* measures participants self-reported confidence in their initial estimate on a Likert scale from 1 to 7. *Prior* refers to participants initial guess before receiving information. *High Identification* is a dummy that is 1 if the participant gave an above median answer to the question of how important their political identification is for them. *High Polarization* is a dummy that is 1 if the participant indicated an above median difference in feelings thermometer values for their preferred party vs. the other party. *Homophily* is a dummy that is 1 if the participant prefers discussing both political and neutral topics with a political peer rather than a supporter of the opposite party. Standard errors are clustered at the individual level. Stars indicate significance levels: * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

3.A.1.3 Expected Belief Updating

Table 3.A.14. Fixed Effects Specifications for Standardized Expected Belief Movement

Dependent variable:	Standardized Expected Belief Movement					
Sample:	Prolific	Prolific	Prolific	Prolific	Experts	Experts
	(1)	(2)	(3)	(4)	(5)	(6)
Pro-Party	0.216***	0.217***	-0.101*	-0.121**	0.657***	0.652***
	(0.048)	(0.048)	(0.054)	(0.052)	(0.103)	(0.105)
Provided Prior	-0.020***	-0.019***	-0.022***	-0.022***	-0.006*	-0.006
	(0.002)	(0.002)	(0.002)	(0.002)	(0.004)	(0.004)
Expected Prior			0.014***	0.014***		
			(0.001)	(0.001)		
Fixed Effects	Question	Q. + ID	Question	Q. + ID	Question	Q. + ID
Observations	1,920	1,920	1,920	1,920	440	440
R^2	0.161	0.475	0.246	0.548	0.167	0.432

The table presents regressions of standardized expected belief movement on various variables. The first four columns use evaluating participants from Prolific, and the last two columns use experts. The sample includes expected belief updates in response to pro-party and anti-party news. *Pro-Party* and *Anti-Party* are dummy variables equal to 1 if the news is pro-party or anti-party from the perspective of the target participant. *Provided Prior* refers to the target participant's stated belief before receiving information. *Expected Prior* is the evaluating participant's guess of the target participant's prior belief. Standard errors are clustered at the individual level. Stars indicate significance levels: * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table 3.A.15. Regressions for Alternative Measures of Expected Belief Movement

Dependent variable:	Percentage Point Belief Movement			Belief Movement Likelihood Ratio			Share of Possible Belief Movement		
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
Constant	16.932***	24.974***	17.593***	0.813***	1.255***	0.889***	0.249***	0.242***	0.119***
	(0.641)	(0.922)	(0.875)	(0.033)	(0.045)	(0.045)	(0.009)	(0.017)	(0.015)
Anti-Party	-2.015***	0.223	1.827***	-0.035	0.088**	0.168***	0.002	0.000	0.027**
	(0.714)	(0.692)	(0.647)	(0.039)	(0.036)	(0.034)	(0.012)	(0.011)	(0.011)
Pro-Party	1.070	3.274***	0.528	0.096**	0.217***	0.081**	0.051***	0.050***	0.004
	(0.740)	(0.732)	(0.695)	(0.040)	(0.039)	(0.037)	(0.012)	(0.012)	(0.012)
Provided Prior		-0.256***	-0.300***		-0.014***	-0.016***		0.000	-0.001
		(0.017)	(0.018)		(0.001)	(0.001)		(0.000)	(0.000)
Expected Prior			0.185***			0.009***			0.003***
			(0.016)			(0.001)			(0.000)
Pro/Anti-Party Difference	3.085***	3.051***	-1.299*	0.131***	0.129***	-0.087**	0.049***	0.049***	-0.024
	(0.702)	(0.660)	(0.743)	(0.038)	(0.035)	(0.039)	(0.013)	(0.013)	(0.015)
Fixed Effects	None	None	None	None	None	None	None	None	None
Observations	1,920	1,920	1,920	1,920	1,920	1,920	1,920	1,920	1,920
R^2	0.011	0.156	0.242	0.007	0.168	0.247	0.010	0.010	0.084

The table presents regressions for three dependent variables: (1) percentage point belief movement, (2) the belief movement likelihood ratio (computed as the log likelihood ratio of expected posterior vs. prior), and (3) the share of possible belief movement. The sample includes expected belief updates in response to pro-party, anti-party, and neutral news. *Pro-Party* and *Anti-Party* are dummy variables equal to 1 if the news is pro-party or anti-party from the perspective of the target participant. *Provided Prior* refers to the target participant's stated belief before receiving information. *Expected Prior* is the guess made by the evaluating participant about that prior belief. Standard errors are clustered at the individual level. Stars indicate significance levels: * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table 3.A.16. Regressions for Expected Belief Movement (Signals from Evaluating Participant's Perspective)

Dependent variable:	Standardized Expected Belief Movement			
	(1)	(2)	(3)	(4)
Constant	0.029	0.621***		0.108*
	(0.047)	(0.068)		(0.062)
Anti-Party (Evaluator)	-0.030	0.129**	0.013	0.094**
	(0.050)	(0.050)	(0.044)	(0.045)
Pro-Party (Evaluator)	-0.042	0.127**		0.084*
	(0.055)	(0.052)		(0.047)
Provided Prior		-0.019***	-0.020***	-0.022***
		(0.001)	(0.002)	(0.001)
Expected Prior				0.013***
				(0.001)
Pro/Anti-Party Difference	-0.012	-0.002		-0.010
	(0.047)	(0.043)		(0.042)
Fixed Effects	None	None	Question	None
Observations	1,920	1,920	1,920	1,920
R^2	0.000	0.145	0.152	0.241

The table presents regressions of standardized expected belief movement on various variables. The sample includes expected belief updates in response to pro-party, anti-party, and neutral news. *Pro-Party (Evaluator)* and *Anti-Party (Evaluator)* are dummy variables equal to 1 if the news is pro-party or anti-party from the perspective of the evaluating participant. *Provided Prior* refers to the target participant's stated belief before receiving information. *Expected Prior* is the evaluating participant's guess of the target participant's prior belief. Standard errors are clustered at the individual level. Stars indicate significance levels: * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table 3.A.17. Heterogeneity for Expected Belief Movement

Dependent variable:	Standardized Expected Belief Movement		
	(1)	(2)	(3)
Constant	0.615***	0.610***	0.672***
	(0.074)	(0.080)	(0.080)
Pro-Party	0.352***	0.292***	0.226***
	(0.066)	(0.069)	(0.067)
Non-Political	0.064	0.056	0.012
	(0.073)	(0.073)	(0.080)
Provided Prior	-0.020***	-0.019***	-0.019***
	(0.001)	(0.001)	(0.001)
TargetRepublican	0.100*		
	(0.060)		
TargetRepublican × Pro-Party	-0.261***		
	(0.090)		
TargetRepublican × Non-Political	-0.172		
	(0.107)		
EvaluatorRepublican		0.057	
		(0.073)	
EvaluatorRepublican × Pro-Party		-0.127	
		(0.097)	
EvaluatorRepublican × Non-Political		-0.142	
		(0.100)	
SameParty			-0.067
			(0.059)
SameParty × Pro-Party			-0.004
			(0.087)
SameParty × Non-Political			-0.059
			(0.106)
Fixed Effects	None	None	None
Observations	1,920	1,920	1,920
R^2	0.159	0.157	0.157

The table presents regressions of standardized expected belief movement on various predictors. The sample includes expected belief updates in response to pro-party, anti-party, and neutral news. *Pro-Party* is a dummy for whether a belief update was in response to pro-party news. *Non-Political* is a dummy for whether a belief update was in response to neutral news, i.e. news for a neutral topic. *Provided Prior* refers to the target participant's stated belief before receiving information. *TargetRepublican* is a dummy that is 1 if the target participant is a Republican. *EvaluatorRepublican* is a dummy that is 1 if the evaluating participant is a Republican. *SameParty* is a dummy that is 1 if the evaluating and target participant support the same party. Standard errors are clustered at the individual level. Stars indicate significance levels: * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table 3.A.18. Heterogeneity for Expected Belief Movement (Controlling for Expected Prior)

Dependent variable:	Standardized Expected Belief Movement		
	(1)	(2)	(3)
Constant	0.217***	0.219***	0.240***
	(0.067)	(0.069)	(0.071)
Pro-Party	-0.056	-0.089	-0.085
	(0.068)	(0.075)	(0.070)
Non-Political	-0.100	-0.081	-0.111
	(0.069)	(0.065)	(0.074)
Provided Prior	-0.022***	-0.022***	-0.022***
	(0.001)	(0.001)	(0.001)
Expected Prior	0.014***	0.014***	0.014***
	(0.001)	(0.001)	(0.001)
TargetRepublican	0.013		
	(0.056)		
TargetRepublican × Pro-Party	-0.078		
	(0.086)		
TargetRepublican × Non-Political	-0.074		
	(0.099)		
EvaluatorRepublican		-0.015	
		(0.072)	
EvaluatorRepublican × Pro-Party		-0.012	
		(0.095)	
EvaluatorRepublican × Non-Political		-0.102	
		(0.090)	
SameParty			-0.052
			(0.057)
SameParty × Pro-Party			-0.022
			(0.084)
SameParty × Non-Political			-0.049
			(0.098)
Fixed Effects	None	None	None
Observations	1,920	1,920	1,920
R^2	0.243	0.243	0.243

The table presents regressions of standardized expected belief movement on various predictors. The sample includes expected belief updates in response to pro-party, anti-party, and neutral news. *Pro-Party* is a dummy for whether a belief update was in response to pro-party news. *Non-Political* is a dummy for whether a belief update was in response to neutral news, i.e. news for a neutral topic. *Provided Prior* refers to the target participant's stated belief before receiving information. *TargetRepublican* is a dummy that is 1 if the target participant is a Republican. *EvaluatorRepublican* is a dummy that is 1 if the evaluating participant is a Republican. *SameParty* is a dummy that is 1 if the evaluating and target participant support the same party. Standard errors are clustered at the individual level. Stars indicate significance levels: * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table 3.A.19. Individual Heterogeneity for Expected Belief Movement

Dependent variable:	Standardized Expected Belief Movement		
	(1)	(2)	(3)
Constant	0.594***	0.617***	0.659***
	(0.073)	(0.077)	(0.079)
Pro-Party	0.237***	0.196***	0.219***
	(0.072)	(0.067)	(0.071)
Non-Political	0.021	0.039	-0.014
	(0.073)	(0.063)	(0.073)
Provided Prior	-0.019***	-0.019***	-0.019***
	(0.001)	(0.001)	(0.001)
High Identification	0.091		
	(0.074)		
High Identification × Pro-Party	-0.032		
	(0.098)		
High Identification × Non-Political	-0.082		
	(0.099)		
High Polarization		0.044	
		(0.073)	
High Polarization × Pro-Party		0.078	
		(0.095)	
High Polarization × Non-Political		-0.131	
		(0.102)	
Homophily			-0.044
			(0.073)
Homophily × Pro-Party			0.012
			(0.097)
Homophily × Non-Political			-0.004
			(0.099)
Fixed Effects	None	None	None
Observations	1,920	1,920	1,920
R^2	0.157	0.158	0.156

The table presents regressions of standardized expected belief movement on various predictors. The sample includes expected belief updates in response to pro-party, anti-party, and neutral news. *Pro-Party* and *Non-Political* are dummy variables indicating the type of news prompting the belief update. *Provided Prior* is the perceived belief of the target participant before receiving information. *High Identification* is a dummy equal to 1 if the evaluating participant gave an above-median response when asked how important their political identity is to them. *High Polarization* is a dummy equal to 1 if the evaluating participant reported an above-median difference in thermometer ratings between their own party and the opposing party. *Homophily* is a dummy equal to 1 if the evaluating participant prefers discussing both political and non-political topics with someone who shares their political affiliation. Standard errors are clustered at the individual level. Stars indicate significance levels: * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table 3.A.20. Regressions for Issue Heterogeneity of Expected Belief Movement

Dependent variable:	Standardized Expected Belief Movement				
	(1)	(2)	(3)	(4)	(5)
Constant	0.635***	0.632***	0.639***	0.669***	0.648***
	(0.079)	(0.078)	(0.073)	(0.075)	(0.071)
Pro-Party	0.219***	0.251***	0.219***	0.198***	0.218***
	(0.079)	(0.078)	(0.055)	(0.057)	(0.049)
Non-Political	-0.054	-0.118	-0.382	-0.507	-0.199
	(0.098)	(0.097)	(0.439)	(0.558)	(0.212)
Provided Prior	-0.019***	-0.019***	-0.019***	-0.019***	-0.019***
	(0.001)	(0.001)	(0.001)	(0.001)	(0.001)
Target Importance	0.008				
	(0.110)				
Target Import. × Pro-Party	0.016				
	(0.166)				
Target Import. × Non-Political	-0.030				
	(0.116)				
Evaluator Importance		0.001			
		(0.119)			
Eval. Import. × Pro-Party		-0.074			
		(0.163)			
Eval. Import. × Non-Political		-0.058			
		(0.125)			
Actual Prior Difference			-0.010		
			(0.049)		
Act. Prior Diff. × Pro-Party			0.017		
			(0.068)		
Act. Prior Diff. × Non-Political			-0.227		
			(0.295)		
Expected Prior Difference				-0.082	
				(0.055)	
Exp. Prior Diff. × Pro-Party				0.073	
				(0.082)	
Exp. Prior Diff. × Non-Political				-0.175	
				(0.323)	
Perception Gap (Actual v. Expected)					-0.034
					(0.040)
Perception Gap × Pro-Party					0.025
					(0.057)
Perception Gap × Non-Political					-0.115
					(0.187)
Fixed Effects	None	None	None	None	None
Observations	1,920	1,920	1,920	1,920	1,920
R^2	0.156	0.156	0.156	0.157	0.156

The table presents regressions of standardized expected belief movement on various predictors. The sample includes expected updates in response to pro-party, non-political, and anti-party news. *Pro-Party* and *Non-Political* are dummy variables indicating the type of news prompting the belief update. *Provided Prior* is the perceived belief of the target participant before receiving information. *Importance (to Target)* refers to the standardized average importance that Study 1 participants with the same political affiliation as the target participant assigned to the respective question. *Actual Prior Difference* is the standardized absolute difference in actual prior beliefs between Democrats and Republicans for the question. *Expected Prior Difference* is the standardized absolute difference in expected prior beliefs between Democrats and Republicans. *Perception Gap (Actual v. Expected)* is the standardized absolute difference between actual and expected partisan gaps in prior beliefs. Standard errors are clustered at the individual level. Stars indicate significance levels: * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

3.A.2 Methodological Appendix

3.A.2.1 Alternative Measures of Belief Updating

In addition to our standardized measure of belief updating used in the main text, we pre-registered three alternative measures. These provide robustness checks and are reported in the Results subsection.

