Self-serving biases in beliefs about collective outcomes

Shimon Kogan, Florian H. Schneider and Roberto A. Weber*

March 2, 2021

Beliefs about collective outcomes, such as economic growth or firm profitability, play an important role in many contexts. We study biases in the formation of such beliefs. Specifically, we explore whether over-optimism and self-serving biases in information processing—documented for beliefs about individual outcomes—affect beliefs about collective outcomes. We find that people indeed exhibit self-serving biases for collective outcomes, and that such biases are similar to biases for individual outcomes. In addition, we investigate whether collective self-delusion is mitigated by market institutions. If anything, biases in information processing are more pronounced in the presence of a market.

Keywords: beliefs, Bayes' rule, asymmetric updating, overconfidence, motivated reasoning

JEL Classification: I18, C93, Z13

---

* Kogan: Sloan School of Management, Massachusetts Institute of Technology, and Interdisciplinary Center Herzliya (IDC), skogan@mit.edu. Schneider: Department of Economics, University of Zurich, florian.schneider2@econ.uzh.ch. Weber: Department of Economics, University of Zurich, roberto.weber@econ.uzh.ch. We thank participants at several seminars and conferences for helpful comments.

1. Introduction

There is overwhelming evidence that individuals tend to maintain overly positive beliefs about their abilities (e.g. Svenson, 1981; Quattrone and Tversky, 1984), the likelihood of desired future life events (e.g. Irwin, 1953; Weinstein, 1980; Mayraz, 2013) and their own morality (e.g. Bénabou and Tirole, 2016; Gino, Norton and Weber, 2016). Such over-optimism often appears to result from biases in how people acquire and process information: individuals avoid information that challenges their overly positive beliefs (e.g. Dana, Weber and Kuang, 2007; Oster, Shoulson and Dorsey, 2013; Ganguly and Tasoff, 2017), update less in response to bad news than to good news (e.g. Eil and Rao, 2011; Möbius et al., 2017), and are less likely to remember negative past signals than positive past signals (Zimmermann, 2020; Saucet and Villeval, 2019). Such biases are supported by the psychological and motivational benefits from optimism and maintaining a positive self-image (Akerlof and Dickens, 1982; Rabin, 1994; Brunnermeier and Parker, 2005; Bénabou and Tirole, 2006 and 2011; Bénabou, 2013).

Most existing work on over-optimism and self-delusion focuses on beliefs about desirable *individual* outcomes. However, in many important economic contexts, people have to form beliefs about *collective* outcomes, such as future economic growth, firm profitability or success in containing the outbreak of a pandemic.[1] As with the benefits from maintaining a desirable self-image, people may often be motivated to maintain beliefs that collective outcomes will also be positive. For example, individuals may benefit psychologically and may be motivated to engage in productive activities like work and investment when they believe that future economic conditions will be positive. Similarly, employees in a firm may find it desirable and motivating to believe that the company is and will continue to perform well. The same motives that underlie self-delusion regarding individual outcomes could, therefore, also affect information processing about collective outcomes. This could, in turn, yield important economic consequences, as collective over-optimism and wishful-thinking about the price growth of widely held assets are believed to play an important role in the formation of speculative bubbles

[1] Multiple studies provide evidence that beliefs about macroeconomic expectations matter for economic decision making (e.g., Armona, Fuster and Zafar, 2019; Roth and Wohlfart, 2020; Andre et al, 2019). There is some previous work that investigates the role of personal experience in belief formation for collective outcomes (Malmendier and Nagel, 2011, 2016; Kuchler and Zafar, 2019; Cotofan, Cassar, Dur and Meier, 2021).

(Shiller, 2002; Foote, Gerardi and Willen, 2012; Cheng, Raina and Xiong, 2014) and in explaining market anomalies such as the equity home bias puzzle (Strong and Xu, 2003).[2]

In this paper, we study whether self-serving biases in information processing exist for beliefs about collective outcomes, and directly compare them to the beliefs that agents form over individual outcomes.[3] A challenge in investigating this question using naturally occurring collective outcomes is that the information generating process in such contexts is typically unknown, and, as a consequence, we do not know the rational benchmark for belief updating. To overcome this challenge, we design a laboratory experiment in which we can construct the information generating process and hold key aspects of it fixed when comparing belief updating for collective outcomes with belief updating for individual outcomes. We find that participants exhibit self-serving biases for collective outcomes, and that biases are remarkably similar to biases for beliefs about individual outcomes. We thereby provide important evidence that motivated reasoning matters for the formation of beliefs about collective outcomes.

In our experiment, subjects perform a task based on reasoning ability, with their relative performance determining whether or not they receive a monetary prize. After completing the task, individuals receive noisy signals about their relative performance. We elicit their beliefs about relative performance, both before and after receiving the noisy signals. By relying on individual ability, this task creates the conditions that have been found to give rise to self-serving information processing, whereby positive information is over weighted relative to negative information (Eil and Rao, 2011; Möbius, et al., 2017).

In our first principal treatment condition, *Individual*, participants work independently on the above task. Each individual is matched with another participant, and relative individual performance on the task in the pair determines who wins a monetary prize. We elicit participants' beliefs of the relative likelihood that they outperformed their competitor, both before and after providing them with a noisy signal of their relative performance. The *Individual*

[2] There is some suggestive evidence that motivated beliefs might affect market outcomes. Strong and Xu (2003) show that—in line with wishful thinking—fund managers are significantly more optimistic towards their home equity market. Cheng, Raina and Xiong (2014) present evidence on the role of overconfidence in financial crises. Ma (2015) find that banks with CEOs who were more optimistic about future payoffs of housing investments had worse crisis performance. Collective self-delusion is also believed to affect corporate behavior. Anecdotal evidence suggests that collective denial of unethical and illegal business practices played a role in corporate scandals (Anand, Ashforth and Joshi, 2005; Bénabou, 2013).

[3] A recent paper, somewhat related to ours, investigates how individuals' beliefs are influenced by observing the beliefs of other individuals of similar (true) ability (Oprea and Yuksel, 2020). They find that exposure to such beliefs increases the degree of positivity bias. A key difference between our studies and theirs is that we directly investigate beliefs over collective outcomes, while their study retains a focus on beliefs about individual ability.

condition provides us with a benchmark of self-serving biases in forming and updating beliefs about individual outcomes. Consistent with earlier work, we find both initial overconfidence and evidence of asymmetric updating—with individuals overweighting positive signals relative to negative ones—for beliefs about individual outcomes. While not the main focus of our project, this replication is of interest in itself as there is conflicting evidence on the existence of asymmetric updating (Buser, et al., 2018; Coutts, 2019; Schwardmann and Van der Weele, 2020; Oprea and Yuksel, 2020).[4]

Our main alternative condition, *Collective*, is identical except that the task and relative beliefs now involve the performance of six-person groups. Group members perform the task together, submitting a common response, and communicate with each other via a chat box. Each group competes with another six-person group, with the higher-performing group obtaining a prize. Hence, whether or not the group outperforms the other group is a collective outcome, determined by the joint performance of group members. We elicit incentivized beliefs about the group's relative performance, both before and after providing a noisy signal.

Our results provide clear evidence of self-serving biases in the formation of collective beliefs. We document such biases in two ways. First, the priors reveal that subjects exhibit overconfidence over collective outcomes. Second, we find that subjects update their beliefs asymmetrically, that is, they update less in response to bad news about their group's relative performance than to good news. Moreover, we also document other biases in the formation of beliefs about collective outcomes that are also present in the formation of beliefs about individual performance. For example, we find that subjects exhibit substantial conservatism when updating beliefs about collective performance—that is, they do not react as strongly to signals as they should according to the rational (Bayesian) benchmark—and base-rate neglect. All of the above biases are qualitatively similar to those we observe in the *Individual* condition. We thus conclude that biased information processing exists for collective outcomes in a manner similar to how it occurs for individual outcomes.

[4] Eil and Rao (2011), Sharot, Korn and Dolan (2011), Garrett and Sharot (2014), and Wiswall and Zafar (2015) find evidence for asymmetric updating. Ertac (2011) find that people put more weight on negative signals. These studies differ substantially from the Möbius et al. (2017) framework, making comparisons across studies challenging. Cacault and Grieder (2019) find evidence for asymmetric updating about others' ability. Coutts, Gerhards and Murad (2020) investigate self-attribution bias in the context of belief updating. Barron (2021), Coutts (2019) and Gotthard-Real (2017) look at updating of non-ego-relevant information, and find no asymmetric updating.