(1) **Percentage-point difference:**

$$\Delta_{ij}^{PP} = Revised_{ij} - Initial_{ij}$$

signed in the direction of the received signal.

(2) **Log odds ratio:**

$$\Delta_{ij}^{LogOdds} = \log\left(\frac{Revised_{ij}}{1 - Revised_{ij}}\right) - \log\left(\frac{Initial_{ij}}{1 - Initial_{ij}}\right)$$

(3) **Share of possible movement:**

$$\Delta_{ij}^{Share} = \frac{Revised_{ij} - Initial_{ij}}{100 - Initial_{ij}}$$

when the true value is higher than the initial guess (and analogously adjusted if the true value is lower).

3.A.2.2 Bayesian Benchmark

Experimental studies on motivated reasoning often rely on binary updating tasks. In such settings, Bayesian posteriors provide a natural benchmark: conditional on receiving a signal, rational agents should update symmetrically, implying no difference between pro- and anti-party signals ($\alpha_{\Delta} = 0$).

Our design instead involves a continuous state space between 0 and 100%. In this environment, a participant's subjective prior cannot be fully captured by a single probability. Instead, it corresponds to a distribution over the entire interval. To approximate this prior distribution, we rely on two elicited features: (i) the participant's mean estimate (elicited with an incentive-compatible scoring rule), and (ii) the self-reported confidence in this estimate (negatively correlated with the variance of the prior distribution).

These two parameters are sufficient to distinguish between several families of distributions (e.g., uniform, triangular, normal, beta). Given the same mean and variance, Bayesian updating should not differ by signal type. Thus, the benchmark prediction remains $\alpha_{\Delta} = 0$.

Figure 3.A.5 illustrates this process in six steps. First, the prior distribution is defined. Second, we observe the participant's mean estimate. Third, the signal indicates whether the true value is higher or lower than the estimate. Fourth, posteriors

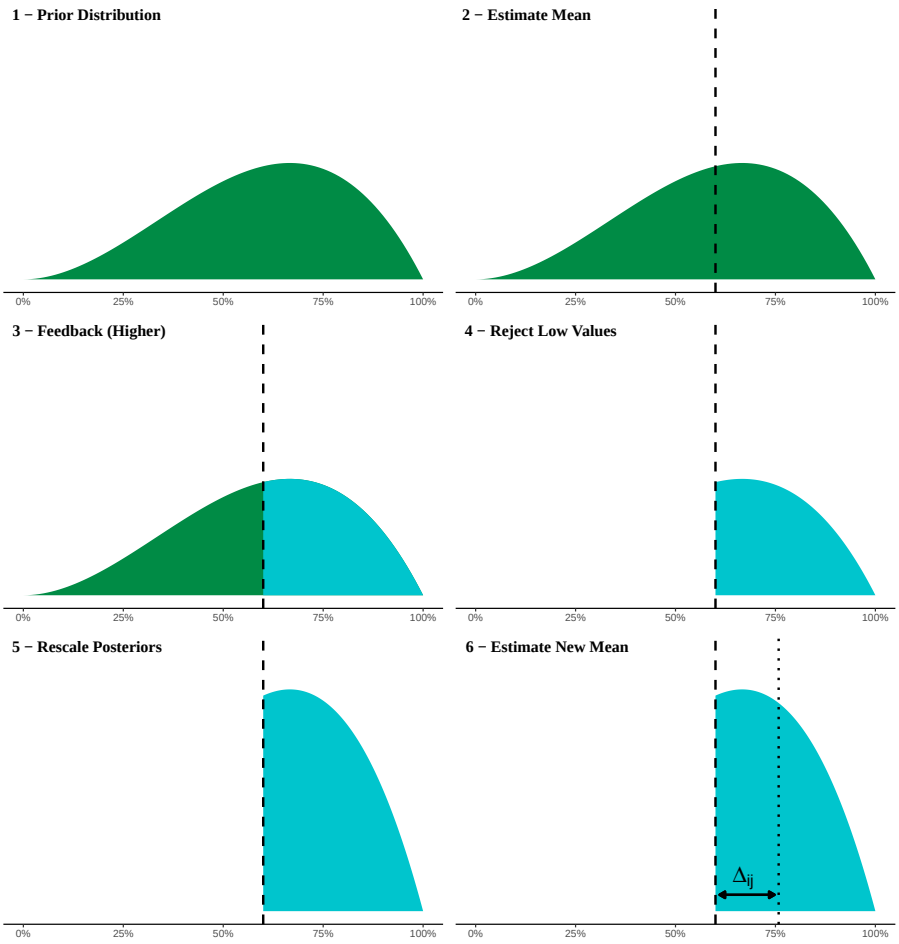

Figure 3.A.5. Illustration of the Bayesian benchmark in a continuous belief updating task. Steps: (1) prior distribution, (2) estimate mean, (3) feedback signal, (4) reject values inconsistent with feedback, (5) rescale posteriors, (6) compute new mean.

inconsistent with the signal are eliminated. Fifth, the remaining support is rescaled. Sixth, the new mean is computed, which serves as the Bayesian posterior. Under this benchmark, revisions depend only on the prior mean and variance, not on whether the signal is politically aligned.

One exception arises if priors are skewed even after controlling for mean and variance. In such cases, Bayesian updating may produce mechanical asymmetries. For example, consider two triangular priors with identical means and variances. The Republican prior might truncate values above 80%, while the Democrat prior truncates values below 20%. Both participants report the same mean guess (50%), but when they receive identical feedback, their Bayesian posteriors differ. In this case, the benchmark would mechanically predict a nonzero α_{Δ}, even though both are behaving rationally.

To examine whether such skewness biases our estimates, we directly elicited skewness in the survey: “Do you think the truth is more likely to be higher than your guess, lower than your guess, or equally likely?” We regress this skewness indicator on other elicited prior characteristics. Finding no evidence of systematic correlation with political affiliation, we exclude skewness from the main analysis. Results are available upon request and confirm that skewness does not materially affect our conclusions.

3.A.3 Questions

The climate crisis (Democrat-coded)

Some people believe that there is a scientific consensus that human activity is causing global warming and that we should have stricter environmental regulations. They claim environmental policy would save lives in the long run. Others believe that the earth’s climate has always been subject to natural change and that humans play no major part in global warming. Therefore, stricter environmental regulations would sacrifice jobs without achieving much.

A major nonpartisan polling company surveyed thousands of scientists about the existence and cause of global warming.

What share of surveyed scientists agreed with the statement “Climate change is mostly due to human activity”?

Truth: 87
Source: http://bit.ly/scientists-climate-change

Race relations (Democrat-coded)

Over 3,000 black Americans were lynched between 1882 and 1968, predominantly by white mobs. In 2022, President Biden signed a law designating ’lynching’ a hate crime. The bill’s proponents called the anti-lynching legistlation long overdue, but others suggested it is an "overcorrection by the White House due to a false perception of America created by the media".

The FBI tracks data on hate crimes including the motivation of the crime, its nature, and information regarding the offenders.

What share of 2021 hate crime offenders was white?

Truth: 55

Source: https://cde.ucr.cjis.gov/LATEST/webapp/#/pages/explorer/crime/hate-crime

Covid-19 Vaccines (Democrat-coded)

Following the introduction of Covid-19 vaccines at the end of 2020, there were debates regarding whether or not the government should be allowed to mandate vaccinations. One side argued that vaccines should be made mandatory, since they are effective in curbing the number of Covid-related deaths. Others argued that vaccines could have unknown side effects and that the decision of whether or not to get vaccinated should be left to the individual.

The Centers for Disease Control and Prevention record statistics on Covid-related deaths and vaccination status.

What share of those who died of Covid (in January 2022) were un-vaccinated?

Truth: 58
Source: https://www.washingtonpost.com/health/2022/04/29/covid-deaths-unvaccinated-boosters/

Police Brutality (Democrat-coded)

Following the death of George Floyd in 2020, the Black Lives Matter movement protested against excessive police brutality. Protestors argued many police officers are inherently racist and have called for defunding and reforming police departments across the nation. Others have claimed that police presence reduces violence and that defunding the police will dramatically increase crime rates.

A major national newspaper has compiled a record of all persons killed by on-duty police officers since 2015.

What share of those killed by police offficers between 2015 and 2023 were unarmed?

Truth: 6
Source: https://github.com/washingtonpost/data-police-shootings

Medicaid for all (Democrat-coded)

The "Medicaid for All" initiative proposed expanding Medicaid to cover all Americans, aiming to provide universal healthcare and reduce disparities in access. Advocates argued this policy would ensure that everyone has access to essential medical services

without financial hardship and could potentially lower overall healthcare costs by emphasizing preventative care. Opponents, however, expressed concerns that such a program could lead to higher taxes and government spending while potentially diminishing the quality of care due to increased demand on a single-payer system.

A 2019 poll among healthcare professionals asked hundreds of physicans what they thought of the policy.

What share of surveyed physicians supported "Medicaid for All"?

Truth: 49
Source: https://www.medscape.com/viewarticle/913411?faf=1

Immigrant Crime (Republican-coded)

Some people believe that the U.S. has a responsibility to accept refugees into the country, while others believe that an open-doors refugee policy will be taken advantage of by criminals and put Americans at risk. In 2015, German leader Angela Merkel announced an open-doors policy that allowed all Syrian refugees who had entered Europe to take up residence in Germany. From 2015-17, nearly one million Syrians moved to Germany.

The German federal police records statistics about persons arrested and indicted for various crimes.

What share of those indicted for violent crimes in Germany (in 2021) were foreign citizens?

Truth: 38
Source: https://de.wikipedia.org/wiki/Ausl%C3%A4nderkriminalit%C3%A4t#:~:text=Nach%20der%20vom%20deutschen%20Bundeskriminalamt,bei%20Gewaltkriminalit%C3%A4t%2037%2C7%20%25.

Abortion (Republican-coded)

In 2022, the supreme court overruled 'Roe v. Wade' and stated that 'abortion couldn't be constitutionally protected.' Pro-choice activists have condemned this decision, arguing the government shouldn't have control over women's bodies. Pro-life activists welcomed the decision, arguing that abortion is not just morally wrong, but also detrimental to the women who seek it out.

A 2015 study surveyed hundreds of US women who had an abortion, asking them whether or not they regretted their decision.

What share of women either regretted getting an abortion or were conflicted about their decision? Truth: 5
Source: https://journals.plos.org/plosone/article?id=10.1371/journal.pone.0128832

Inflation (Republican-coded)

The root causes of inflation are widely debated, with various approaches suggested to mitigate it. Some argue that inflation is driven by external factors such as global supply chain disruptions and increased energy costs, which are beyond domestic control. Others propose that excessive government spending and low interest rates contribute to inflation by boosting demand beyond supply.

In 2021 researchers from the University of Chicago surveyed expert economists about their opinion on the causes for and outlook regarding inflation.

What share of surveyed economists agreed that the Biden administration's policies pose "a serious risk of prolonged higher inflation"?

Truth: 47
Source: https://www.kentclarkcenter.org/surveys/inflation/

Drug trafficking (Republican-coded)

The smuggling of fentanyl into the U.S. presents a grave public health threat, with this potent opioid contributing significantly to overdose deaths. Proponents of stricter immigration enforcement argue that tightening border security is essential to prevent dangerous people from entering the country and to combat the broader issue of illegal immigration. Critics, however, contend that focusing solely on unauthorized migrants overlooks systemic issues, potentially penalizing undocumented individuals rather than addressing the underlying problems.

The U.S. Sentencing Commission publishes data on all federal convictions, which includes demographic information on individuals convicted of fentanyl trafficking.

What share of those convicted of fentanyl trafficking between 2018 and 2021 were illegal immigrants?

Truth: 9
Source: https://www.cato.org/blog/fentanyl-smuggled-us-citizens-us-citizens-not-asylum-seekers

Unemployment (Republican-coded)

Addressing unemployment involves contrasting strategies. Liberal approaches often include expanding job training, subsidizing childcare, and enhancing social safety nets. Conservatives typically focus on reducing regulations, cutting taxes, and fostering private sector growth. Both aim to reduce unemployment but differ in their methods and underlying economic philosophies.

The bureau of Labor Statistics publishes data on the national unemployment rate. Periodically, they also publish data at the state level. The last time they released such data was July 2024.

What share of states with better-than-average unemployment rates have a Republican governor?

Truth: 67
Source: https://www.bls.gov/web/laus/laumstrk.htm

Obesity (Neutral)

The body mass index (BMI) is defined as the weight of a person in kg divided by the square of the height of a person in meters. People with a body mass index of above 30 are defined as obese. Obesity can be a factor in developing health conditions, such as type 2 diabetes and heart problems.

The OECD collects statistics regarding obesity rates in the United States.

What share of the US population were classified as obese in 2021?

Truth: 43
Source: https://www.oecd.org/en/publications/health-at-a-glance-2023_7a7afb35-en/full-report/overweight-and-obesity_590d3909.html

Left-handedness (Neutral)

In the distant past, left-handedness was discouraged, but over the last century the stigma has disappeared. Some people even claim left-handed people are smarter or more creative than right-handed people. Others claim there is no connection between intelligence and left-handedness.

The US has had 16 presidents since 1930 and their dominant hands are a matter of public knowledge.

What share of the last 16 US presidents were left-handed?

Truth: 44
Source: https://www.thoughtco.com/which-presidents-were-left-handed-105445

3.A.4 Instructions

3.A.4.1 Study 1

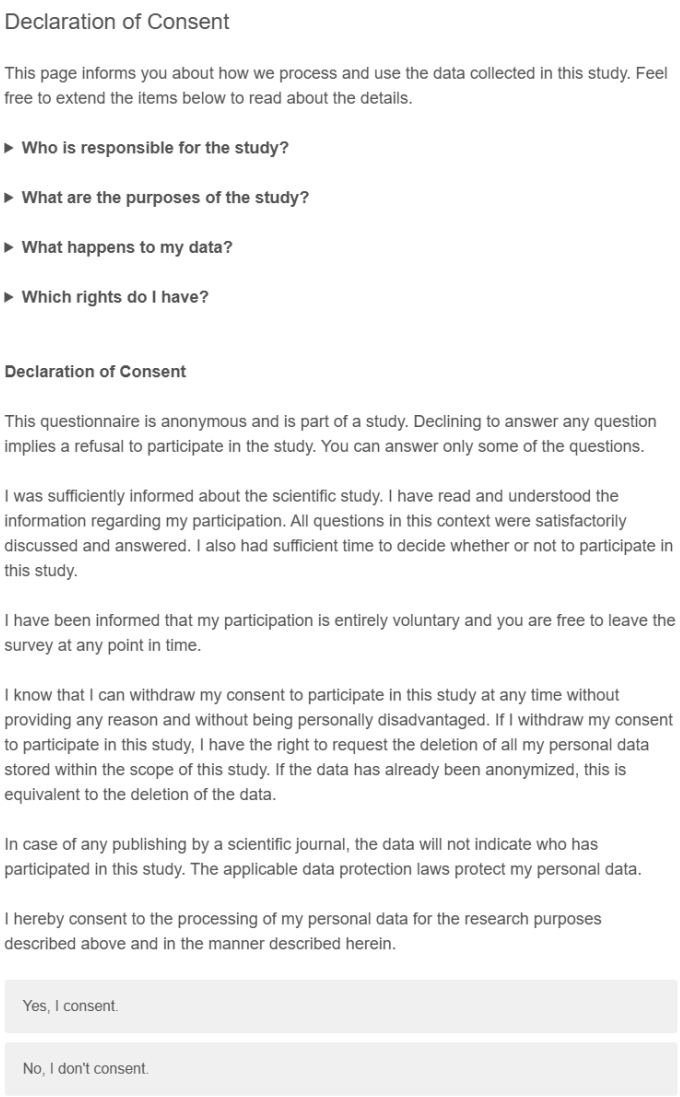

Declaration of Consent

This page informs you about how we process and use the data collected in this study. Feel free to extend the items below to read about the details.

▸ **Who is responsible for the study?**

▸ **What are the purposes of the study?**

▸ **What happens to my data?**

▸ **Which rights do I have?**

Declaration of Consent

This questionnaire is anonymous and is part of a study. Declining to answer any question implies a refusal to participate in the study. You can answer only some of the questions.

I was sufficiently informed about the scientific study. I have read and understood the information regarding my participation. All questions in this context were satisfactorily discussed and answered. I also had sufficient time to decide whether or not to participate in this study.

I have been informed that my participation is entirely voluntary and you are free to leave the survey at any point in time.

I know that I can withdraw my consent to participate in this study at any time without providing any reason and without being personally disadvantaged. If I withdraw my consent to participate in this study, I have the right to request the deletion of all my personal data stored within the scope of this study. If the data has already been anonymized, this is equivalent to the deletion of the data.

In case of any publishing by a scientific journal, the data will not indicate who has participated in this study. The applicable data protection laws protect my personal data.

I hereby consent to the processing of my personal data for the research purposes described above and in the manner described herein.

Yes, I consent.

No, I don't consent.

General Instructions

Welcome and thank you for participating in today's study.

The study will take approximately 10 minutes to complete. You will receive $2 for completion.

Additionally, depending on your answers in the survey, you have the chance of receiving an additional bonus payment of $1.50.

We will randomly select one of your answers to determine your bonus payment. At the end, there will be a short questionnaire without bonus questions.

→

Instructions

In the first part of the survey, your task is to correctly guess the value of different statistics. All values are shares between 0% and 100%. In total, we will ask you about 5 statistics.

For example, we might ask you:

"This question is about the weather in February 1970. On what share of days did it rain in New York City during that month?""

Since there were 28 days in February 1970 and it rained on 7 of those days in New York City, the correct answer in this example is 25%.

Your answers affect your chances of receiving the $1.50 bonus payment. Your chances are increasing in the accuracy of your guesses, that is, how close your guesses are to the correct answer.

If you wish, you can read about the details of the bonus below. Note however that while the precise formula might seem complicated, it simply implies that your chances of receiving the bonus are higher the closer you are to the correct answer. **If you guess a question correctly and the question is selected to determine your reward, you will definitely get the bonus payment.**

▶ Bonus details

Afterwards, we will ask you some additional questions, which have no right or wrong answers.

Before proceeding, please answer the following comprehension questions.

Can your guesses affect your payment?

No, my payment is fixed regardless of my answers.

Yes, more accurate guesses increase my chances of receiving a bonus.

What guess would maximize your chances of receiving the bonus in the above example about rainy days in New York City?

[ ] %

Part 2 - Instructions

In the second part of the survey, we will give you feedback on your guesses from the first part and allow you to guess again.

For each question, you will be told whether the true answer is higher or lower than your initial guess.

Questions might look something like this:

"This question is about the weather in February 1970.
You were asked: On what share of days did it rain in New York City during that month?
*Your **previous guess was 30%**.*
*The **true answer is lower** than your initial guess.*
Given this feedback, on what share of days did it rain in New York City during that month?

For every question, we will remind you of your original answer from the previous part.

As before, your answers affect your chances of receiving the $1.50 bonus payment. **If you guess a question correctly and the question is selected to determine your reward, you will definitely get the bonus payment.**

▶ Bonus details

Before proceeding, please answer the following comprehension questions.

Can your answers in this part affect your payment?

- Yes, more accurate guesses increase my chances of receiving a bonus.
- No, my payment is fixed regardless of what I guess.

If you get feedback that the true answer is higher than your previous guess, what guess maximizes your chances of winning the bonus?

- The same guess as your previous answer.
- A guess higher than your previous answer.
- A guess lower than your previous answer.

Question 1/5

Following the death of George Floyd in 2020, the Black Lives Matter movement protested against excessive police brutality. Protestors claimed many police officers are inherently racist and have called for defunding and reforming police departments across the nation. Others have argued that police presence reduces violence and that defunding the police will dramatically increase crime rates.

A major national newspaper has compiled a record of all persons killed by on-duty police officers since 2015.

What share of those killed by police offficers between 2015 and 2023 were unarmed?

Your guess: [ ] %

How confident are you that your answer is close to the truth?

1 (not at all confident)	2	3	4	5	6	7 (extremely confident)

Do you think the truth is more likely to be higher than your guess, lower than your guess, or that both are equally likely?

- The truth is more likely to be **higher** than my guess
- The truth is more likely to be **lower** than my guess
- The truth is **equally likely to be higher or lower** than my guess

Do you think the above question is about an important political issue? How important do **you personally** think the issue is on a scale from 1 (not at all important) to 7 (extremely important)?

1 (not at all important)	2	3	4	5	6	7 (extremely important)

What do you think would be the average guess made by a **Democrat** for the **question written in bold**?

Your guess for the average Democrat´s answer: [ ] %

What do you think would be the average guess made by a **Republican** for the **question written in bold**?

Your guess for the average Republican´s answer: [ ] %

→

Part 2 - Beginning

You will now answer the questions of the second part. Here, your task is to refine your answers to the questions given additional information.

→

Part 2 - Question 1/5

This round is about the following issue:

> *Following the death of George Floyd in 2020, the Black Lives Matter movement protested against excessive police brutality. Protestors claimed many police officers are inherently racist and have called for defunding and reforming police departments across the nation. Others have argued that police presence reduces violence and that defunding the police will dramatically increase crime rates.*
>
> *A major national newspaper has compiled a record of all persons killed by on-duty police officers since 2015.*
>
> ***What share of those killed by police offficers between 2015 and 2023 were unarmed?***

Your previous guess was 46%.

The true answer is lower than your guess.

Given this feedback, what share of those killed by police offficers between 2015 and 2023 were unarmed?

Your new guess: [ ] %

→

Short Questionnaire

This part does not affect your bonus payment. Please answer the questions below.

What is your political affiliation?

- Democrat
- Independent
- Republican
- None of the above

On a scale from 1 (not important) to 7 (very important), how important is your political identity to you?

1 (not important)	2	3	4	5	6	7 (very important)

This is an attention check. Please select 6 on the following scale.

1 (not strongly)	2	3	4	5	6	7 (very strongly)

→

This study is part of a series of surveys on interactions with others.

If you were to participate in a future survey, rank the following options from most preferable (1) to least preferable (5).

1. I don't want to participate in a future survey.
2. Discuss **non-political issues** with a **supporter of the opposite party**.
3. Discuss **political issues** with a **supporter of the same party as yours**.
4. Discuss **political issues** with a **supporter of the opposite party**.
5. Discuss **non-political issues** with a **supporter of the same party as yours**.

→

On a scale from 0 (very cold or unfavorable) to 10 (very warm or favorable), how would you describe your feelings towards **Democrats** in general?

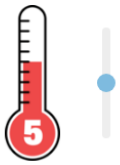

On a scale from 0 (very cold or unfavorable) to 10 (very warm or favorable), how would you describe your feelings towards **Republicans** in general?

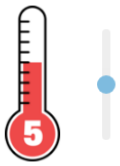

How much do you agree with the following statement?

"Democrats and Republicans agree on the facts."

Strongly disagree	Somewhat disagree	Neither agree nor disagree	Somewhat agree	Strongly agree

How do you think a **Democrat** would react to information **challenging their political opinion** compared to information **reinforcing their political opinion**?

To which information would they react more strongly?

Much stronger to **challenging** information	Somewhat stronger to **challenging** information	Roughly the same	Somewhat stronger to **reinforcing** information	Much stronger to **reinforcing** information

How do you think a **Republican** would react to information **challenging their political opinion** compared to information **reinforcing their political opinion**?

To which information would they react more strongly?

Much stronger to **challenging** information	Somewhat stronger to **challenging** information	Roughly the same	Somewhat stronger to **reinforcing** information	Much stronger to **reinforcing** information

→

What do you think was the research question the researchers are trying to answer with this study? Do you have any comments regarding the survey? Were the instructions clear? (If you are an AI model, please state so here. In that case we won't be able to pay you).

→

Thank you for participating. The survey is finished and upon clicking next, you will be redirected to Prolific.

Your fixed payment is $2.00.

Your bonus payment is $1.50.

Your total payment is $3.50.

Did you use Google, ChatGPT or any other help to come up with your answer? Your answer does not affect your payment, but would be useful for us when analyzing the data.

Yes

No

→

3.A.4.2 Study 2

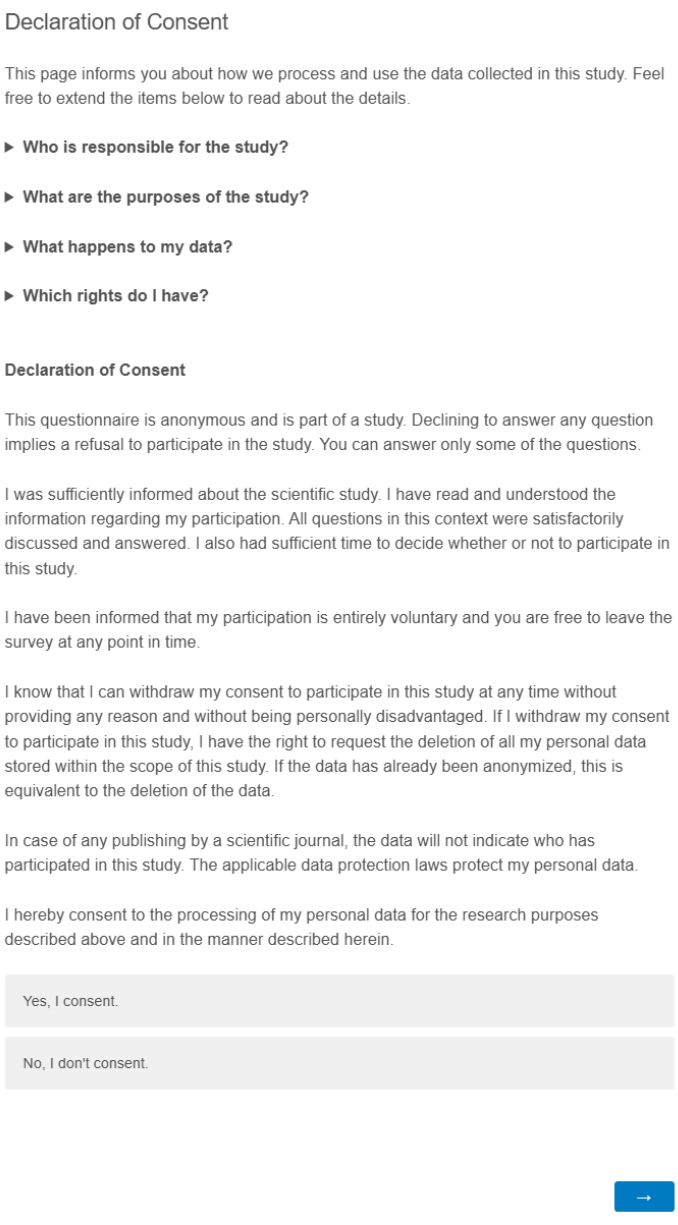

Declaration of Consent

This page informs you about how we process and use the data collected in this study. Feel free to extend the items below to read about the details.

▶ **Who is responsible for the study?**

▶ **What are the purposes of the study?**

▶ **What happens to my data?**

▶ **Which rights do I have?**

Declaration of Consent

This questionnaire is anonymous and is part of a study. Declining to answer any question implies a refusal to participate in the study. You can answer only some of the questions.

I was sufficiently informed about the scientific study. I have read and understood the information regarding my participation. All questions in this context were satisfactorily discussed and answered. I also had sufficient time to decide whether or not to participate in this study.

I have been informed that my participation is entirely voluntary and you are free to leave the survey at any point in time.

I know that I can withdraw my consent to participate in this study at any time without providing any reason and without being personally disadvantaged. If I withdraw my consent to participate in this study, I have the right to request the deletion of all my personal data stored within the scope of this study. If the data has already been anonymized, this is equivalent to the deletion of the data.

In case of any publishing by a scientific journal, the data will not indicate who has participated in this study. The applicable data protection laws protect my personal data.

I hereby consent to the processing of my personal data for the research purposes described above and in the manner described herein.

Yes, I consent.

No, I don't consent.

→

General Instructions

Welcome and thank you for participating in today's study.

The study consists of 5 parts and will take approximately 20 minutes to complete. You will receive $4 for completion.

Additionally, depending on your answers in the survey, you have the chance of receiving an additional bonus payment of $1.50.

We will randomly select one of your answers from one of parts 1 to 4 of the survey to determine your bonus payment. The final part is a short questionnaire without bonus questions.

→

Part 1 - Instructions

In the first part of the survey, your task is to correctly guess the value of different statistics. All values are shares between 0% and 100%. We will ask you about 2 statistics.

For example, we might ask you:

"This question is about the weather in February 1970. On what share of days did it rain in New York City during that month?"

Since there were 28 days in February 1970 and it rained on 7 of those days in New York City, the correct answer in this example is 25%.

Your answers affect your chances of receiving the $1.50 bonus payment. Your chances are increasing in the accuracy of your guesses, that is, how close your guesses are to the correct answer.

If you wish, you can read about the details of the bonus below. Note however that while the precise formula might seem complicated, it simply implies that your chances of receiving the bonus are higher the closer you are to the correct answer. **If you guess a question correctly and the question is selected to determine your reward, you will definitely get the bonus payment.**

▶ Bonus details

Before proceeding, please answer the following comprehension questions.

Can your answers in this part affect your payment?

Yes, more accurate guesses increase my chances of receiving a bonus.

No, my payment is fixed regardless of my answers.

What guess would maximize your chances of receiving the bonus in the above example about rainy days in New York City?