We also conduct an additional treatment condition to investigate whether an information-aggregation institution, such as a market, influences the formation of collective self-delusion. Self-delusion about collective outcomes might be particularly important in market contexts, such as asset trading. While individual biases might be mitigated by the collective judgment produced in markets (Camerer, 1987; Camerer et al., 1989; Forsythe et al., 1992), market interactions with others that have similar motives for self-deception could also reinforce biased beliefs (Shiller, 2002; Seybert and Bloomfield, 2009; Kogan, Kwasnica and Weber, 2011; Bénabou, 2013).

The *Market* condition is identical to *Collective*, except that after subjects receive feedback regarding their group's relative performance, they participate in an asset market in which they trade assets with the other members of their group. Each asset pays a positive dividend if their group wins the competition, but no dividend if their group loses the competition. This reflects a market situation where traders may engage in wishful-thinking about the value of an asset, either because they are financially committed to it (Shiller, 2002) or because they are elsewise affected by the profitability of the issuer, for example in the case of assets from a major local employer. After subjects have participated in the market, we elicit beliefs about their group's relative performance.

We find that market prices depart substantially from fundamentals, reflecting overconfidence, asymmetric updating, and reluctance to bet against the occurrence of desired outcomes. This suggests that collective self-delusion might indeed play an important role in market contexts. More importantly, we do not find that markets mitigate biases in information processing. If anything, we find that the market institution exacerbates biases: compared to the *Collective* condition, subjects underreact even more to bad signals in the *Market* condition than in the *Collective* condition, resulting in higher degrees of asymmetric updating.

Our findings add an important dimension to the growing body of evidence on the formation of self-serving beliefs. Given the widespread relevance of beliefs about collective outcomes, our findings that self-serving biases in belief formation extend to collective settings and that markets do not seem to substantially mitigate these biases are important.

The next section provides a detailed description of the design of our study. In Section 3, we present our results. Finally, section 4 concludes.

2. Study design

Our design builds on the framework of Möbius et al. (2017). In each round of the experiment, individuals, or groups, first compete by performing a task and then submit beliefs about their relative performance in this task. Next, they receive a noisy signal about their relative performance and submit their updated beliefs. The above steps are repeated in four rounds.[5]

This context allows us to investigate both overconfidence and updating behavior in response to good and bad signals. To study how overconfidence and biased information processing differ for individual and group outcomes, we manipulate whether the task and beliefs are about individual performance or collective performance. We also conduct an additional treatment in which we introduce a market to study how this institution influences the formation of collective beliefs.

2.1 The Task

At the beginning of the experiment, directly after groups are formed (in the *Collective* and *Market* conditions), each individual/group is randomly matched with another individual/group. This pairing is constant across the four rounds.

Across all treatment conditions (described below) and rounds of the experiment, subjects are first asked to work on solving a knapsack problem (Murawski and Bossaerts, 2016; Tang, et al., 2017) in a limited amount of time. Knapsack problems consist of deciding which objects in a finite set to select (i.e., to "put into the knapsack"). Each object has a value and a weight. The goal is to choose the set of objects that maximize the total value subject to some weight constraint:

$$max_{x_i \in \{0,1\}} \sum_{i}^{I} v_i x_i \quad s.t. \quad \sum_{i}^{I} w_i x_i \leq C$$

where I is the total number of objects, v_i is the value of object i, w_i is the weight of object i and C is the weight constraint. Finding the optimal solution is a combinatorial optimization problem that has no single solution approach. More formally, these problems are NP-hard, which means that no know algorithm solves these problems in an efficient manner as the size of the problem increases.[6] Practically, that means that there is no approach that ensures an optimal solution of

[5] In the design of Möbius et al. (2017), subjects receive multiple signals for each task and therefore repeatedly update their beliefs. We opt for only one belief update per task to decrease the number of markets in the *Market* condition.
[6] Efficient means that computational time increases in polynomial time as the size of the problem increases.

these problems across all instances. For example, the intuitive approach of calculating the value/weight ratio of each object and then inserting objects into the knapsack in descending order until the weight capacity is reached results in an optimal solution in some but not all problem instances.

The experiment introduced knapsack problems with differing levels of difficulty in random order across rounds. We chose these knapsack problems for several reasons. First, they are easy to explain to subjects. Second, knapsack problems allow a group of subjects to collaborate electronically. Finally, given that there is no global solution approach and many possible permutations, it is difficult for subjects to verify whether they have a "correct" answer. If, instead, subjects could easily verify whether they submitted an optimal response, beliefs regarding relative performance would be less clear in such cases. Finally, each instance of the problem presents a new challenge in which earlier solutions may not be helpful, meaning that there is some independence of performance across rounds.

In the *Individual* condition, subjects work independently on the knapsack problem on their computer screen. They observe the parameters for that specific problem, can click on items to include in the knapsack and observe the score of their current selection. The interface also records the best solution found thus far, allowing subjects to easily implement this solution. Figure C1 in the Appendix provides an example of the interface employed when working on the task.

In both the *Collective* and *Market* conditions, subjects are randomly organized into groups of six and work jointly on one knapsack problem, submitting one final solution for the group. Subjects see the current knapsack problem on their screen and can search for solutions. When they find a promising solution, they can share it with the other members of their group. Subjects observe the best solution found yet by any of their group members. To facilitate collaboration, subjects can communicate via a chat interface with their teammates. This communication includes free text as well as proposed solutions to the knapsack problem. Figure C2 in the Appendix provides an example of the interface employed when solving the task. Group composition is constant across the four rounds. The only difference between the *Collective* and the *Market* conditions is the information subjects receive between rounds, which we explain later.

Across all conditions, subjects get 90 seconds to initially inspect the knapsack problem for that round and 60 additional seconds to try out different solutions. In the *Collective* and *Market* conditions, subjects can communicate with their group members during the full 150 seconds.

2.2 Incentives, beliefs and markets

In the *Individual* condition, participants who provide a better solution than their matched counterpart receive 40 CHF (≈ \$40) while those with the inferior solution receive 10 CHF. Given the discrete nature of the solution, ties are resolved by observing the time it took to submit the final (best) solution. The *Collective* and *Market* conditions generate the same per-person payoffs from providing an inferior or a better solution in the knapsack problem: the group that provides the better solution receives 240 CHF (paid in equal shares to each member) while the group with the inferior solution receives 60 CHF.

Next, subjects are asked to report their belief, *p* (on a scale of 0-100), that their solution is better than that of the matched individual or group. We incentivize accuracy with a quadratic scoring rule:[7]

$$Payment\ (in\ CHF) = 10 - 10 * \left(\mathbf{1}(better)\ - \frac{p}{100}\right)^2$$

where $\mathbf{1}(better)$ equals 1 if a subject's solution is better in that round (and 0 otherwise) and p is the probability estimate.

After submitting the first confidence report, subjects receive a signal about their relative performance in that round. In the *Individual* condition, the signal ($\{negative, positive\}$) is drawn from a distribution that reveals the true relative performance in that round with 2/3 probability and gives the wrong relative performance with 1/3 probability. The same signal is drawn, at the group level, in both the *Collective* and *Market* conditions; all members of a group receive the same signal.[8]

[7] Möbius et al. (2017) and the subsequent literature use the probabilistic crossover method (also called matching probabilities) to incentivize the belief elicitation. This method, however, potentially depends on ambiguity attitudes and the ambiguity of the event evaluated (see e.g. Baillon, Cabantous and Wakker, 2012). In our experiment, we worried about potential differences between treatment conditions in ambiguity about relative performance, so we therefore chose a quadratic scoring rule instead. This rule is also easy to explain to participants.

[8] We explain to subjects that there is an urn with 3 balls: 2 balls correspond to the participant (group) that won and 1 ball corresponds to the participant (group) that lost. We explain that 1 ball is drawn from the urn, and both the subject (group), and the

In the *Market* condition, subjects next participate in a double-auction asset market.[9] Each group forms an independent market; that is, subjects trade assets *within their group*. The group trades a single Arrow-Debreu security that pays off 2 CHF if the group obtained the better knapsack solution in that round and 0 CHF otherwise. Subjects' security endowments in the market session are designed to add up to zero, as to neutralize aggregate incentive effects from the market.[10] Figure C3 in the Appendix provides an example of the market interface.

In all treatment conditions, after observing the additional feedback (and after participating in the market in the *Market* condition) and before proceeding to the next round, subjects are asked to submit an additional confidence report (posterior beliefs) in the same format as their initial confidence report. The payment scheme for the second confidence report is the same as the first one.