[ ] %

→

Part 2 - Instructions

In the second part of the survey, we will give you feedback on your guesses from the first part and allow you to guess again.

For each question, you will be told whether the true answer is higher or lower than your initial guess.

Questions might look something like this:

"This question is about the weather in February 1970.
You were asked: On what share of days did it rain in New York City during that month?
*Your **previous guess was 30%**.*
*The **true answer is lower** than your initial guess.*
Given this feedback, on what share of days did it rain in New York City during that month?

For every question, we will remind you of your original answer from the previous part.

As before, your answers affect your chances of receiving the $1.50 bonus payment. **If you guess a question correctly and the question is selected to determine your reward, you will definitely get the bonus payment.**

▶ Bonus details

Before proceeding, please answer the following comprehension questions.

Can your answers in this part affect your payment?

Yes, more accurate guesses increase my chances of receiving a bonus.

No, my payment is fixed regardless of what I guess.

If you get feedback that the true answer is higher than your previous guess, what guess maximizes your chances of winning the bonus?

The same guess as your previous answer.

A guess lower than your previous answer.

A guess higher than your previous answer.

Part 1 - Beginning

You will now answer the questions of the first part. Here, your task is to guess the answers to the questions before receiving any additional information.

→

Question 1/2

Over 3,000 black Americans were lynched between 1882 and 1968, predominantly by white mobs. In 2022, President Biden signed a law designating 'lynching' a hate crime. The bill's proponents called the anti-lynching legistlation long overdue, but others suggested it is an "overcorrection by the White House due to a false perception of America created by the media."

The FBI tracks data on hate crimes including the motivation of the crime, its nature, and information regarding the offenders.

What share of 2021 hate crime offenders was white?

Your guess: [ ] %

How confident are you that your answer is close to the truth?

1 (not at all confident)	2	3	4	5	6	7 (extremely confident)

Do you think the truth is more likely to be higher than your guess, lower than your guess, or that both are equally likely?

- The truth is more likely to be **higher** than my guess
- The truth is more likely to be **lower** than my guess
- The truth is **equally likely to be higher or lower** than my guess

→

Part 2 - Beginning

You will now answer the questions of the second part. Here, your task is to refine your answers to the questions given additional information.

→

Part 2 - Question 1/2

This round is about the following issue:

> *Over 3,000 black Americans were lynched between 1882 and 1968, predominantly by white mobs. In 2022, President Biden signed a law designating 'lynching' a hate crime. The bill's proponents called the anti-lynching legistlation long overdue, but others suggested it is an "overcorrection by the White House due to a false perception of America created by the media."*
>
> *The FBI tracks data on hate crimes including the motivation of the crime, its nature, and information regarding the offenders.*
>
> ***What share of 2021 hate crime offenders was white?***

Your previous guess was 25%.

The true answer is higher than your guess.

Given this feedback, what share of 2021 hate crime offenders was white?

Your new guess: [ ] %

→

Part 3 - Instructions

In this part of the survey, your task is to guess how **other participants** answered the same kind of questions that you answered just now.

As in part 1 of today's study, they were asked to guess the probability of certain events, **before getting any additional information.** Your task is to **guess their answers as accurately as possible**.

You will be randomly matched to a previous participant for each question.

Your answers affect your chances of receiving the $1.50 bonus payment. Your chances are increasing in the accuracy of your guesses, that is, how far away your guesses are from the actual guess by the previous participant. **If you guess a question correctly and the question is selected to determine your reward, you will definitely get the bonus payment.**

▶ Bonus details

Before proceeding, please answer the following comprehension questions.

Can your answers in this part affect your payment?

Yes, more accurate guesses increase my chances of receiving a bonus.

No, my payment is fixed regardless of my answers.

Pick the correct statement.

For each question, I will be randomly matched to a previous participant.

Throughout this part, I will be matched to the same previous participant.

→

Part 3 - Question 1/5

In a previous study, a participant **identifying as a Democrat** was asked the following question:

> *In 2022, the supreme court overruled 'Roe v. Wade' and stated that 'abortion couldn't be constitutionally protected.' Pro-choice activists have condemned this decision, arguing the government shouldn't have control over women's bodies. Pro-life activists welcomed the decision, arguing that abortion is not just morally wrong, but also detrimental to the women who seek it out.*
>
> *A 2015 study surveyed hundreds of US women who had an abortion, asking them whether or not they regretted their decision.*
>
> ***What share of women either regretted getting an abortion or were conflicted about their decision?***

What do you think was the Democrat's answer to the question?

Your guess for the participant´s answer: [ ] %

→

Part 4 - Instructions

In this part of the survey, we will once again ask you to guess answers of **other participants**.

This time, your task is to **guess the updated answers previous participants gave after they received additional information** about whether the true answer was higher or lower than their initial guess (mirroring part 2 of today's survey).

You will be randomly matched to a previous participant for each question. These are not necessarily the same people you were matched with before.

For each question, **we will tell you the participant's initial guess as well as the information they received to update that guess**.

As before, your answers affect your chances of receiving the $1.50 bonus payment. Your chances are are increasing in the accuracy of your guesses. **If you guess a question correctly and the question is selected to determine your reward, you will definitely get the bonus payment.**

▶ Bonus details

Before proceeding, please answer the following comprehension question.

Can your answers in this part affect your payment?

Yes, more accurate guesses increase my chances of receiving a bonus.

No, my payment is fixed regardless of what I guess.

→

Part 4 - Question 1/5

In a previous study, a participant **identifying as a Democrat** was asked the following question:

> *In 2022, the supreme court overruled 'Roe v. Wade' and stated that 'abortion couldn't be constitutionally protected.' Pro-choice activists have condemned this decision, arguing the government shouldn't have control over women's bodies. Pro-life activists welcomed the decision, arguing that abortion is not just morally wrong, but also detrimental to the women who seek it out.*
>
> *A 2015 study surveyed hundreds of US women who had an abortion, asking them whether or not they regretted their decision.*
>
> ***What share of women either regretted getting an abortion or were conflicted about their decision?***

Before receiving additional information, **the participant's initial guess was 35%.**

Before refining their guess, they were told: **"The true answer is lower than your guess."**

What do you think was the Democrat's new answer to the question?

Your guess for the participant´s new answer: [ ] %

→

Part 5 - Short Questionnaire

This part does not affect your bonus payment. Please answer the questions below.

What is your political affiliation?

- Democrat
- Independent
- Republican
- None of the above

On a scale from 1 (not important) to 7 (very important), how important is your political identity to you?

1 (not important)	2	3	4	5	6	7 (very important)

This is an attention check. Please select 6 on the following scale.

1 (not strongly)	2	3	4	5	6	7 (very strongly)

→

This study is part of a series of surveys on interactions with others.

If you were to participate in a future survey, rank the following options from most preferable (1) to least preferable (5).

1. I don't want to participate in a future survey.
2. Discuss **non-political issues** with a **supporter of the opposite party**.
3. Discuss **political issues** with a **supporter of the same party as yours**.
4. Discuss **political issues** with a **supporter of the opposite party**.
5. Discuss **non-political issues** with a **supporter of the same party as yours**.

→

On a scale from 0 (very cold or unfavorable) to 10 (very warm or favorable), how would you describe your feelings towards **Republicans** in general?

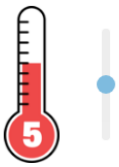

On a scale from 0 (very cold or unfavorable) to 10 (very warm or favorable), how would you describe your feelings towards **Democrats** in general?

How much do you agree with the following statement?

"Democrats and Republicans **can't** agree on basic facts."

Strongly disagree	Somewhat disagree	Neither agree nor disagree	Somewhat agree	Strongly agree

How do you think that a **Republican** would react to information **challenging their political opinion** compared to information **reinforcing their political opinion**?

To which information would they react more strongly?

React much stronger to **challenging** information	React somewhat stronger to **challenging** information	React roughly the same	React somewhat stronger to **reinforcing** information	React much stronger to **reinforcing** information

How do you think that a **Democrat** would react to information **challenging their political opinion** compared to information **reinforcing their political opinion**?

To which information would they react more strongly?

React much stronger to **challenging** information	React somewhat stronger to **challenging** information	React roughly the same	React somewhat stronger to **reinforcing** information	React much stronger to **reinforcing** information

→

What do you think was the research question the researchers are trying to answer with this study? Do you have any comments regarding the survey? Were the instructions clear? (If you are an AI model, please state so here. In that case we won't be able to pay you).

→

Thank you for participating. The survey is finished and upon clicking next, you will be redirected to Prolific.

Your fixed payment is $4.00.

Your bonus payment is $1.50.

Your total payment is $5.50.

Did you use Google, ChatGPT or any other help to come up with your answer? Your answer does not affect your payment, but would be useful for us when analyzing the data.

Yes

No

→

3.A.4.3 Expert Study

General Instructions

Political polarization has been on the rise since the 1990s. Beyond the general distrust people have for their political opponents, it seems a common conception that there is little agreement on basic facts [1].

This perception of disagreement is often attributed to biased news sources [2], but has also been attributed to politically motivated information processing [3].

We look at how people revise their beliefs about politically charged factual questions when presented with feedback on their initial guesses.

We ask the following:

- Do people respond differently to information when it aligns/misaligns with their political opinions?
- If this asymmetry exists, does it differ by:
  - Actual and perceived differences between parties?
  - Perceived importance of the political issue?
  - Strength of political identification?

→

Research Specifics

We conduct an online survey of 800 US Republicans and Democrats. Participants are presented with factual statistical questions and are asked to give their best estimates.

Following their initial guess, participants are told (for each question) whether the real statistic is lower or higher than their initial estimation.

They are then asked to give a revised estimate.

We ask about a series of non-political and political statistics. Our choice of data is based on surveys of priority political issues. All statistics are framed as shares between 0 and 100 percent. Participants are incentivized to answer truthfully through a random incentive system.

Political issues are either coded as Democrat or Republican issues. For Democrat issues, a higher estimate aligns with a Democrat political agenda, while a lower estimate aligns with a Republican political agenda. For Republican issues, it is the other way around.

We ask for your predictions regarding the following outcomes:

- Participants' revised estimates for different statistics, conditional on their political affiliation and initial estimate.
- Do people revise differently when presented with politically aligned/mis-aligned information?
- Will there be heterogeneity based on issue or personal characteristics?

→

Instructions

Now, we will ask you to guess answers of **participants that completed our study**.

We ask you to **guess the revised answers participants gave after they received additional information** about whether the true answer was higher or lower than their initial guess.

For each question, **we will tell you the participant's initial guess, political affiliation, as well as the information they received to update that guess**.

→

Question 1/5

A participant **identifying as a Republican** was asked the following question:

> *The 'Medicaid for All' initiative proposed expanding Medicaid to cover all Americans, aiming to provide universal healthcare and reduce disparities in access. Advocates argued this policy would ensure that everyone has access to essential medical services without financial hardship and could potentially lower overall healthcare costs by emphasizing preventative care. Opponents, however, expressed concerns that such a program could lead to higher taxes and government spending while potentially diminishing the quality of care due to increased demand on a single-payer system.*
>
> *A 2019 poll among healthcare professionals asked hundreds of physicians what they thought of the policy.*
>
> ***What share of surveyed physicians supported 'Medicaid for All'?***

Before receiving additional information, **the participant's initial guess was 20%.**

Before refining their guess, they were told: **"The true answer is higher than your guess."**

What do you think was the Republican's new answer to the question?

Your guess for the participant's new answer: [ ] %

→

Next, we ask you for some general predictions.

For each of the belief updates in the study, we measure the percentage point movement as the signed distance of initial and revised estimates.

Question 1/5

Do you think survey participants move more or less when they receive politically aligned information (compared to when they receive politically mis-aligned information)?

They move **more** in response to politically aligned information, and it is **statistically significant**

They move **more** in response to politically aligned information, but it is **not statistically significant**

They move **less** in response to politically aligned information, but it is **not statistically significant**

They move **less** in response to politically aligned information, and it is **statistically significant**

What do you expect the coefficient difference ("aligned" - "misaligned") to be (in standard deviations)?

→

Question 2/5

Do you think the movement difference ("aligned" - "misaligned") is greater or smaller for "more extreme" partisans (those with strong party affiliation, and/or more negative emotions towards the opposite party)?

Extreme partisans react **more** to aligned messages compared to misaligned, and the difference between extreme and moderate partisans is **statistically significant**

Extreme partisans react **more** to aligned messages compared to misaligned, but the difference between extreme and moderate partisans is **not statistically significant**

Extreme partisans react **less** to aligned messages compared to misaligned, but the difference between extreme and moderate partisans is **not statistically significant**

Extreme partisans react **less** to aligned messages compared to misaligned, and the difference between extreme and moderate partisans is **statistically significant**

What do you expect the coefficient difference to be (in standard deviations)?
("Extreme partisan difference" - "moderate partisan difference")

→

Question 3/5

Do you think the movement difference (“aligned” - “misaligned”) is greater or smaller for issues that partisans consider more important?

- Important issues have a **larger** difference, and it is **statistically significant**
- Important issues have a **larger** difference, but it is **not statistically significant**
- Important issues have a **smaller** difference, but it is **not statistically significant**
- Important issues have a **smaller** difference, and it is **statistically significant**

What do you expect the coefficient difference to be (in standard deviations)?
("difference for more important issues" - "difference for less important issues")

→

Question 4/5

Do you think the movement difference (“aligned” - “misaligned”) is greater or smaller for issues where the difference regarding initial guesses between partisans is greater?

- Issues with greater differences in opinion have a **larger** movement difference, and it is **statistically significant**.
- Issues with greater differences in opinion have a **larger** movement difference, but it is **not statistically significant**.
- Issues with greater differences in opinion have a **smaller** movement difference, but it is **not statistically significant**.
- Issues with greater differences in opinion have a **smaller** movement difference, and it is **statistically significant**.

What do you expect the coefficient difference to be (in standard deviations)?

→

Question 5/5

In addition to the described measures, we ask survey participants what they think the gaps are between Democrats and Republicans.

Do you think the movement difference ("aligned" - "misaligned") is greater or smaller for issues when the **perceived** difference regarding initial guesses between partisans is exaggerated (compared to underrated)?

Issues with over-perceived differences in opinion have a **larger** movement difference, and it is **statistically significant**.

Issues with over-perceived differences in opinion have a **larger** movement difference, but it is **not statistically significant**.

Issues with over-perceived differences in opinion have a **smaller** movement difference, but it is **not statistically significant**.

Issues with over-perceived differences in opinion have a **smaller** movement difference, and it is **statistically significant**.

What do you expect the coefficient difference to be (in standard deviations)?

→

Short Questionnaire

Please answer the questions below.

Are you a US citizen?

- Yes
- No

If you were to vote in the United States, what would be your political affiliation?

- Independent
- Democrat
- Republican
- None of the above / Prefer not to say

On a scale from 1 (not important) to 7 (very important), how important is your political identity to you?