After playing all four rounds we measure risk-aversion, using the method by Gneezy and Potters (1997)—subjects receive an initial 3 CHF balance and decide how much of it to invest in a project that generates 6 times the investment with 25 percent probability and loses the investment with 75 percent probability. We also measure ambiguity attitudes by eliciting the certainty equivalent of a bet that pays 5 CHF if a color of the participant's choice (red or black) is drawn from an urn that consists of 10 red and black balls of unknown composition. To measure ambiguity aversion, we compare this certainty equivalent with the certainty equivalent of a lottery that pays 5 CHF with a probability of 50 percent. Finally, subjects fill out a survey eliciting various demographic characteristics.

To determine payments, we draw one of the four rounds and pay for the relative performance in that round's knapsack problem. Then a different round is randomly selected. One of the two belief estimates in this round is randomly drawn to count for the payment. Finally, in the *Market* condition, a third and different round is randomly selected to count for the market payment. This procedure limits possibilities for hedging within a round. In addition to these

"matched subject" ("matched group") observe to which person (group) this ball corresponds. This means that so the subject and the "matched subject" receive perfectly negative correlated signals about their own relative performance.

[9] In the *Market* condition, we only measure beliefs before the signal and after the market stage, but not between these two stages. We do so to keep the treatment conditions as similar as possible.

[10] In each round and market, three randomly drawn subjects start with 5 assets and 0 CHF cash, the three other subjects start with -5 assets and 10 CHF cash. In addition, each subject receives a loan of 6 CHF for trading that has to be paid back. We allow short selling. Trading is restricted in that trades are not allowed if they result in a negative cash balance, or if they potentially generate losses of more than the loan received.

payments, subjects receive all their payoffs from the tasks eliciting risk aversion and ambiguity attitudes.

Subjects were informed about the full procedure before the experiment started.

2.3 Procedures

Subjects were students from the joint subject pool of the University of Zurich and the Swiss Federal Institute of Technology (ETH). The experiment was programmed in z-Tree (Fischbacher, 2007) and participants were recruited with hroot (Bock et al., 2014).

At the beginning of a session, subjects received detailed instructions on knapsack problems, on the procedures for performing the task, on how payments would be determined, and, in the *Market* condition, on the market. All instructions were delivered both on paper and with pre-recorded audio files. Instructions and materials are available in Appendix D.[11] To ensure that subjects understood the instructions, they answered comprehension questions before the start of the experiment. Subjects also saw a trial knapsack problem to familiarize them with the structure of the problem and the solution interface. To familiarize subjects with the market interface in the *Market* condition, they participated in a trial market. In this trial market, subject traded assets whose payoffs depended on a virtual coin flip. We did not incentivize this trial period.

We collected data from a total of 324 subjects, 48 subjects in the *Individual* condition, 144 subjects in the *Collective* condition, and 132 subjects in the *Market* condition. The sessions lasted about 75 minutes in the *Individual* and *Collective* conditions and 120 minutes in the *Market* condition. Average earnings were CHF 44.46 (sd=CHF 17.71), or around US$45.

2.4 Econometric specification

To investigate biases in belief updating, we compare the observed updating to the Bayesian benchmark. Previous studies on processing of ego-relevant information focus on two deviations from Bayesian updating: asymmetric updating and conservatism. Asymmetric updating means that people react more strongly to positive than negative signals, a mechanism facilitating the preservation of positive self-image. Conservatism means that people react less strongly to the signals than predicted by Bayes' rule. We follow Möbius et al. (2017) and the subsequent

[11] We follow the use of voice recordings to deliver instructions, as in Bartling, Engl and Weber (2015). This, combined with standardized instructions and computerized interfaces, ensures highly replicable environments across sessions.

literature by estimating the following regression model, a linearized version of Bayes' rule (see also Grether, 1980; Augenblick and Rabin, 2021):

$$logit(posterior_i) = \delta_t logit(prior_i) + \beta_{L,t}\lambda_L \mathbf{1}(s_i = neg.) + \beta_{H,t}\lambda_H \mathbf{1}(s_i = pos.) + \varepsilon_i \quad (1)$$

where t denotes the treatment condition (*Individual*, *Collective* or *Market*), $\mathbf{1}(s_i = pos.)$ is an indicator of a positive signal, and $\lambda_H = -\lambda_L = \ln(2)$ is the log likelihood ratio of a positive signal. Note that δ_t measures the weight placed on prior beliefs, and $\beta_{L,t}$ and $\beta_{H,t}$ measures how strongly subjects react to negative and positive signals, respectively. If subjects update according to Bayes' rule, then $\delta_t = \beta_{L,t} = \beta_{H,t} = 1$. Asymmetric updating favoring a positive self-image is defined as $\beta_{L,t} < \beta_{H,t}$ and conservatism as $\beta_{L,t}, \beta_{H,t} < 1$. Finally, $\delta_t < 1$ can be interpreted as base-rate neglect and $\delta_t > 1$ as confirmation bias (Augenblick and Rabin, 2021). We exclude observations where subjects update in the wrong direction. We discuss such mistakes in belief updating in the next section.

3. Results

We first briefly assess the degree to which subjects' updating process represents an at least partially sensible response to information. We then discuss overconfidence in priors and study biases in belief updating for individual outcomes, replicating earlier work. Next, we address the main questions of this paper: Do people update beliefs about collective outcomes in a self-serving way? And, if so, do such biases in beliefs about collective outcomes differ from biases for individual outcomes? In the last section, we discuss the results of the *Market* condition.[12]

3.1 Updating mistakes

In this section, we discuss mistakes in belief updating, in particular updates in the wrong direction. We also investigate whether participants change their beliefs in response to the

[12] Since actual performance on the knapsack problems is not our focus, we do not analyze this measure in detail. In the *Individual* condition, subjects successfully find the optimal solution to the knapsack problem in 29.2% of all cases. The success rate is lower than in Murawski and Bossaerts (2016). This difference is likely due to subjects having less time to solve the knapsack problem in our experiment (up to 90s time difference). Groups' success rates are 73.7% in the *Collective* condition, and 78.4% in the *Market* condition. In the Collective and Market conditions, each subject sent on average 1.1 messages per round while working on the task.

(informative) signal.[13] Updating mistakes are indicators of loss of experimental control that can be due, for example, to unclear instructions.

In the *Collective* condition, 17.9% of subjects update at least once in the wrong direction and a total of 6.7% of the updating decisions go in the wrong direction. Moreover, subjects do not update their beliefs in response to the signal in 25.6% of all updating decisions and 4.2% of subjects do not update their beliefs in any of the four rounds. We find similar frequencies in the *Individual* condition (see Table A1 in the Appendix). Previous studies (e.g., Möbius et al, 2017) find more frequent updating mistakes, indicating the high data quality of our study.

Subjects in the *Market* condition receive multiple signals: the signal from the urn and signals from the market. It is therefore less clear what is a parsimonious definition of incorrect updating.[14] We consider an update as a mistake if the subject increases her belief in response to a negative signal from both the urn and the market, or if the subject decreases her belief in response to a positive signal from both the urn and the market.[15] Mistakes are somewhat more common in the *Market* condition than in the other treatment conditions: 23.5% of subjects update at least once in the wrong direction and 8.3% of subjects do not update their beliefs in any of the four rounds (see Table A1 in the Appendix). This is likely due to the updating decision being more complicated in the *Market* condition.

3.2 Priors: Confidence about individual and collective outcomes

Figure 1 shows the cumulative distribution of prior beliefs about winning the competition for all three treatment conditions. On average, the prior belief is 57.9% in the *Collective* condition, 56.6% in the *Individual* condition and 59.3% in the *Market* condition.[16] In each condition, average beliefs are statistically higher than 50% (p-value<0.01), indicating moderate degrees of overconfidence.[17] Biases do not differ between treatment conditions: Subjects exhibit similar

[13] On average, subjects react to the signal. Table 2 shows that the posterior beliefs are substantially and significantly lower for subjects who receive a positive signal compared to subjects who receive a negative signal.

[14] Suppose, for example, that a subject first receives a negative signal, but the market price reflects more confidence than this subject's prior. As the subject receives a negative signal from the urn, and a positive signal from the market, the subject might increase or decrease her belief.

[15] To define the market signal, we calculate the price of the last 10 trades, divided by 2. This is a measure of the optimism manifested in the market price. If this normalized price is higher than a subject's prior, it is considered a positive market signal for the subject; if the price is lower than the prior, it is considered a negative market signal for the subject.