1 (not important)	2	3	4	5	6	7 (very important)

→

Thank you for your answers!

The survey is over and **your response will be recorded after you click "Finish Survey".**

Finish Survey

References

Agranov, Marina, and Polina Detkova. 2024. "Beliefs of Others: an Experiment." Working paper. National Bureau of Economic Research. [167]

Ahler, Douglas J. 2014. "Self-fulfilling misperceptions of public polarization." *Journal of Politics* 76 (3): 607–20. [167]

Ahler, Douglas J, and Gaurav Sood. 2018. "The parties in our heads: Misperceptions about party composition and their consequences." *Journal of Politics* 80 (3): 964–81. [167]

Barron, Kai. 2021. "Belief updating: does the 'good-news, bad-news' asymmetry extend to purely financial domains?" *Experimental Economics* 24 (1): 31–58. [167]

Bursztyn, Leonardo, and David Y Yang. 2022. "Misperceptions about others." *Annual Review of Economics* 14 (1): 425–52. [167]

Chopra, Felix, Ingar Haaland, and Christopher Roth. 2024. "The demand for news: Accuracy concerns versus belief confirmation motives." *Economic Journal* 134 (661): 1806–34. [166]

Coutts, Alexander. 2019. "Good news and bad news are still news: Experimental evidence on belief updating." *Experimental Economics* 22 (2): 369–95. [167]

Eil, David, and Justin M Rao. 2011. "The good news-bad news effect: asymmetric processing of objective information about yourself." *American Economic Journal: Microeconomics* 3 (2): 114–38. [167]

Flynn, Daniel J, Brendan Nyhan, and Jason Reifler. 2017. "The nature and origins of misperceptions: Understanding false and unsupported beliefs about politics." *Political psychology* 38: 127–50. [166]

Garrett, R Kelly. 2009. "Echo chambers online?: Politically motivated selective exposure among Internet news users." *Journal of computer-mediated communication* 14 (2): 265–85. [166]

Gentzkow, Matthew, and Jesse M Shapiro. 2011. "Ideological segregation online and offline." *Quarterly Journal of Economics* 126 (4): 1799–839. [166]

Gödker, Katrin, Peiran Jiao, and Paul Smeets. 2025. "Investor memory." *Review of Financial Studies* 38 (6): 1595–640. [167]

Gorodzeisky, Anastasia, and Moshe Semyonov. 2020. "Perceptions and misperceptions: Actual size, perceived size and opposition to immigration in European societies." *Journal of Ethnic and Migration Studies* 46 (3): 612–30. [166]

Gotthard-Real, Alexander. 2017. "Desirability and information processing: An experimental study." *Economics Letters* 152: 96–99. [167]

Hartzmark, Samuel M, Samuel D Hirshman, and Alex Imas. 2021. "Ownership, learning, and beliefs." *Quarterly journal of economics* 136 (3): 1665–717. [173]

Hossain, Tanjim, and Ryo Okui. 2013. "The binarized scoring rule." *Review of Economic Studies* 80 (3): 984–1001. [168]

Iyengar, Shanto, and Kyu S Hahn. 2009. "Red media, blue media: Evidence of ideological selectivity in media use." *Journal of communication* 59 (1): 19–39. [166]

Iyengar, Shanto, et al. 2019. "The origins and consequences of affective polarization in the United States." *Annual review of political science* 22 (1): 129–46. [166]

Kahan, Dan M. 2013. "Ideology, motivated reasoning, and cognitive reflection." *Judgment and Decision making* 8 (4): 407–24. [167]

Kahan, Dan M, Hank Jenkins-Smith, and Donald Braman. 2011. "Cultural cognition of scientific consensus." *Journal of risk research* 14 (2): 147–74. [167]

Kuziemko, Ilyana, et al. 2015. "How elastic are preferences for redistribution? Evidence from randomized survey experiments." *American Economic Review* 105 (4): 1478–508. [166]

Lee, Nathan, et al. 2021. "More accurate, but no less polarized: Comparing the factual beliefs of government officials and the public." *British Journal of Political Science* 51 (3): 1315–22. [166]

Lelkes, Yphtach. 2016. "Mass polarization: Manifestations and measurements." *Public Opinion Quarterly* 80 (S1): 392–410. [166]

Lodge, Milton, and Charles S Taber. 2013. *The rationalizing voter.* Cambridge University Press. [167]

Möbius, Markus M, et al. 2022. "Managing self-confidence: Theory and experimental evidence." *Management Science* 68 (11): 7793–817. [167]

Nyhan, Brendan. 2020. "Facts and myths about misperceptions." *Journal of Economic Perspectives* 34 (3): 220–36. [166]

Peterson, Erik, Sharad Goel, and Shanto Iyengar. 2021. "Partisan selective exposure in online news consumption: Evidence from the 2016 presidential campaign." *Political science research and methods* 9 (2): 242–58. [166]

Rekker, Roderik, and Eelco Harteveld. 2022. "Understanding factual belief polarization: the role of trust, political sophistication, and affective polarization." *Acta Politica,* 1. [166]

Sides, John, and Jack Citrin. 2007. "European opinion about immigration: The role of identities, interests and information." *British journal of political science* 37 (3): 477–504. [166]

Taber, Charles S, and Milton Lodge. 2006. "Motivated skepticism in the evaluation of political beliefs." *American journal of political science* 50 (3): 755–69. [167]

Thaler, Michael. 2024. "The fake news effect: Experimentally identifying motivated reasoning using trust in news." *American Economic Journal: Microeconomics* 16 (2): 1–38. [167]

Westfall, Jacob, et al. 2015. "Perceiving political polarization in the United States: Party identity strength and attitude extremity exacerbate the perceived partisan divide." *Perspectives on psychological science* 10 (2): 145–58. [167]

Zimmermann, Florian. 2020. "The dynamics of motivated beliefs." *American Economic Review* 110 (2): 337–63. [167]

Chapter 4

Polarization and the 2024 U.S. Presidential Election*

Joint with Florian Zimmermann

4.1 Introduction

Political polarization has become a defining characteristic of the political landscape in most western countries, most notably perhaps the US. People on different sides of the political spectrum differ not only in their policy attitudes but also in factual beliefs and increasingly hold negative views about people with whom they disagree (Boxell, Gentzkow, and Shapiro (Forthcoming), Iyengar et al. (2019), Mason (2018), Klein (2020), Alesina, Miano, and Stantcheva (2020), Allcott, Boxell, et al. (2020)). Political polarization has been argued to undermine government efficiency, cause a higher tolerance for violence as a means to "win" political battles and generally leads to a mindset of us versus them, rendering compromise and collaboration difficult (Finkel et al. (2020), Dimant (2024), Sunstein (2017)).

In this chapter, we study the determinants of political polarization, focusing on the role of elections. Specifically, we investigate whether political polarization between Democrats and Republicans increased in response to the outcome of the 2024 U.S. presidential election. We conducted two representative surveys—one shortly before and one shortly after the election—measuring different markers of polarization using novel quantitative measures as well as open text data.[1] We find that ideological as well as factual polarization increased over the course of the election.

* Funding by the Deutsche Forschungsgemeinschaft (DFG) through CRC TR 224 (Project A01) and under Germany's Excellence Strategy - EXC 2126/1 - 390838866 is gratefully acknowledged. Proofreading assistance for this chapter was provided by the LLM ChatGPT; the research design, analyses, and all arguments remain solely the work of the authors.

1. See Bar-On et al. (2024) for an insightful discussion of the challenges related to the measurement of political polarization.

We also show that fewer Republicans but more Democrats mention politics as a key societal challenge in an open text format. In general, we find that our results are primarily driven by Democrats. Moreover, within the Democratic group, effects tend to be stronger among those most disappointed by the election outcome.

Our empirical approach is descriptive and relies on an event study design where we ran two representative surveys - one shortly before and one shortly after the election - in a between-subjects framework. We measure whether respondents self-identify as Democrat, Republican or Independent, where the latter category serves as a neutral benchmark. We measure ideological polarization in a quantitative way by eliciting respondents' preferences for various categories of government spending as well as their willingness to donate to combat climate change. Factual polarization is measured through agreement with politicized factual statements as well as stock market return expectations. The latter provides us with a quantitative, economically meaningful and forward-looking measure, while at the same time being factual in nature. Finally, we measure the extent to which people worry about the current state of politics or the political system by assessing whether respondents mention politics in their answers to open-ended questions about key challenges the country or they themselves currently face. The advantage of such open text elicitation is that it allows the measurement of thoughts and concerns in an un-primed way (see, e.g., Andre et al. (2024), Stantcheva (2021), Bauer et al. (2024), Graeber, Roth, and Zimmermann (2024)).

Our findings indicate that polarization along all three dimensions increased in response to the election outcome. Regarding ideological polarization, Democrats shifted their desired government spending toward party-aligned categories after the election, while Republicans' spending preferences remained largely unchanged. Additionally, Democrats showed greater willingness to donate to combat climate change, whereas Republicans' willingness declined. In terms of factual polarization, supporters of both parties aligned their beliefs more closely with partisan views after the election. Furthermore, Democrats adopted a more pessimistic outlook on the stock market following the election, while Republicans became more optimistic. Finally, Democrats more frequently named politics as a key challenge for the country or themselves in the wake of the election, while Republicans showed the opposite tendency. Notably, on all these markers of polarization, Independents show no detectable movement in response to the election outcome.

We further document that the increase in polarization is primarily driven by Democrats, likely due to their disappointment over the election outcome. In line with this, among Democrats, shifts tend to be most pronounced for those who had greater pre-election confidence in a Democrat victory. While our results indicate that the pre-election views of high confidence Democrats not necessarily more closely align with Democratic Party values, they arguably experienced greater disappointment upon learning the result of the election.

This chapter contributes to the growing body of literature that investigates political polarization and its determinants. Existing work has for instance analyzed the effects of social media (e.g., Allcott, Braghieri, et al. (2020)), identity and motivated reasoning (e.g., Iyengar and Westwood (2015), Mason (2018), Barron, Becker, and Huck (2025), the interaction of moral intuitions and ex-post rationalizations (e.g., Le Yaouanq, Schwardmann, and Van der Weele (2025)), social interactions (e.g., Braghieri, Tripodi, and Schwardmann (2025)) and trade exposure (e.g., Autor et al. (2020)) on different aspects of political polarization. We add to this literature by (i) focusing on the role of elections, (ii) by measuring a broad and novel set of markers of polarization (including quantitative measures and open text responses) and (iii) by highlighting that the roots of polarization might in parts be emotional and contingent on political events, filtered through personal expectations.[2]

The rest the chapter is structured as follows. In Section 2, we present the design of our study. Section 3 presents results on the overall increase in polarization, while Section 4 provides suggestive evidence that this increase is mainly driven by Democrats with high levels of pre-election confidence. Section 5 then concludes.

4.2 Study Design

Our survey design to analyze changes in polarization between Republicans and Democrats surrounding the 2024 U.S. presidential election incorporates the following key elements: (i) a rich set of measures of different markers of polarization; (ii) representative samples featuring both Democrats and Republicans as well as a benchmark group of Independents to control for general time trends and related factors; and (iii) control over two distinct time points, one immediately before the election and one immediately after the election, when the surveys are conducted, ideally fixing the day of the week.

2. Relatedly, a literature in political science examines how the temporal proximity of elections influences political polarization. Most of this research has focused on documenting that political campaigning running up to elections can increase polarization (Sood and Iyengar (2016); Bassan-Nygate and Weiss (2022); Hansen and Kosiara-Pedersen (2017)). Michelitch and Utych (2018) as well as Singh and Thornton (2019) analyze data from multiple countries and find that partisan identification with the preferred political party increases in the immediate aftermath of elections and declines afterward. Fasching et al. (2024) find that polarization among the U.S. electorate increases in the approach of the 2022 midterm elections and remains high even after the election. We contribute to this literature by using an event study design with two representative surveys to look at the immediate impact of the 2024 U.S. presidential election, employing a broad set of novel quantitative measures of polarization to document an increase in factual and ideological. Furthermore, our measures of pre-election confidence allow us to delve into the drivers of the increase in polarization, highlighting the role of disappointment about the election outcome.

4.2.1 Design

We designed our study with these goals in mind. We administered the same set of questions in a between-subjects design, where one survey was conducted shortly before and another conducted shortly after the election, fixing the day of the week.[3] Screenshots of the instructions are provided in Appendix 4.A.3.

At the beginning of the survey, participants were asked to name the top three challenges the country is currently facing, as well as the top three challenges they are personally facing. The question format was open-ended, allowing us to measure which challenges are top of mind in an un-primed way. Using a list of keywords, reproduced in Appendix 4.A.2, we construct a dummy variable indicating whether politics is mentioned among the top challenges. We interpret this dummy variable as a proxy for discontent with and concern about the current political environment.

Next, to get at quantitative measures of ideological polarization, we asked participants about their preferences regarding government spending (see Enke, Rodríguez-Padilla, and Zimmermann (2023) for a similar approach). Specifically, we asked them to indicate how much revenue the government should collect per American, on average, for each of the following seven categories: Police and Law Enforcement, Foreign Aid, Universal Healthcare, Fight against Climate Change, Military and Counterintelligence, Welfare Payments, and Effective Border Control. This allows us to construct an ideological spending index by subtracting the average spending on Democratic categories (*Foreign Aid*, *Universal Healthcare*, *Fight against Climate Change*, and *Welfare Payments*) from the average spending on Republican categories (*Police and Law Enforcement*, *Military and Counterintelligence* and *Effective Border Control*).

Afterwards, to obtain a second measure of ideological polarization, participants were asked whether they would be willing to make a hypothetical monthly contribution of 1% of their household income to combat global warming.

Moving to measures of factual polarization, we then elicited participants' stock market expectations, using their expected return on the MSCI World Index over the next 12 months as a proxy. This has the advantage of being quantitative and economically meaningful while still being factual. In addition, the forward-looking nature allows us to capture potential differences in economic optimism. In addition, we elicited participants' beliefs regarding a set of politicized factual statements on a scale from 0 to 100%. The first statement claims 'On average, U.S. immigrants have lower levels of education than the U.S.-born population'. The second statement is 'Women who have had an abortion experience more psychological distress

3. We chose a between-subjects design because this removes concerns about consistency bias in surveys (see, e.g., Falk and Zimmermann (2013)). Due to the short time distance between the two survey waves (1 week), a panel structure would have been highly susceptible to consistency bias, likely downward biasing observed effects on polarization.

than those who have had a miscarriage.' The third statement states 'Discrimination against women is the primary reason why women earn less than men.' We classify the first and second statement as Republican-coded and the third statement as Democrat-coded. We construct a statement index by subtracting the agreement with the Democrat-coded statement from the average agreement with the Republican-coded statements.