[16] Priors do not differ significantly between rounds nor between knapsack problems. Prior beliefs are predictive of actual relative performance: a one percentage point higher belief translates into a 0.15% percentage point higher probability to actually be better (p-value=0.027).

[17] In the following analysis, standard errors are clustered on the matched-individual level (the two matched individuals form a cluster) in the *Individual* condition and on the matched-group level (the two matched groups form a cluster) in the *Collective* and *Market* conditions.

levels of confidence about collective and individual outcomes. We conclude that, in the aggregate, overconfidence is as prevalent for collective outcomes as for individual outcomes.

Figure 1: Distributions of prior beliefs by condition

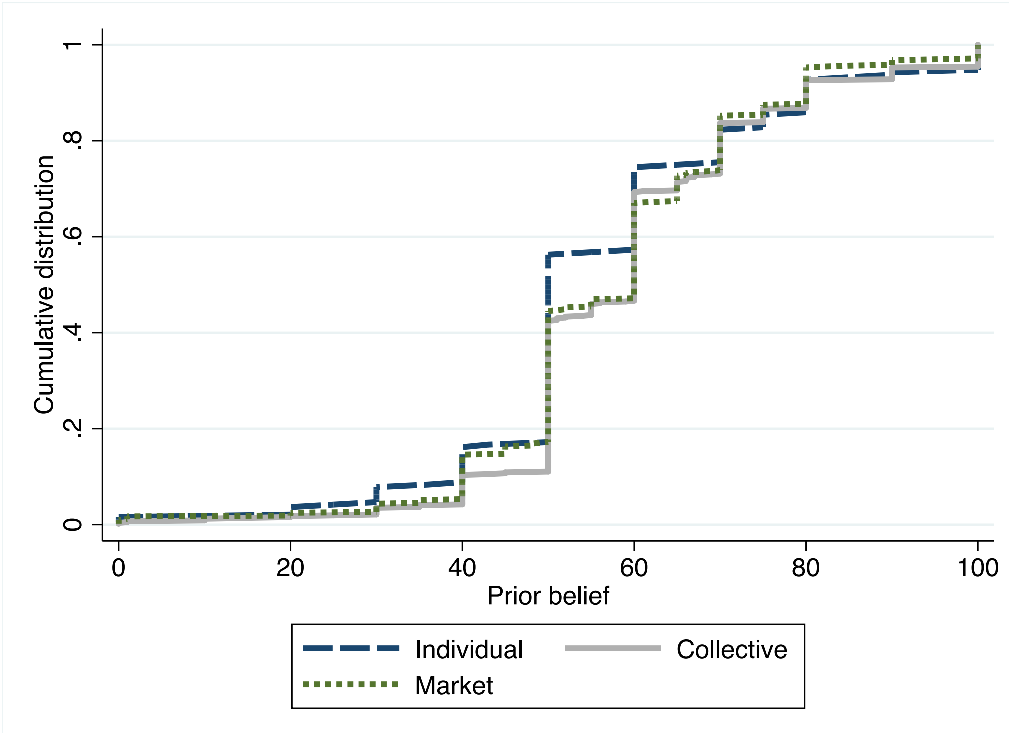

3.3 Updating beliefs about individual outcomes

Before we study how participants update beliefs about collective outcomes, we investigate whether our results replicate earlier work on self-serving biases in updating beliefs about individual outcomes. Figure 2 shows the mean absolute belief update conditional on the prior and the signal for the *Individual* condition, using a similar approach to Möbius et al. (2017). The figure compares subjects who received a positive signal and had a prior belief μ with subjects who received a negative signal and had a prior belief of $1-\mu$. If participants are Bayesians, the absolute magnitude of the belief update should be the same for both groups. If people update asymmetrically, however, the belief update should be larger for positive signals than for negative signals. For intermediate priors, we find that subjects update asymmetrically. For the very extreme prior category 0-9%, we find that subjects react more strongly to negative than to positive signals. This difference, however, is based on only six observations, and is not statistically significantly different from zero. If we compare the mean absolute belief revision with the Bayesian benchmark (red crosses), we see that subjects with low priors overreact to the signal, in particular if they receive a good signal. Subjects with intermediate priors tend to be conservative, particularly when they receive a negative signal.

Figure 2: Asymmetric updating in the Individual condition

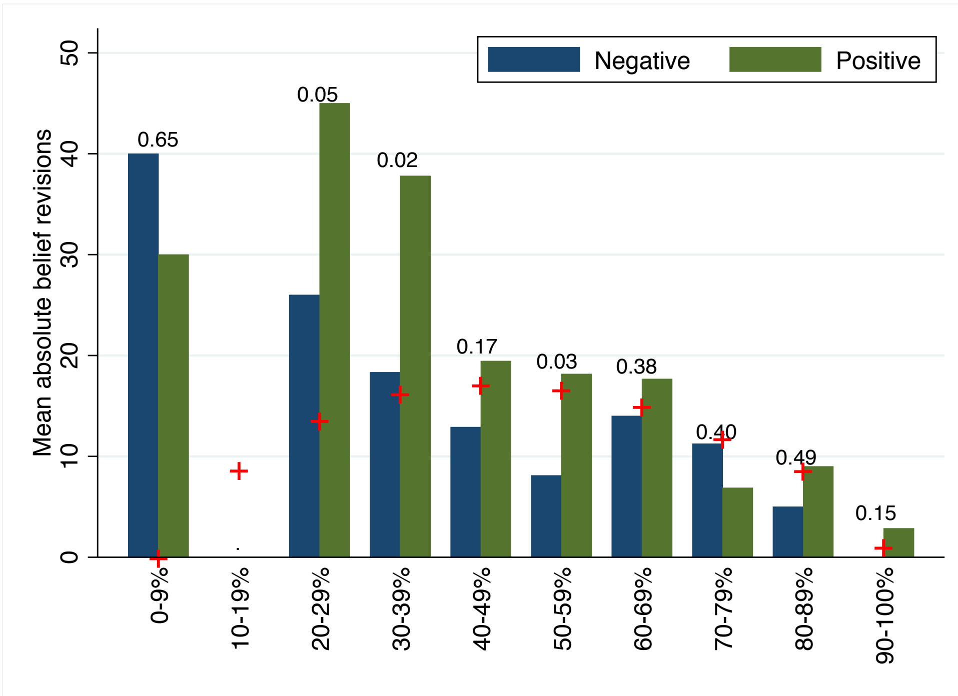

Mean absolute belief revisions by decile of prior belief in being of type equal to the signal received (following Möbius et al., 2017). + indicates the rational benchmark of Bayesian updating. Observations where people updated in the wrong direction are excluded. The numbers on top of the bars indicate p-values, testing whether the update in response to the negative and positive signals is equal.

The first column of Table 1 reports the estimated coefficients from model (1) for the *Individual* condition.[18] We find evidence for base-rate neglect ($\hat{\delta}_{Individual} < 1$, p-value<0.01). Moreover, we find that subjects react conservatively in response to negative signals ($\hat{\beta}_{L,Individual}$<1, p-value=0.014). For positive signals, however, behavior corresponds to the Bayesian benchmark ($\hat{\beta}_{H,Individual}$ is close to 1). Hence, subjects put more weight on positive signals than negative signals; that is, they update asymmetrically (p-value=0.003). Given that the non-parametric analysis revealed different updating patterns for extreme priors, we also report

[18] Note that the logit does not exist for the priors at the boundary (0 and 100). We follow the previous literature by excluding these observations.

estimated coefficients when restricting the sample to observations with priors in [20,80]. Estimates are similar (see Table 1, column 4). Table A2 in the Appendix demonstrates that the results are robust to other sample restrictions.

Comparing our estimates to previous studies, we find similar degrees of conservatism as Schwardmann and Van der Weele (2020) and Coutts (2019). However, unlike some earlier replications of Möbius et al. (2017) (see Buser et al., 2018; Schwardmann and Van der Weele, 2020; Coutts, 2019) we find robust evidence of asymmetric updating.[19]

Table 1: Updating behavior in Individual and Collective conditions

Subsample:	**Priors in (0,100)**			**Priors in [20,80]**		
	Individual	**Collective**	**Difference**	**Individual**	**Collective**	**Difference**
$\hat{\boldsymbol{\delta}}_t$	0.627***	0.710***	-0.082	0.542***	0.642***	-0.100
	(0.094)	(0.057)	(0.108)	(0.120)	(0.075)	(0.140)
$\hat{\boldsymbol{\beta}}_{L,t}$	0.671**	0.678***	-0.006	0.642***	0.597***	0.045
	(0.123)	(0.054)	(0.133)	(0.112)	(0.041)	(0.118)
$\hat{\boldsymbol{\beta}}_{H,t}$	1.133	0.834**	0.299**	1.090	0.881*	0.209*
	(0.111)	(0.073)	(0.131)	(0.102)	(0.060)	(0.117)
N	160	475		159	459	
$\mathbf{p}(\hat{\boldsymbol{\delta}}_t == \mathbf{1})$	0.001	0.000		0.001	0.001	
$\mathbf{p}(\hat{\boldsymbol{\beta}}_{L,t} == \mathbf{1})$	0.014	0.000		0.004	0.000	
$\mathbf{p}(\hat{\boldsymbol{\beta}}_{H,t} == \mathbf{1})$	0.243	0.045		0.391	0.072	
$\mathbf{p}(\hat{\boldsymbol{\beta}}_{L,t} == \hat{\boldsymbol{\beta}}_{H,t})$	0.003	0.154		0.005	0.000	

*Note: Estimated coefficients of model (1). Priors in (0,100): Observations where people updated in the wrong direction or with a prior=0 or =100 are excluded. Priors in [20,80]: Observations where people updated in the wrong direction or with a prior<20 or >80 are excluded. Standard errors clustered at matched-individual/group level; Standard errors in parentheses; Coefficient is significantly different from 1 (Bayesian benchmark) at * - p < 0.1; ** - p < 0.05; *** - p < 0.01. p(H) gives the p-value for testing hypothesis H.*

3.4 Updating beliefs about collective outcomes

Do these biases in the formation of beliefs about individual outcomes extend to beliefs about collective outcomes? In the following, we replicate the above analysis for the *Collective* condition. Figure 3 illustrates the mean absolute belief revision for the *Collective* condition, in a manner similar to Figure 2. We find a similar pattern as in the *Individual* condition: subjects with intermediate priors exhibit asymmetric updating and conservatism.

[19] Our study differs in some aspects from the previous studies: subjects update their prior only once in each round, and we use a different task (the knapsack task).

Figure 3: Asymmetric updating in the Collective condition

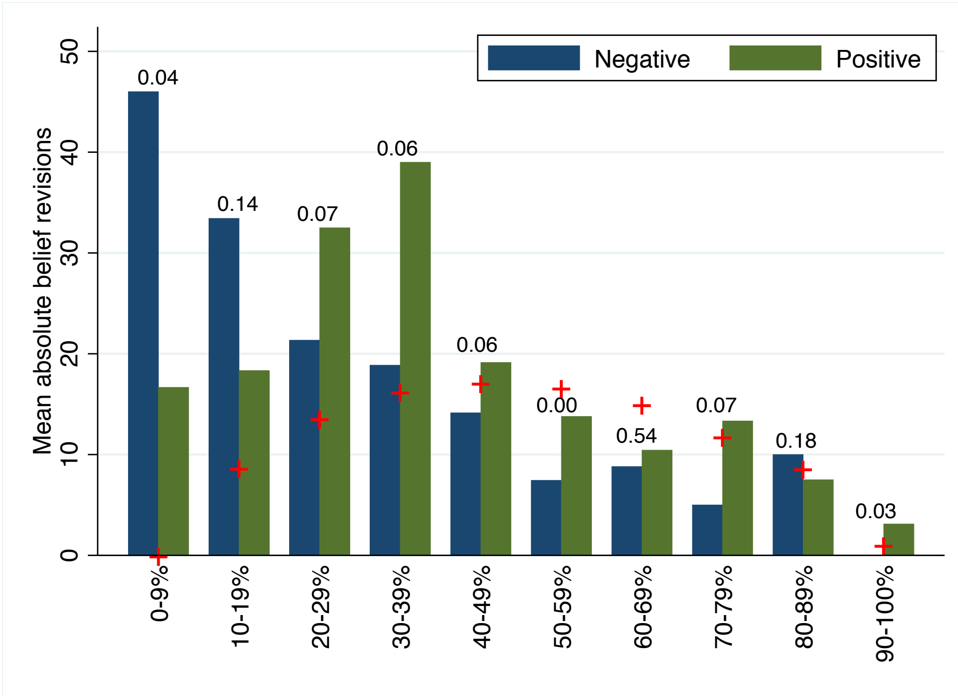

Mean absolute belief revisions by decile of prior belief in being of type equal to the signal received (following Möbius et al., 2017). + indicates the rational benchmark of Bayesian updating. Observations where people updated in the wrong direction are excluded. The number on top of the bars indicate p-values.

Table 1, columns 2 and 5, give the parameter estimates of model (1) for the *Collective* condition. As in the *Individual* condition, we find base-rate neglect, conservatism and asymmetry in belief updating. However, the difference between $\hat{\beta}_{L,Group}$ and $\hat{\beta}_{H,Group}$ is only statistically significant if we exclude observations with extreme priors (see Table A2 in the Appendix for different sample restrictions).

Columns 3 and 6 compare estimates between the *Collective* and the *Individual* conditions. There is some evidence for treatment differences in updating: subjects put slightly less weight on positive signals in the *Collective* condition compared to the *Individual* condition (p-value=0.029, entire sample). However, we cannot reject the joint hypothesis that all three coefficients are the same across the treatment conditions (p-value=0.159).

Instead of comparing individual belief updating, we can also compare average posteriors between treatment conditions. We do this in Table 2. In line with our analysis of individual belief updating, we do not find substantial differences in aggregate beliefs between the *Individual* and *Collective* conditions. There is weak evidence that, after receiving a negative signal, subjects in the *Collective* condition are more confident than subjects in the *Individual* condition (p-value=0.09). However, we can not reject the joint hypothesis that both differences are zero (p-value=0.12). We therefore conclude that biased information processing exists for beliefs about collective outcomes in a manner similar to how they occur for individual outcomes.

Table 2: Posterior beliefs

	Individual (I)	Collective (C)	Market (M)	Difference I and C	Difference C and M
Mean	57.5***	57.9***	55.1***	0.44	-2.80*
	(1.62)	(1.22)	(0.99)	(1.99)	(1.52)
Negative signal	42.5***	46.5***	47.0***	4.0*	0.44
	(1.86)	(1.48)	(1.41)	(2.34)	(1.97)
Positive signal	72.4***	69.5***	63.2***	2.9	-6.29***
	(2.43)	(1.62)	(1.26)	(2.88)	(1.99)
Difference	29.9***	23.0***	16.2***		
	(2.89)	(2.00)	(1.80)		

Note: Standard errors clustered at matched-individual/group level; Standard errors in parentheses;
** - p < 0.1; ** - p < 0.05; *** - p < 0.01.*

3.5 Market condition

In this section we analyze behavior in the *Market* condition.[20] The first row of Table 3 shows the trading volume. Subjects interacted frequently in the markets; on average, a subject traded about 4 assets per market round. Trading reflects subjects' beliefs about their group's performance: at the end of the market round, more optimistic subjects owned more assets than less optimistic subjects (see Table A3 in the Appendix). The trading volume is independent of the group's signal.

The second row of Table 3 gives the average market prices conditional on the signal.[21] The price is normalized such that a risk-neutral trader would buy the asset if she believed that the probability of winning the competition was higher than the normalized price, and she would sell the asset if she believed that it was lower than the price. Market prices incorporate signals: the average market price is substantially higher, by 16.5, if participants receive a positive signal.

[20] Behavior in the trial market round suggests that subjects understand how the market works. In the trial market round, subjects trade a risky asset (not incentivized), which is not connected to subjects' self-images. Prices in this trial market are close to CHF 1.00, the equilibrium prediction: the average price is CHF 1.10 and the median price is CHF 1.00.
[21] Figure A1 in the Appendix gives the cumulative distributions for market prices. Figure A2 shows average market prices for the four rounds separately; prices do not substantially differ between rounds. Figure A3 in the Appendix gives the price development over trading rounds.

The third row of Table 3 shows the objective expected value of the assets conditional on the signal.[22] Average prices are substantially higher than the expected value, in particular for markets that received a negative signal. Prices express even more optimism than the posterior beliefs documented in the *Collective* condition (69.5 for positive signals and 46.5 for negative signals, see Table 2), an estimate for what the beliefs in the *Market* condition might have been before subjects start trading. These high prices are consistent with evidence that people are reluctant to bet against the occurrence of desired outcomes (Seybert and Bloomfield, 2009; Morewedge, Tang and Larrick, 2018).