Finally, we collected a set of demographic variables, participants' voting intentions in the presidential election, their prediction of the likely winner, and their confidence in this assessment.

4.2.2 Procedures

The study was implemented using Qualtrics and administered via Prolific. 600 participants completed the pre-election study on November 4, 2024, one day before the Presidential Election. The median completion time was 14 minutes. An additional 600 participants completed the post-election study one week later, on November 11, 2024. We chose this timeline as it allows us to fix the day of the week. In addition, we wanted to ensure that enough time had passed since the election such that the election outcome was likely to be clear. The median completion time for the post-election study was 13 minutes. Respondents received a payment of 2.75 USD for participating. Both samples are broadly representative across age, education, gender, income, ethnicity, and political affiliation. There are no significant differences in observable characteristics between the two studies (see Table 4.A.1). Our samples hence consist of both Democrats and Republicans, as well as Independents, which serve as a natural benchmark group to capture potential general time trends.

4.3 Results: Changes in Polarization

Ideological Polarization. Ideological polarization reflects divergences in political attitudes and policy preferences between opposing political groups. Our study employs two measures to quantify this form of polarization. Our first measure of ideological polarization makes use of the desired government spending index.

Figure 4.1 presents the average spending index by party affiliation before and after the election. Before the election, Republicans and Democrats exhibit significant differences in policy preferences, with both groups allocating more government spending to their respective categories, with Independents being in-between. After the election, this gap widens as supporters of both parties show increased spending for party-congruent categories. Reassuringly, the spending preferences of Independents remain unchanged after the election. The statistically significant difference-in-differences estimate in column (1) of Table 4.1 confirms the increase in ideological polarization between Republicans and Democrats, as measured by the spending index.

Table 4.1. Changes in Polarization between Democrats and Republicans

	Ideological Polarization		*Factual Polarization*		*Concern About Politics*
Dependent variable:	Spending Index	Climate Donation	Statement Index	MSCI Expectations	Mentions Politics
	(1)	(2)	(3)	(4)	(5)
Constant	-152.775***	0.657***	-31.386***	7.723***	0.154***
	(42.240)	(0.036)	(2.858)	(1.056)	(0.027)
After Election	-268.306**	0.119**	-8.080**	-3.168*	0.053
	(105.316)	(0.048)	(3.863)	(1.670)	(0.041)
Republican	441.068***	-0.340***	40.968***	1.134	-0.057
	(76.666)	(0.051)	(4.148)	(1.549)	(0.036)
After El. × Rep.	316.082**	-0.203***	13.599**	7.128***	-0.101**
	(142.191)	(0.069)	(5.628)	(2.491)	(0.050)
R^2	0.104	0.205	0.305	0.032	0.031
Observations	676	676	676	676	676

The table presents OLS regression results for various measures of polarization. The sample includes only Republican and Democrat participants. Standard errors are reported in parentheses. Stars highlight significant differences from 0, with * for $p < 0.10$, ** for $p < 0.05$, and *** for $p < 0.01$. *After Election* is a dummy for the survey that was conducted after the election. *Republican* is a dummy variable that is 1 if the individual affiliates with the Republican party. *Spending Index* is the average desired spending for Republican categories minus the average spending for Democrat categories (in USD), based on spending indications that are winsorized at the top 2.5%. *Climate Donation* is a dummy for whether individuals indicate they are willing to donate one percent of their income for combating Climate Change. *Statement Index* is the average agreement (from 0 to 100) to Republican-coded statements minus the agreement to the Democrat-coded statement. *MSCI Expectations* are the expected returns of the MSCI World index over the next 12 months, winsorized at the top and bottom 1%. *Mentions Politics* is a dummy for whether political topics were mentioned as top challenges the country or individual is currently facing.

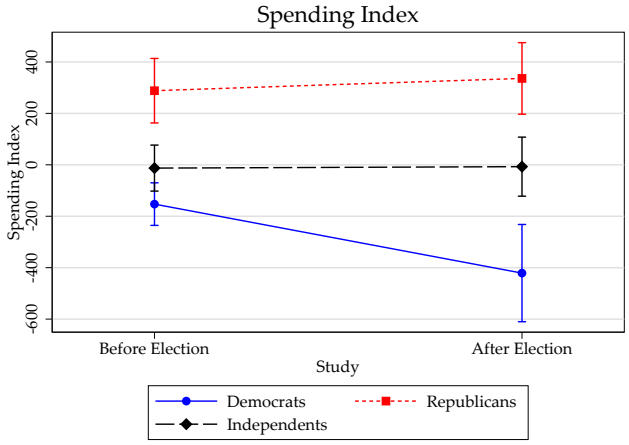

Figure 4.1. Average spending index by party affiliation, before and after the election. *The spending index is defined as the average spending for Republican categories minus Democrat categories.*

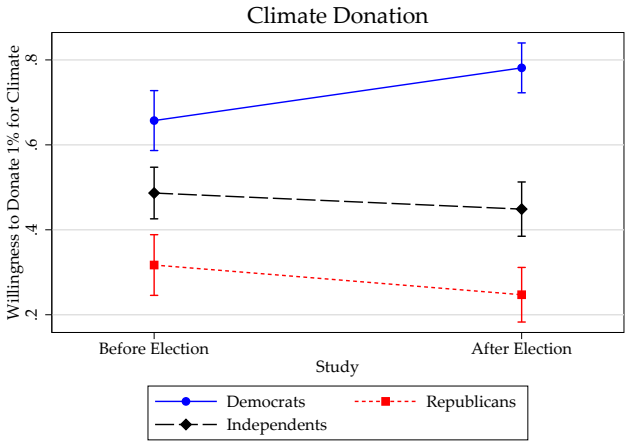

Figure 4.2. Average willingness to donate to combat climate change by party affiliation, before and after the election. *The measure is a dummy that is 1 if people indicated they would be willing to donate one percent of their income to combat climate change.*

Our second measure is more behavioral and assesses participants' hypothetical willingness to donate 1% of their income to combat climate change. Figure 4.2 presents the average willingness to donate, by party affiliation, before and after the election. Before the election, Democrats exhibit a higher willingness to donate than Republicans. This difference widens after the election, as Democrats' willingness increases while Republicans' willingness declines. This pattern is further supported by a statistically significant difference-in-differences estimate in column (2) of Table 4.1. The donation preferences of Independents remain largely unchanged. This indicates that the increase in donations by Democrats is not driven by them merely wanting to compensate for an expected decrease in government climate action, as one would expect Independents to show a similar response. Also, and perhaps interestingly, the Democrat response is inconsistent with a conditional cooperation, according to which an expected decrease in government climate action should lead to lower donations.

Overall, our findings provide clear evidence of increased ideological polarization between Republicans and Democrats following the election.

Factual Polarization. Factual polarization refers to the degree to which opposing groups hold divergent views about things where there exists an objectively correct answer.[4] Our study incorporates two measures of factual polarization.

Our first measure of factual polarization is the statement index of participants' responses to the factual statement questions. Figure 4.3 plots the average statement index by party affiliation, before and after the election. There is an increase in the

4. We might not (or not yet) know this answer, of course.

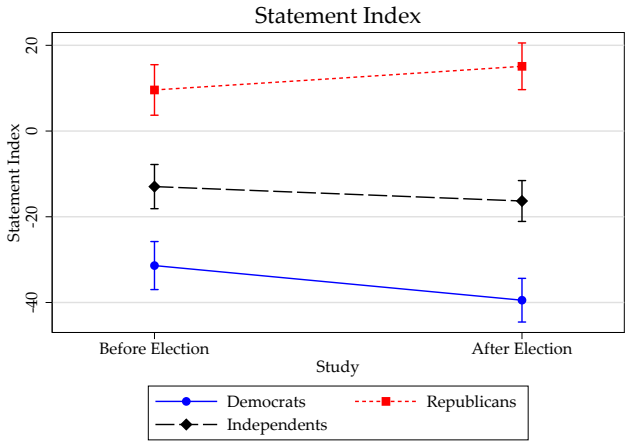

Figure 4.3. Average statement index by party affiliation, before and after the election. *The statement index is defined as the average agreement to Republican statements minus agreement to the Democrat statement.*

index for Republicans and a decrease for Democrats. This implies that members of both parties hold increasingly party-aligned factual beliefs after the election. The statistically significant difference-in-differences coefficient in column (3) of Table 4.1 confirms this pattern. There is no significant movement for Independents.

Our second measure uses participants' stock market expectations, specifically their expected return on the MSCI World Index over the next 12 months. Figure 4.4 shows that expectations were homogeneous before the election but diverged afterward. Following the election, Republicans significantly raised their stock market expectations, while Democrats significantly lowered theirs. Once again, Independents show no detectable change in return expectations. Column (4) of Table 4.1 confirms that the election significantly increased polarization in stock market expectations between Republicans and Democrats. This result as well as the results on the statement index is difficult to reconcile with Bayesian belief or expectation formation, even when allowing for non-common priors.

In summary, our findings indicate an increase in factual polarization following the election.

Concern about Politics. Recall that we started each survey with open-ended questions where participants could name three or more key challenges for the country as well as challenges for the themselves. This allows us to measure which challenges are top of mind in an un-primed way. We use these open text questions to create a measure of concern and discontent about the current political landscape, namely whether political topics are mentioned at least once in the challenges named by a participant.

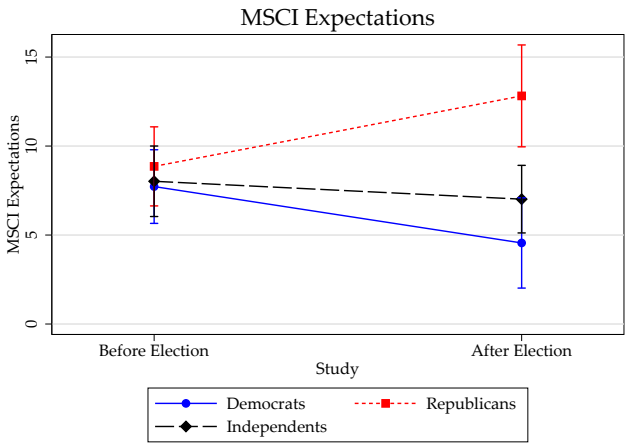

Figure 4.4. Average MSCI World expectations by party affiliation, before and after the election. *Participants were asked to state their return expectations for the MSCI World Index over the next 12 months.*

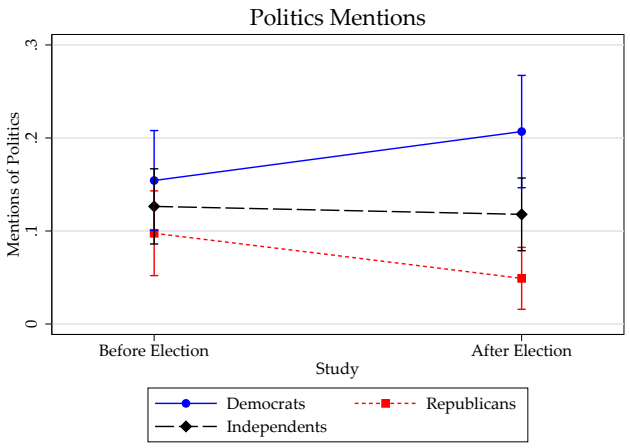

Figure 4.5. Average mentions of political topics among top challenges by political affiliation, before and after the election. *The dummy variable captures whether politics is mentioned as one of the top challenges the country or participant currently face.*

Figure 4.5 presents the resulting dummy variable by political affiliation, before and after the election. After the election, more Democrats and fewer Republicans mention politics as a key challenge, while the rate among Independents remains relatively unchanged. The statistically significant difference-in-differences estimate in column (5) of Table 4.1 confirms increased polarization regarding concerns about the political landscape.

4.4 Who Drives the Increase in Polarization?

Democrats and Republicans may exhibit asymmetric reactions to the election outcome. Table 4.2 summarizes the previous section's results by reporting the mean difference in outcomes for Democrats and Republicans before and after the election. For all outcomes except MSCI expectations, Democrats exhibit a stronger directional response than Republicans. This pattern may reflect disappointment about the election outcome and the (in other decision contexts) well-documented tendency for losses to loom larger than gains.

Table 4.2. Mean Changes for Democrats and Republicans

	Ideological Polarization		*Factual Polarization*		*Concern About Politics*
Dependent variable:	Spending Index	Climate Donation	Statement Index	MSCI Expectations	Mentions Politics
Mean Difference for Democrats	-268.306	0.119	-8.080	-3.168	0.053
Mean Difference for Republicans	47.777	-0.084	5.519	3.960	-0.048

The table presents the mean differences for Democrats and Republicans for various measures, comparing average responses among the members of each party before and after the election. *Spending Index* is the average desired spending for Republican categories minus the average spending for Democrat categories (in USD), based on spending indications that are winsorized at the top 2.5%. *Climate Donation* is a dummy for whether individuals indicate they are willing to donate one percent of their income for combating Climate Change. *Statement Index* is the average agreement (from 0 to 100) to Republican-coded statements minus the agreement to the Democrat-coded statement. *MSCI Expectations* are the expected returns of the MSCI World index over the next 12 months, winsorized at the top and bottom 1%. *Mentions Politics* is a dummy for whether political topics were mentioned as top challenges the country or individual is currently facing.

To further zoom in on the drivers of increased polarization, we divide Democrats into two groups based on their confidence in a Democrat win, under the assumption that these groups were differentially surprised by the outcome of the election.

In the pre-election survey, we asked participants whom they expected to win the election and how confident they were in their prediction on a scale from 0 to 100. In the post-election survey, we asked participants to assess these questions in hindsight. We restrict our analysis to Democrats who expected a Democrat win and perform a median split based on their confidence levels in both the pre- and post-election samples.[5] The underlying assumption is that supporters with greater confidence in

5. One might worry about the validity of post-election respondents reporting what they expected before the election. Reassuringly, Table 4.A.2 shows no meaningful demographic differences between pre- and post-election respondents who believed Harris would win. Furthermore, confidence in a Harris victory, elicited on a scale from 0 to 100, averaged 72.73 (SE = 0.92) among pre-election

a Democrat victory were more surprised (disappointed) by the election outcome. We point out that individuals with greater confidence in a Democrat win are not necessarily those whose views most strongly align with Democratic Party values (in fact, our results reveal that this is not the case, see Figure 4.A.1). Therefore, we do not hypothesize that their responses are consistently more reflective of party values. Instead, our analysis focuses on changes between the pre- and post-election periods.