Table 3: Trading volume and asset prices

	Group signal positive	**Group signal negative**	**Difference**
Number of assets traded (per participant)	3.83 (0.25)	4.03 (0.59)	-0.2 (0.51)
Asset price, normalized	80.4 (2.56)	63.9 (4.02)	16.5*** (3.28)
Expected value of asset	66.7	33.3	

*Note: Price asset normalized is the average market price for the last 10 trades, divided by 2 (normalized). Expected value of asset is the normalized expected value of the asset = 2/3*200/2 resp. 1/3*200/2.*

Are the individual biases that we observe in the *Collective* condition mitigated by the collective judgment produced in markets? Or, does the collective over-optimism expressed in market prices bias participants' beliefs even more?[23] To study the impact of the market on subjects' beliefs, we estimate the parameters of model (1) for the *Market* condition; we regress subjects' posteriors (that is, their beliefs after they observed signals and then interacted in the market) on their priors and the signal from the urn.[24] We then compare the resulting coefficients to the *Collective* condition.[25] Table 4 gives the estimates. We also report estimates for the sample restricted to observations with a prior in [20,80].

[22] The average empirical value of the assets is almost the same as the expected value.
[23] After controlling for the groups' signals, the market price is not predictive of winning the competition (p-value=0.982). This suggests that, from a normative point of view, subjects should not react to the market prices.
[24] We therefore do not explicitly study how subjects incorporate market signals, such as prices, in their updating behavior. Instead, we estimate a similar model as for the *Collective* condition. This approach allows us to study whether the market interaction affects belief updating. In Appendix B, we incorporate market signals into model (1) and provide the corresponding estimates.
[25] Figure A4 in the Appendix replicates Figure 2, that is, the non-parametric analysis, for the Market condition. We find a similar pattern as in the *Collective* condition.

Table 4: Updating behavior in Market condition

Subsample:	Priors in (0,100)			Priors in [20,80]		
	Collective	**Market**	**Difference**	**Collective**	**Market**	**Difference**
$\hat{\boldsymbol{\delta}}_{\boldsymbol{t}}$	0.710*** (0.057)	0.643** (0.137)	0.067 (0.145)	0.642*** (0.075)	0.455*** (0.050)	0.187** (0.089)
$\hat{\boldsymbol{\beta}}_{\boldsymbol{L,t}}$	0.678*** (0.054)	0.411*** (0.130)	0.267* (0.137)	0.597*** (0.041)	0.317*** (0.075)	0.280*** (0.083)
$\hat{\boldsymbol{\beta}}_{\boldsymbol{H,t}}$	0.834** (0.073)	0.679*** (0.095)	0.155 (0.117)	0.881* (0.060)	0.813*** (0.050)	0.068 (0.076)
N	475	453		459	441	
p-value $\hat{\boldsymbol{\delta}}_{\boldsymbol{t}} == \mathbf{1}$	0.000	0.026		0.001	0.000	
p-value $\hat{\boldsymbol{\beta}}_{\boldsymbol{L,t}} == \mathbf{1}$	0.000	0.001		0.000	0.000	
p-value $\hat{\boldsymbol{\beta}}_{\boldsymbol{H,t}} == \mathbf{1}$	0.045	0.007		0.072	0.004	
p-value $\hat{\boldsymbol{\beta}}_{\boldsymbol{L,t}} == \hat{\boldsymbol{\beta}}_{\boldsymbol{H,t}}$	0.154	0.225		0.000	0.000	

*Note: Estimated coefficients of model (1). Priors in (0,100): Observations where people updated in the wrong direction or with a prior=0 or =100 are excluded. Priors in [20,80]: Observations where people updated in the wrong direction or with a prior<20 or >80 are excluded. Standard errors clustered at matched-individual/group level; Standard errors in parentheses; Coefficient is significantly different from 1 at * - p < 0.1; ** - p < 0.05; *** - p < 0.01.*

For both the full and the restricted samples, we can reject the joint hypothesis that all three estimates are the same across the *Market* and *Collective* conditions (p-value=0.013 and p-value=0.002, respectively). For subjects with intermediate priors, we find a larger degree of base-rate neglect in the *Market* condition than in the *Collective* condition ($\hat{\delta}_{Market} < \hat{\delta}_{Group}$). A potential explanation is that subjects update their beliefs not only in response to the signal from the urn but also in response to signals from the markets. These additional updates reduce the weight of the initial prior (see Appendix B).

Compared with the *Collective* condition, subjects in the *Market* condition react less to negative signals. For intermediate priors, the difference in $\hat{\beta}_{L,t}$ is statistically significant at the 1%-level. This finding is robust to different sample restrictions (see Table A4 in the Appendix). After receiving a bad signal, the optimism reflected in market prices seems to help subjects to partly restore their confidence.[26] We do not find a treatment difference for positive signals. This potentially reflects the fact that the assets exhibit less overpricing after receiving a positive signal.

Instead of comparing individual belief updating, we can also compare average posteriors between treatment conditions. We do this in Table 2. The average posterior of subjects who

[26] We find some evidence that subjects incorporated market signals in their beliefs (see Appendix B for details).

received a positive signal is 6.3 percentage points lower than the posterior in the *Collective* condition (p-value=0.005). There is no treatment difference for subjects who received a negative signal. At first glance, this seems to be in contradiction with the finding that individual updating differs only in response to *negative* signals ($\beta_{L,t}$), not in response to positive signals ($\beta_{H,t}$). However, these two findings can be explained by the fact that we observe a second treatment effect, a stronger degree of base-rate neglect in the *Market* condition.

We conclude that the market does not mitigate biases in belief formation regarding collective outcomes. If anything, we find that biases are exacerbated: compared to the *Collective* condition, subjects underreact even more to negative signals in the *Market* condition.

4. Conclusion

We explore whether over-optimism and self-serving biases in information processing exist for collective outcomes in a manner similar to how they occur for individual outcomes. We first show that subjects exhibit such biases for beliefs about individual outcomes: subjects are overconfident and update their beliefs asymmetrically in response to new information. That is, they put more weight on good news than on bad news. This replicates patterns found in many—but not all—previous studies that investigate this question.

We then investigate biases in beliefs about collective outcomes. As with beliefs about individual outcomes, subjects also exhibit self-serving beliefs about collective outcomes, and magnitudes are remarkably similar to those for biases about individual outcomes. Thus, our first main novel contribution is to document that the tendency to overweight positive information more than negative information, i.e., asymmetric updating, also extends to the formation of beliefs about collective outcomes. Given the importance of beliefs about such outcomes—from macroeconomic performance to firm profitability—for a wide variety of economic behaviors, this observation is important.

We also investigate how such belief formation is influenced by the presence of an information aggregation institution, specifically, a market. Collective self-delusion potentially plays an important role for market outcomes. However, it has been argued that markets can mitigate individual biases. We find that the market institution, if anything, exacerbates biases. We also observe that market prices depart substantially from fundamentals, in manner consistent with positive self-delusion. Thus, our findings suggest that, rather than reducing the tendency to engage in self-serving information processing, aggregating beliefs through an institution such as a market may have the opposite effect. This is consistent with the observation that many instances of such collective self-delusion occur in market contexts, as in speculative bubbles.

6. Literature

Akerlof, G. A., and Dickens, W. T. (1982) "The Economic Consequences of Cognitive Dissonance," American Economic Review, 72(3): 307–319.

Anand, V., Ashforth, B. E. and Joshi, M. (2005) "Business as Usual: The Acceptance and Perpetuation of Corruption in Organizations," Academy of Management Executive, 19(4): 9-23.

Andre, P., Pizzinelli, C., Roth, C. and Wohlfart, J. (2019) "Subjective Models of the Macroeconomy: Evidence from Experts and a Representative Sample," working paper.

Armona, L., Fuster, A. and Zafar, B. (2019) "Home Price Expectations and Behaviour: Evidence from a Randomized Information Experiment," Review of Economic Studies, 86(4): 1371–1410.

Augenblick, N. and Rabin, M. (2021) "Belief Movement, Uncertainty Reduction, and Rational Updating," Quarterly Journal of Economics, forthcoming.

Baillon, A., Cabantous, L. and Wakker, P. P. (2012) "Aggregating imprecise or conflicting beliefs: an experimental investigation using modern ambiguity theories," Journal of Risk and Uncertainty, 44: 115-147.

Barron, K. (2021) "Belief Updating: Does the `Good-news, Bad-news´ Asymmetry Extend to Purely Financial Domains?," Experimental Economics, forthcoming.

Bartling, B., Engl, F. and Weber, R. A. (2015) "Game form misconceptions are not necessary for a willingness-to-pay vs. willingness-to-accept gap," Journal of the Economic Science Association, 1(1): 72-85.

Bénabou, R. (2013) "Groupthink: Collective Delusions in Organizations and Markets," Review of Economic Studies, 80: 429-462.

Bénabou, R., and Tirole, J. (2006) "Incentives and Prosocial Behavior," American Economic Review, 96(5): 1652-1678.

Bénabou, R., and Tirole, J. (2011). "Identity, Morals, and Taboos: Beliefs as Assets," Quarterly Journal of Economics, 126: 805-855.

Bénabou, R., and Tirole, J. (2016) "Mindful Economics: The Production, Consumption, and Value of Beliefs," Journal of Economic Perspectives, 30(3): 141-164.

Bock, O., Baetge, I., and Nicklisch, A. (2014) "hroot: Hamburg registration and organization online tool," European Economic Review, 71: 117-120.

Brookins, P., Lucas, A., Ryvkin, D. (2014) "Reducing within-group overconfidence through group identity and between-group confidence judgments," Journal of Economic Psychology, 44, 1-12.

Brunnermeier, M. K., and Parker, J. A. (2005) "Optimal Expectations," American Economic Review 95(4): 1092-1118.

Buser, T., Niederle, M. and Oosterbeek, H. (2014) "Gender, Competitiveness, and Career Choices," Quarterly Journal of Economics, 129(3): 1409-1447.

Buser, T., Gerhards, L., and Van der Weele, J. J. (2018) "Responsiveness to feedback as a personal trait," Journal of Risk and Uncertainty, 56(2): 165–192.

Cacault, M. P. and Grieder, M. (2019) "How group identification distorts beliefs," Journal of Economic Behavior & Organization, 164: 63-76.

Camerer, C. F. (1987) "Do Biases in Probability Judgment Matter in Markets? Experimental Evidence," American Economic Review, 77(5): 981-997.

Camerer, C., F. Loewenstein, G., and Weber, M. (1989) "The Curse of Knowledge in Economic Settings: An Experimental Analysis," Journal of Political Economy, 97(5): 1232-1254.

Cheng, I., Raina, S., and Xiong, W. (2014) "Wall Street and the Housing Bubble," American Economic Review, 104(9): 2797–2829.

Cotofan, M., Cassar, L., Dur, R., and Meier, S. (2020) "Macroeconomic Conditions When Young Shape Job Preferences for Life," working paper.

Coutts, A. (2019) "Good News and Bad News are Still News: Experimental Evidence on Belief Updating," Experimental Economics, 22(2): 369-395.

Coutts, A., Gerhards, L., and Murad, Z. (2020) "Who to blame? Self-serving attribution bias with multi-dimensional uncertainty," working paper.

Dana, J., Weber, R. A., and Kuang, J. X. (2007) "Exploiting moral wiggle room: experiments demonstrating an illusory preference for fairness," Economic Theory, 33(1): 67-80.

Eil, D. and Rao, J. M. (2011) "The Good News-Bad News Effect: Asymmetric Processing of Objective Information about Yourself," American Economic Journal: Microeconomics, 3(2): 114–138.

Ertac, S. (2011) "Does self-relevance affect information processing? Experimental evidence on the response to performance and non-performance feedback," Journal of Economic Behavior and Organization, 80(3): 532-545.

Fischbacher, U. (2007) “z-Tree: Zurich toolbox for ready-made economic experiments,” Experimental Economics, 10(2): 171–178.

Foote, C. L., Gerardi, K. S., and Willen, P. S. (2012) “Why Did So Many People Make So Many Ex Post Bad Decisions? The Causes of the Foreclosure Crisis,” NBER Working Paper.

Forsythe, R., Nelson, F., Neumann, G. R., and Wright, J. (1992) “Anatomy of an Experimental Political Stock Market,” American Economic Review, 82(5): 1142-1161.

Ganguly, A. R., and Tasoff, J. (2017) “Fantasy and Dread: The Demand for Information and the Consumption Utility of the Future.” Management Science, 63(12): 3999-4446.

Garrett, N., Sharot, T. (2014) “How Robust Is the Optimistic Update Bias for Estimating Self-Risk and Population Base Rates?,” PLoS ONE, 9(6).

Gino, F., Norton, M. I., and Weber, R. A. (2016) “Motivated Bayesians: Feeling Moral While Acting Egoistically,” Journal of Economic Perspectives, 30(3): 189-212.

Gneezy, U., and Potters, J. (1997) “An Experiment on Risk Taking and Evaluation Periods.” Quarterly Journal of Economics, 112(2): 631–645.

Gotthard-Real, A. (2017) “Desirability and Information Processing: An Experimental Study,” Economics Letters, 152: 96-99.

Grether, D. M. (1980) “Bayes Rule as a Descriptive Model: The Representativeness Heuristic,” Quarterly Journal of Economics, 95(3): 537-557.

Irwin, F. W. (1953) “Stated Expectations as Functions of Probability and Desirability of Outcomes,” Journal of Personality, 21(3): 329-335.

Kogan, S., Kwasnica, A. M. and Weber, R. A. (2011) “Coordination in the presences of asset markets,” American Economic Review, 101(2): 927-947.

Kuchler, T. and Zafar, B. (2019) “Personal Experiences and Expectations about Aggregate Outcomes,” Journal of Finance, 74: 2491-2542.

Ma, Y. (2015) “Bank CEO Optimism and the Financial Crisis,” working paper.

Malmendier, U., and Nagel, S. (2011) “Depression babies: Do macroeconomic experiences affect risk taking?,” Quarterly Journal of Economics, 126: 373–416.

Malmendier, U., and Nagel, S. (2016) “Learning from inflation experiences,” Quarterly Journal of Economics, 131: 53–87.

Mayraz, G. (2013) “Wishful Thinking,” working paper.

Möbius, M. M., Niederle, M., Niehaus, P. and Rosenblat, T. S. (2017) "Managing Self-Confidence," working paper.

Morewedge, C. K., Tang, S. and Larrick, R. P. (2018) "Betting Your Favorite to Win: Costly Reluctance to Hedge Desired Outcomes," Management Science, 64(3): 983-1476.

Murawski, C. and Bossaerts, P. (2016) "How Humans Solve Complex Problems: The Case of the Knapsack Problem," Scientific Reports 6.

Niederle, M. and Vesterlund, L. (2007) "Do Women Shy Away from Competition? Do Men Compete Too Much?," Quarterly Journal of Economics, 122(3): 1065–1101.

Oprea, R. and Yuksel, S. (2020) "Social Exchange of Motivated Beliefs," working paper.

Oster, E., Shoulson, I. and Dorsey, E. R. (2013) "Optimal expectations and limited medical testing: Evidence from Huntington disease." American Economic Review, 103(2): 804–830.

Quattrone, G. A., and Tversky, A. (1984) "Causal versus diagnostic contingencies: On self-deception and on the voter's illusion," Journal of Personality and Social Psychology, 46(2): 237-248.

Rabin, M. (1994) "Cognitive dissonance and social change," Journal of Economic Behavior and Organization, 23: 177-194.

Roth, C. and Wohlfart, J. (2020) "How Do Expectations about the Macroeconomy Affect Personal Expectations and Behavior?," Review of Economics and Statistics, 102(4), 731-748.

Saucet, C. and Villeval, M. C. (2019) "Motivated memory in dictator games," Games and Economic Behavior, 117, 250-275.

Seybert, N., and Bloomfield, R. (2009) "Contagion of wishful thinking in markets," Management Science, 55(5): 738-751.

Schwardmann, P. and Van der Weele, J. J. (2020) "Deception and Self-deception," Nature Human Behavior, 3: 1055-1061.

Sharot, T., Korn, C. W. and Dolan, R .J. (2011) "How unrealistic optimism is maintained in the face of reality," Nature Neuroscience, 14(11): 1475–1481.

Shiller, R. J. (2002) "Bubbles, Human Judgment, and Expert Opinion," Financial Analysts Journal, 58(3): 18-26.

Strong, N. and Xu, X. (2003) "Understanding the Equity Home Bias: Evidence from Survey Data," Review of Economics and Statistics, 85(2): 307-312.

Svenson, O. (1981) "Are we all less risky and more skillful than our fellow drivers?," Acta Psychologica, 47: 143-148.

Tang, S., Huang, H., Bowman, E., Yadav, N., Murawski, C. and Bossaerts, P. (2017) "The Efficient Markets Hypothesis Does Not Hold When Securities Valuation is Computationally Hard," working paper.