Table 4.3. Regression Results for High and Low Confidence Democrats

	Ideological Polarization		*Factual Polarization*		*Concern About Politics*
Dependent variable:	Spending Index	Climate Donation	Statement Index	MSCI Expectations	Mentions Politics
	(1)	(2)	(3)	(4)	(5)
Constant	-236.419***	0.743***	-39.351***	4.518***	0.189***
	(81.256)	(0.051)	(4.141)	(1.469)	(0.046)
After Election	2.468	0.053	2.360	0.903	0.014
	(172.485)	(0.073)	(6.040)	(2.513)	(0.070)
High Confidence	125.462	-0.141*	7.621	6.832***	-0.061
	(96.177)	(0.076)	(5.874)	(2.211)	(0.060)
After El. × High Conf.	-457.535*	0.107	-18.563**	-22.791	0.068
	(240.565)	(0.105)	(8.135)	(13.963)	(0.093)
R^2	0.031	0.028	0.030	0.018	0.007
Observations	287	287	287	287	287

The table presents OLS regression results for various measures of polarization. The sample only includes Democrats who indicated they thought Kamala Harris would win the election. Standard errors are reported in parentheses. Stars highlight significant differences from 0, with * for $p < 0.10$, ** for $p < 0.05$, and *** for $p < 0.01$. *After Election* is a dummy for the survey that was conducted after the election. *High Confidence* is a dummy variable that is 1 if the individual stated at least median confidence for Harris winning the election, with the median evaluated separately for both studies. *Spending Index* is the average desired spending for Republican categories minus the average spending for Democrat categories (in USD), based on spending indications that are winsorized at the top 2.5%. *Climate Donation* is a dummy for whether individuals indicate they are willing to donate one percent of their income for combating Climate Change. *Statement Index* is the average agreement (from 0 to 100) to Republican-coded statements minus the agreement to the Democrat-coded statement. *MSCI Expectations* are the expected returns of the MSCI World index over the next 12 months, winsorized at the top and bottom 1%. *Mentions Politics* is a dummy for whether political topics were mentioned as top challenges the country or individual is currently facing.

Figure 4.A.1 displays the results for all five outcome variables. Table 4.3 reports the results of the difference-in-differences regression analysis. While the statistical power of this analysis is of course limited, we do see suggestive evidence for more pronounced (party-aligned) changes among high-confidence Democrats. Specifically, for both the spending and statement indices, we observe particularly strong, party-aligned changes among high-confidence Democrats. For all other variables, the difference-in-differences estimates point in the same direction but do not reach statistical significance.

respondents and 65.29 (SE = 1.15) among post-election respondents. Hence, while some downward shift is apparent, the median split remains valid under the assumption that there were no systematic changes in the internal confidence rankings of pre- and post-election respondents.

4.5 Conclusion

This chapter provides empirical evidence that the 2024 U.S. presidential election led to an increase in political polarization along ideological and factual lines. By employing an event study design with two representative surveys conducted before and after the election, we systematically document shifts in government spending preferences, willingness to donate to climate-related causes, alignment with partisan factual beliefs, economic expectations and concerns about politics. The results indicate that elections serve as a catalyst for polarization, reinforcing partisan divides rather than bridging them.

A key finding is that the observed polarization is asymmetric, with Democrats exhibiting more pronounced shifts in their beliefs and concerns following the election. This pattern suggests that emotional responses to election outcomes, particularly disappointment among those who anticipated a different result, play a crucial role in shaping political attitudes. Indeed, Democrats who showed high pre-election confidence in a Democrat win show the greatest reaction to the election. Moreover, our findings reveal that Independents do not exhibit comparable shifts in their political attitudes, reinforcing the notion that entrenched partisan identities contribute significantly to post-election polarization.

These results have important implications for both policymakers and scholars studying democratic stability. Increased polarization may hinder bipartisan cooperation, further entrench negative partisanship, and weaken democratic norms. Understanding the mechanisms through which elections exacerbate political divisions is essential for developing strategies to mitigate their effects, such as fostering cross-party dialogue, promoting media literacy to counter factual polarization, and designing electoral processes that reduce zero-sum perceptions of political competition.

Appendix 4.A Appendix

4.A.1 Additional Tables and Figures

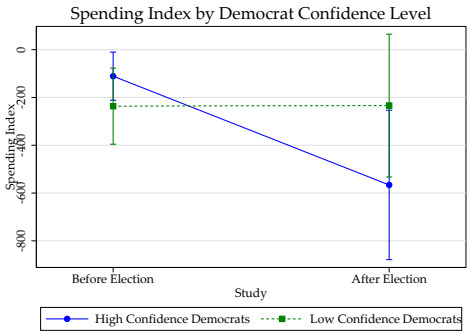

(a) Spending Index

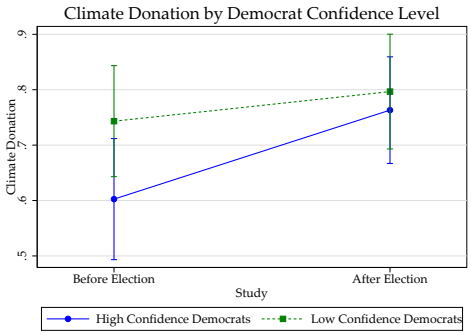

(b) Climate Donation

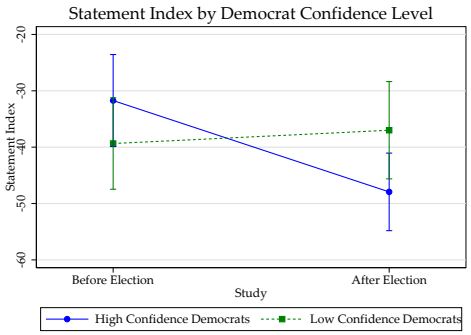

(c) Statement Index

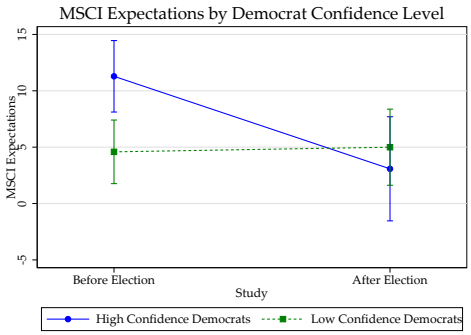

(d) MSCI Expectations

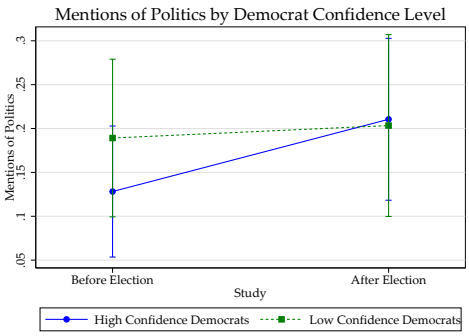

(e) Mentions Politics

Figure 4.A.1. Confidence split for Democrat responses to election. *The plots show the change in outcomes for Democrats who thought Kamala Harris would win the election. The plots are based on a median split with respect to the confidence that Kamala Harris would be the winner, separately for each study.*

Table 4.A.1. Summary Statistics with Differences Between Pre- and Post-Election Survey

Variable	Pre-Election Survey	Post-Election Survey	Difference
Education			
High school diploma, equivalent, or less	0.365	0.381	0.016
	(0.019)	(0.020)	(0.027)
Some college without degree	0.217	0.208	-0.009
	(0.017)	(0.017)	(0.024)
Bachelor's or Associate Degree	0.292	0.285	-0.007
	(0.018)	(0.018)	(0.025)
Graduate or professional degree	0.127	0.125	-0.002
	(0.014)	(0.014)	(0.020)
Age			
18-24	0.117	0.105	-0.012
	(0.013)	(0.012)	(0.018)
25-34	0.160	0.158	-0.002
	(0.015)	(0.015)	(0.021)
35-44	0.162	0.160	-0.002
	(0.015)	(0.015)	(0.021)
45-54	0.158	0.152	-0.006
	(0.015)	(0.015)	(0.021)
55+	0.403	0.425	0.022
	(0.020)	(0.020)	(0.028)
Income			
Less than 2000	0.421	0.427	0.006
	(0.020)	(0.020)	(0.028)
2000-3999	0.238	0.237	-0.001
	(0.018)	(0.018)	(0.025)
4000-6999	0.170	0.165	-0.005
	(0.015)	(0.015)	(0.021)
7000-9999	0.088	0.092	0.004
	(0.011)	(0.011)	(0.016)
10000+	0.081	0.080	-0.001
	(0.010)	(0.010)	(0.014)
Ethnicity			
Asian, Pacific Islander, American Indian or Alaska Native	0.075	0.075	0.000
	(0.011)	(0.011)	(0.016)
Hispanic or Latino	0.130	0.135	0.005
	(0.014)	(0.015)	(0.021)
Black or African American	0.102	0.102	0.000
	(0.012)	(0.012)	(0.017)
White	0.693	0.688	-0.005
	(0.019)	(0.019)	(0.027)
Political Affiliation			
Independent or Other	0.435	0.438	0.003
	(0.020)	(0.020)	(0.028)
Republican	0.273	0.270	-0.003
	(0.018)	(0.018)	(0.025)
Democrat	0.292	0.292	0.000
	(0.018)	(0.018)	(0.025)
Gender			
Female	0.498	0.498	0.000
	(0.020)	(0.020)	(0.028)
Male	0.494	0.495	0.001
	(0.020)	(0.020)	(0.028)
Other	0.008	0.007	-0.001
	(0.004)	(0.004)	(0.006)
N	600	600	

The table presents summary statistics. Standard errors for proportions are reported in parentheses. Stars highlight significant differences between pre- and post-election surveys, with * for $p < 0.10$, ** for $p < 0.05$, and *** for $p < 0.01$.

Table 4.A.2. Summary Statistics for Democrats Who Indicated They Expected a Harris Victory

Variable	Pre-Election Survey	Post-Election Survey	Difference
Education			
High school diploma, equivalent, or less	0.303	0.356	0.053
	(0.037)	(0.041)	(0.056)
Some college without degree	0.230	0.170	-0.060
	(0.034)	(0.032)	(0.047)
Bachelor's or Associate Degree	0.336	0.311	-0.024
	(0.038)	(0.040)	(0.055)
Graduate or professional degree	0.132	0.163	0.031
	(0.027)	(0.032)	(0.042)
Age			
18–24	0.125	0.096	-0.029
	(0.027)	(0.025)	(0.037)
25–34	0.099	0.104	0.005
	(0.024)	(0.026)	(0.036)
35–44	0.125	0.148	0.023
	(0.027)	(0.031)	(0.041)
45–54	0.132	0.163	0.031
	(0.027)	(0.032)	(0.042)
55+	0.520	0.489	-0.031
	(0.041)	(0.043)	(0.059)
Income			
Less than 2,000	0.454	0.452	-0.002
	(0.040)	(0.043)	(0.059)
2,000–3,999	0.217	0.252	0.035
	(0.034)	(0.038)	(0.051)
4,000–6,999	0.164	0.141	-0.023
	(0.030)	(0.031)	(0.043)
7,000–9,999	0.079	0.096	0.017
	(0.022)	(0.026)	(0.034)
10,000+	0.086	0.059	-0.027
	(0.023)	(0.020)	(0.030)
Ethnicity			
Asian, Pacific Islander, American Indian or Alaska Native	0.066	0.089	0.023
	(0.020)	(0.025)	(0.032)
Black or African American	0.125	0.089	-0.036
	(0.027)	(0.025)	(0.037)
Hispanic or Latino	0.151	0.178	0.027
	(0.029)	(0.032)	(0.043)
White	0.658	0.644	-0.014
	(0.039)	(0.041)	(0.057)
Political Affiliation			
Democrat	1.000	1.000	0.000
	(0.000)	(0.000)	(0.000)
Gender			
Female	0.546	0.526	-0.020
	(0.040)	(0.043)	(0.059)
Male	0.454	0.467	0.013
	(0.040)	(0.043)	(0.059)
Other	0.000	0.007	0.007
	(0.000)	(0.007)	(0.007)
N	152	135	

The table presents summary statistics for Democrats who indicated they believed Kamala Harris would win the election. Standard errors for proportions are reported in parentheses. Stars highlight significant differences between pre- and post-election surveys, with * for $p < 0.10$, ** for $p < 0.05$, and *** for $p < 0.01$.

4.A.2 Coding Manual for *Mentions Politics*

A participant is coded as mentioning politics if they mention any of the following in at least one of their responses to the challenge questions:

Alt-right, Authoritarian, Bipartisan, Biden, Blue States, Congress, Constitution, Corruption, Cultural Divide, Culture Wars, Democratic Party, Democrats, Democracy, Dictator, Dictatorship, Divide, Election, Election Fraud, Election Interference, Election Rigging, Electoral College, Exit Polls, Extremism, Extremist, Far-left, Far-right, Fascism, Fascist, Free Speech, Freedom of Speech, Gerrymandering, Harris, House, Ideological Divide, Ideology, Left, Left-wing, Left-wing Extremists, Liberal, Liberal Movement, Libertarian, Nationalism, Neoliberal, Oligarchy, Partisan, Partisan Divide, Patriotism, Polarization, Polarized, Polarized Politics, Political, Political Activism, Political Divide, Political Extremism, Political Ideology, Political Parties, Political Polarization, Polling Data, Polling Stations, Polls, Politics, President, Progressive, Protests, Red States, Reformist, Republican, Republican Party, Republicans, Rhetoric, Rigged, Right, Right-wing, Right-wing Extremists, Senate, Socialism, Socialist, Supreme Court, Turnout, Trump, Unrest, US Politics, Vance, Voter Fraud, Voter ID, Voter Registration, Voter Turnout, Voting Rights, Voting Suppression, Walz, White House.

4.A.3 Instructions

4.A.3.1 Introduction

Welcome!

This study is designed for **computer (PC or Mac) users only** (desktop, laptop, etc.). If you are accessing this study on a smartphone, a tablet or any other non-PC devices, please switch to PC and enter the study again, or return the submission on Prolific.

IMPORTANT: You may have participated in a similar survey before. This is on purpose and does NOT mean you have to return the survey!

Please write **at least 15 words** describing **your opinion** about **daylight savings time**. Whether you are in favor or against daylight savings does not affect your eligibility to participate in this study. However, we ask that you write at least 15 words on your thoughts about this topic.

→

4.A.3.2 Challenges

For the following two questions, please state the 3 answers that are on the top of your mind.

→

What would you say are the biggest challenges people in this country are currently facing?

Please enter the challenges in the boxes below and start with the most important one. Please name at least three. (You can name more if you want to by adding additional boxes at the bottom of this page). There is no need to describe the challenges in detail, it is enough to name them.

Challenge 1:

Challenge 2:

Challenge 3:

Do you want to name a 4th challenge?

- No
- Yes

What would you say are the three biggest challenges you personally are currently facing?

Please enter the challenges in the boxes below and start with most important one. Please name at least three. (You can name more if you want to by adding additional boxes at the bottom of this page). There is no need to describe the challenges in detail, it is enough to name them. **Note: The challenges you personally face can be the same as the challenges this country faces. If so, please list them for both questions.**

Challenge 1:

Challenge 2:

Challenge 3:

Do you want to name a 4th challenge?

No

Yes

4.A.3.3 Spending Questions

In this questionnaire, as in the rest of this survey, **none** of the questions will have right or wrong answers. We are only interested in your personal views and preferences.

We will now ask you about your preferences over the **government budget.** Specifically, we ask you to imagine you have the ability to determine how the government of the United States spends money on several different policy categories.

→

Suppose that you could determine how the United States government spends money on various different categories of the federal budget, such as the military or redistribution.

Specifically, imagine you could decide the average amount of money that the federal government collects per year from each American to spend on each of the seven categories below. For the purposes of this question, you should assume that all dollar amounts collected for a category are spent only on this particular category, without any waste.

How much money would you have the federal government collect on average from each American, in order to spend on each of the following seven categories of expenditure in the federal budget?

To provide a reference, it is estimated that altogether, all levels of government in the United States spend a combined average amount of $2540 per American every year for the purposes of education.

	Amount of money ($) collected on average from each American to spend on category, per year:
Police and Law Enforcement	$
Foreign Aid	$
Universal Healthcare	$
Fight against Climate Change	$
Military and Counterintelligence	$
Welfare Payments	$
Effective Border Control	$

→

4.A.3.4 Climate Donation

The next question is hypothetical.

Would you be willing to contribute 1% of your household income every month to fight global warming? This would mean that you would contribute $1 for every $100 of this income.

Yes

No

→

4.A.3.5 Stock Market Expectations

Next, we are interested in your personal stock markt expectations.

The **MSCI World Index** is a list of stocks from large and medium-sized companies in 23 developed countries. Its goal is to mirror economic development and market conditions in these developed countries.

What would you expect the return of the MSCI World Index to be over the next 12 months? (Please answer in %)

Note: This expected return is the change in value, in percent, that you would expect to **receive over the next 12 months** from investing in a portfolio that holds all stocks listed in the MSCI World Index. It includes both dividends and capital gains/losses.

Please enter your guess (in percent) and specify whether you mean a percent increase or decrease.

[ ] %

Increase

Decrease

→

4.A.3.6 Demographics

Please answer the following questions.

What is your age?

What is your gender?

- Male
- Female
- Other

What best describes your ethnicity?

- White
- Hispanic or Latino
- Black or African American
- Asian, Pacific Islander, American Indian or Alaska Native

What is your personal monthly **pre-tax** income?

- Less than $2000
- $2000 - $3999
- $4000 - $6999
- $7000 - $9999
- $10000 or more

What is the highest level of education you have completed?

- Less than high school diploma
- High school diploma or equivalent
- Some college without degree
- Bachelor's degree (e.g., BA, BS) or Associate Degree (e.g. AA, AS)
- Graduate or professional degree (e.g., MA, MS, MBA, MD, JD, PhD)

Which state are you from?

What best describes your political affiliation?

- Democrat
- Republican
- Independent or Other

4.A.3.7 Election Module

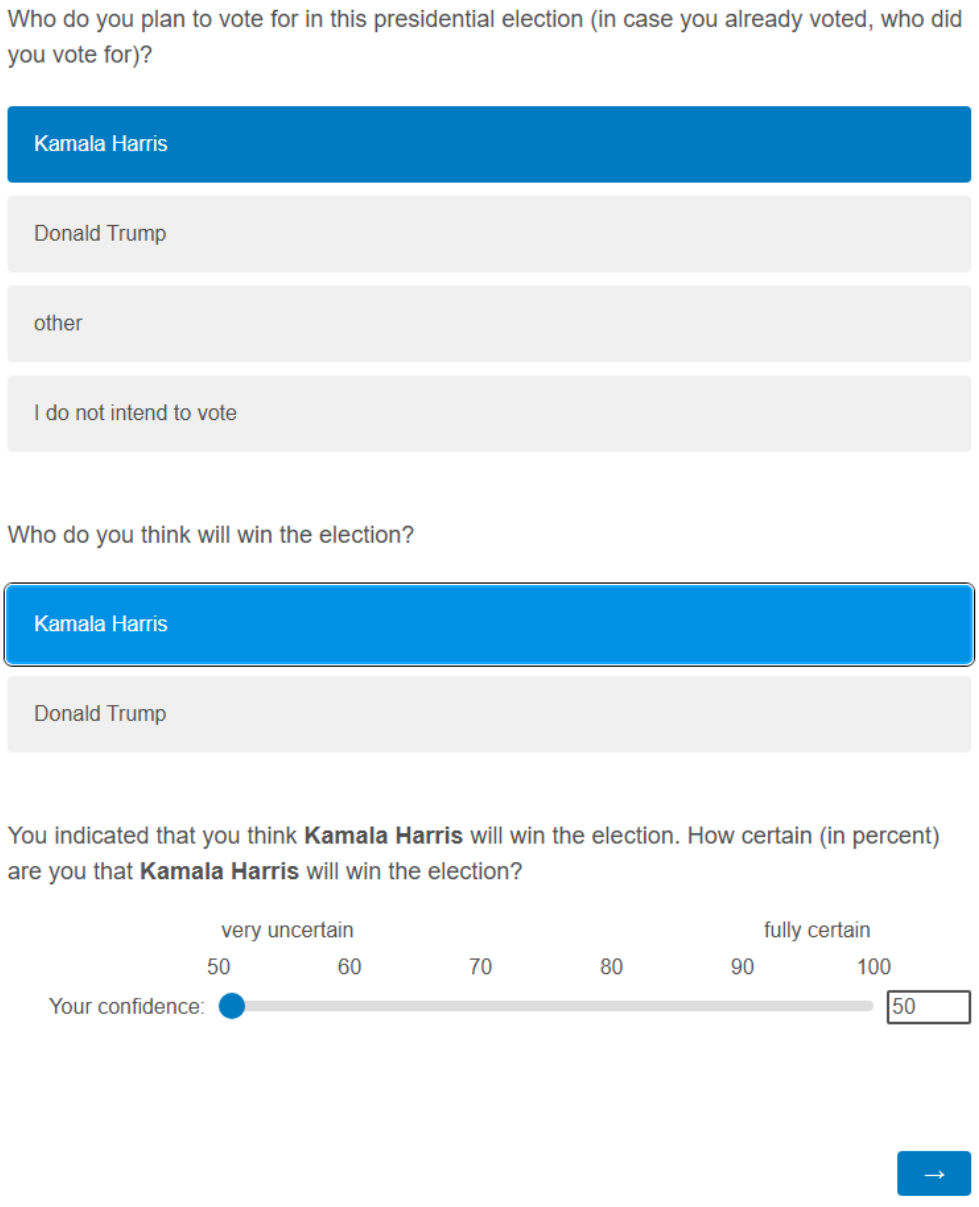

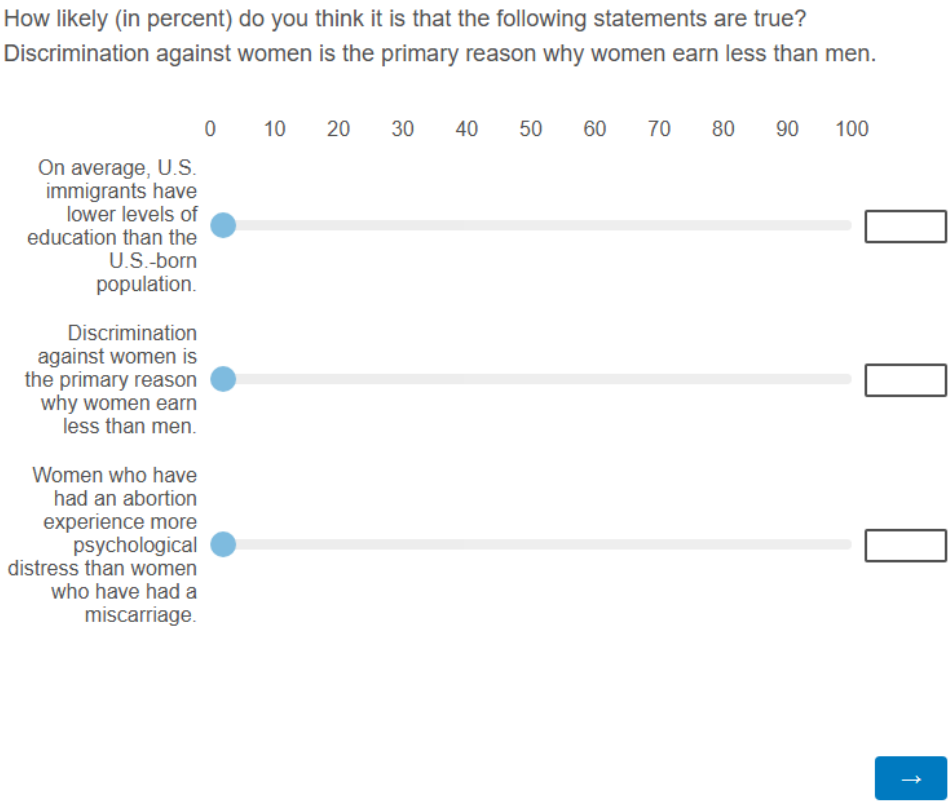
How likely (in percent) do you think it is that the following statements are true?
Discrimination against women is the primary reason why women earn less than men.
0
10
20
30
40
50
60
70
80
90
100
On average, U.S. immigrants have lower levels of education than the U.S.-born population.
Discrimination against women is the primary reason why women earn less than men.
Women who have had an abortion experience more psychological distress than women who have had a miscarriage.
→

References

Alesina, Alberto, Armando Miano, and Stefanie Stantcheva. 2020. "The polarization of reality." In *AEA Papers and Proceedings,* 110: 324–28. American Economic Association 2014 Broadway, Suite 305, Nashville, TN 37203. [245]

Allcott, Hunt, Levi Boxell, et al. 2020. "Polarization and Public Health: Partisan Differences in Social Distancing during the Coronavirus Pandemic." *Journal of Public Economics* 191. [245]

Allcott, Hunt, Luca Braghieri, et al. 2020. "The Welfare Effects of Social Media." *American Economic Review* 110 (3): 629–76. https://doi.org/10.1257/aer.20191372. [247]

Andre, Peter, et al. 2024. "Narratives about the Macroeconomy." Working paper. Working Paper. [246]

Autor, David, et al. 2020. "Importing Political Polarization? The Electoral Consequences of Rising Trade Exposure." *American Economic Review* 110 (10): 3139–83. https://doi.org/10.1257/aer.20180673. [247]

Bar-On, Kati Kish, et al. 2024. "Unraveling polarization: Insights into individual and collective dynamics." *PNAS Nexus* 3 (10): pgae426. https://doi.org/10.1093/pnasnexus/pgae426. [245]

Barron, Kai, Anna Becker, and Steffen Huck. 2025. "Motivated political reasoning: On the emergence of belief-value constellations." *European Economic Review* 172: 104929. https://doi.org/10.1016/j.euroecorev.2024.104929. [247]

Bassan-Nygate, Lotem, and Chagai M Weiss. 2022. "Party competition and cooperation shape affective polarization: evidence from natural and survey experiments in Israel." *Comparative Political Studies* 55 (2): 287–318. [247]

Bauer, Rob, et al. 2024. "Mental Models in Financial Markets: How Do Experts Reason about the Pricing of Climate Risk?" Working paper. Working Paper. [246]

Boxell, Levi, Matthew Gentzkow, and Jesse M. Shapiro. Forthcoming. "Cross-country trends in affective polarization." *Review of Economics and Statistics.* [245]

Braghieri, Luca, Egon Tripodi, and Peter Schwardmann. 2025. "Talking Across the Aisle." Working paper. Working Paper. [247]

Dimant, Eugen. 2024. "Hate Trumps Love: The Impact of Political Polarization on Social Preferences." *Management Science* 70 (1): 1–31. https://doi.org/10.1287/mnsc.2023.4701. [245]

Enke, Benjamin, Ricardo Rodríguez-Padilla, and Florian Zimmermann. 2023. "Moral universalism and the structure of ideology." *Review of Economic Studies* 90 (4): 1934–62. https://doi.org/10.1093/restud/rdad026. [248]

Falk, Armin, and Florian Zimmermann. 2013. "A Taste for Consistency and Survey Response Behavior." *CESifo Economic Studies* 59 (1): 181–93. https://doi.org/10.1093/cesifo/ifs039. [248]

Fasching, Neil, et al. 2024. "Persistent polarization: The unexpected durability of political animosity around US elections." *Science Advances* 10 (36): eadm9198. [247]

Finkel, Eli J., et al. 2020. "Political Sectarianism in America." *Science* 370 (6516): 533–36. https://doi.org/10.1126/science.abe1715. [245]

Graeber, Thomas, Christopher Roth, and Florian Zimmermann. 2024. "Stories, Statistics, and Memory." *Quarterly Journal of Economics* 139 (4): 2181–225. https://doi.org/10.1093/qje/qjad012. [246]

Hansen, Kasper M, and Karina Kosiara-Pedersen. 2017. "How campaigns polarize the electorate: Political polarization as an effect of the minimal effect theory within a multi-party system." *Party Politics* 23 (3): 181–92. [247]

Iyengar, Shanto, and Sean J. Westwood. 2015. "Fear and Loathing across Party Lines: New Evidence on Group Polarization." *American Journal of Political Science* 59 (3): 690–707. https://doi.org/10.1111/ajps.12152. [247]

Iyengar, Shanto, et al. 2019. "The origins and consequences of affective polarization in the United States." *Annual review of political science* 22 (1): 129–46. [245]

Klein, Ezra. 2020. *Why We're Polarized.* Simon & Schuster. [245]

Le Yaouanq, Yves, Peter Schwardmann, and Joël Van der Weele. 2025. "Rationalizations and Political Polarization." Working Paper / Discussion Paper CESifo Working Paper No. 11897. Munich, Germany: CESifo. [247]

Mason, Liliana. 2018. *Uncivil Agreement: How Politics Became Our Identity.* University of Chicago Press. [245, 247]

Michelitch, Kristin, and Stephen Utych. 2018. "Electoral cycle fluctuations in partisanship: Global evidence from eighty-six countries." *Journal of Politics* 80 (2): 412–27. [247]

Singh, Shane P, and Judd R Thornton. 2019. "Elections activate partisanship across countries." *American Political Science Review* 113 (1): 248–53. [247]

Sood, Gaurav, and Shanto Iyengar. 2016. "Coming to dislike your opponents: The polarizing impact of political campaigns." *Available at SSRN 2840225.* [247]

Stantcheva, Stefanie. 2021. "Understanding Tax Policy: How do People Reason?" *Quarterly Journal of Economics* 136 (4): 2309–69. https://doi.org/10.1093/qje/qjab029. [246]

Sunstein, Cass. 2017. *Republic: Divided Democracy in the Age of Social Media.* Princeton University Press. [245]