Weinstein, N. D. (1980) "Unrealistic optimism about future life events," Journal of Personality and Social Psychology, 39(5): 806-820.

Wiswall, M., and Zafar, B. (2015) "How do college students respond to public information about earnings?," Journal of Human Capital, 9(2): 117-169.

Zimmermann, F. (2020) "The Dynamics of Motivated Beliefs," American Economic Review, 110(2): 337-61.

Appendix A: Additional results

Table A1: Updating mistakes

	Individual	**Collective**	**Market**	**Previous studies**
percent no update	28.1	25.6	32.0	36 (M); 41 (C)
percent subjects never update	2.1	4.2	8.3	16 (M)
percent updates wrong direction	3.6	6.7	8.1	10 (M,B); 4.8 (C)
percent subjects with at least one update wrong direction	12.5	17.9	23.5	27 (M)

Notes: "no update" is defined as the prior being equal to the posterior; "update in the wrong direction" is defined as a negative update in response to a positive signal or a positive update in response to a negative signal for the Individual and Collective condition. For the Market condition, it is defined as a negative update in response to both, a positive signal and a "positive market signal" (average price of the last 10 trades/2 > prior) or a positive update in response to both, a negative signal and a "negative market signal" (average price of the last 10 trades/2 < prior). Previous studies gives the percentages from previos studies: (C) refers to Coutts (2019), (B) refers to Buser et al. (2018) and (M) refers to Möbius et al. (2017).

Table A2: Robustness checks updating

	Priors in (0,100)	Priors in [10,90]	Priors in [20,80]	Priors in [30,70]	Priors in [40,60]	Möbius subs.	Möbius subs. 2
				Individual condition			
$\hat{\boldsymbol{\delta}}_t$	0.627***	0.627***	0.542***	0.575***	0.871	0.650***	0.553***
	(0.094)	(0.094)	(0.120)	(0.150)	(0.155)	(0.101)	(0.135)
$\hat{\boldsymbol{\beta}}_{L,t}$	0.671**	0.671**	0.642***	0.597***	0.609***	0.703**	0.669***
	(0.123)	(0.123)	(0.112)	(0.110)	(0.109)	(0.126)	(0.114)
$\hat{\boldsymbol{\beta}}_{H,t}$	1.133	1.133	1.090	1.043	0.951	1.231*	1.175
	(0.111)	(0.111)	(0.102)	(0.103)	(0.098)	(0.115)	(0.109)
N	160	160	159	139	117	142	141
p-value $\hat{\boldsymbol{\delta}}_t == 1$	0.0006	0.0006	0.0009	0.0093	0.4117	0.0021	0.0031
p-value $\hat{\boldsymbol{\beta}}_{L,t} == 1$	0.0136	0.0136	0.0040	0.0013	0.0015	0.0269	0.0078
p-value $\hat{\boldsymbol{\beta}}_{H,t} == 1$	0.2430	0.2430	0.3905	0.6802	0.6202	0.0568	0.1206
p-value $\hat{\boldsymbol{\beta}}_{L,t} == \boldsymbol{\beta}_{H,t}$	0.0033	0.0033	0.0048	0.0084	0.0138	0.0024	0.0040
				Collective condition			
$\hat{\boldsymbol{\delta}}_t$	0.710***	0.597***	0.642***	0.598***	0.619***	0.706***	0.629***
	(0.057)	(0.065)	(0.075)	(0.071)	(0.107)	(0.072)	(0.089)
$\hat{\boldsymbol{\beta}}_{L,t}$	0.678***	0.618***	0.597***	0.558***	0.528***	0.735***	0.643***
	(0.054)	(0.039)	(0.041)	(0.048)	(0.058)	(0.081)	(0.039)
$\hat{\boldsymbol{\beta}}_{H,t}$	0.834**	0.900*	0.881*	0.854***	0.812***	0.901	0.960
	(0.073)	(0.049)	(0.060)	(0.045)	(0.046)	(0.081)	(0.065)
N	475	473	459	414	338	394	380
p-value $\hat{\boldsymbol{\delta}}_t == 1$	0.0004	0.0001	0.0006	0.0001	0.0045	0.0018	0.0016
p-value $\hat{\boldsymbol{\beta}}_{L,t} == 1$	0.0001	0.0000	0.0000	0.0000	0.0000	0.0072	0.0000
p-value $\hat{\boldsymbol{\beta}}_{H,t} == 1$	0.0449	0.0643	0.0715	0.0077	0.0019	0.2481	0.5454
p-value $\hat{\boldsymbol{\beta}}_{L,t} == \boldsymbol{\beta}_{H,t}$	0.1536	0.0001	0.0002	0.0000	0.0003	0.2626	0.0007
				Difference			
Difference $\hat{\boldsymbol{\delta}}_t$	-0.082	0.030	-0.100	-0.023	0.252	-0.057	-0.076
	(0.108)	(0.112)	(0.140)	(0.164)	(0.186)	(0.122)	(0.159)
Difference $\hat{\boldsymbol{\beta}}_{L,t}$	-0.006	0.053	0.045	0.040	0.081	-0.032	0.026
	(0.133)	(0.128)	(0.118)	(0.118)	(0.122)	(0.147)	(0.119)
Difference $\hat{\boldsymbol{\beta}}_{H,t}$	0.299**	0.233*	0.209*	0.189*	0.139	0.331**	0.215*
	(0.131)	(0.120)	(0.117)	(0.111)	(0.107)	(0.139)	(0.125)
p-value asym.	0.083	0.228	0.284	0.355	0.679	0.086	0.274
p-value joint test	0.159	0.220	0.373	0.410	0.452	0.100	0.401

*Note: Estimated coefficients of model (1). Observations where people updated in the wrong direction or with a prior of 0 percent of 100 percent are excluded are excluded. "Priors in [10,90]" means that the sample is restricted to observations with a prior in between 10 percent and 90 percent. Möbius subs. restricts the sample to subjects that never updated in the wrong direction (in the 4 rounds) and to subjects who updated their beliefs at least once (the subsample explored in Möbius et al. (2017)). Möbius subs. 2 restricts Möbius subsample to observations with a prior in [20,80]. "p-value asym." is the p-value from the test on whether there is a difference in asymmetric updating (H0: $\hat{\beta}_{H,Group} - \hat{\beta}_{H,Group} = \hat{\beta}_{H,Individual} - \hat{\beta}_{H,Individual}$). "p-value joint test" is the p-value from a joint test that δ_t, $\beta_{L,t}$ and $\beta_{H,t}$ do not differ between the two treatment conditions (F-test). Standard errors clustered at matched-individual/group level; Standard errors in parentheses; * - $p < 0.1$; ** - $p < 0.05$; *** - $p < 0.01$.*

Table A3: Relationship between beliefs and trading in the market

Rank	**Average number assets**
1 (least optimistic)	-0.76
2	-1.75
3	-.41
4	.09
5	1.19
6 (most optimistic)	1.64

Dependent Variable: Average number assets	
Rank	0.61***
	(0.150)
Constant	-2.13***
	(0.526)
N	528
R^2	.076

*Notes: For each round and group, subjects are ranked according to their posterior belief with rank=1 being the least optimistic subject in the group. We then correlate the rank with the number of assets the participant owns at the end of the market period. The total number of assets sum up to 0 by design. Standard errors clustered at matched-group level; Standard errors in parentheses; * - $p < 0.1$; ** - $p < 0.05$; *** - $p < 0.01$.*

Table A4: Robustness checks updating

	Priors in (0,100)	Priors in [10,90]	Priors in [20,80]	Priors in [30,70]	Priors in [40,60]	Möbius subs.	Möbius subs. 2
				Collective condition			
$\hat{\boldsymbol{\delta}}_{\boldsymbol{t}}$	0.710***	0.597***	0.642***	0.598***	0.619***	0.706***	0.629***
	(0.057)	(0.065)	(0.075)	(0.071)	(0.107)	(0.072)	(0.089)
$\hat{\boldsymbol{\beta}}_{\boldsymbol{L,t}}$	0.678***	0.618***	0.597***	0.558***	0.528***	0.735***	0.643***
	(0.054)	(0.039)	(0.041)	(0.048)	(0.058)	(0.081)	(0.039)
$\hat{\boldsymbol{\beta}}_{\boldsymbol{H,t}}$	0.834**	0.900*	0.881*	0.854***	0.812***	0.901	0.960
	(0.073)	(0.049)	(0.060)	(0.045)	(0.046)	(0.081)	(0.065)
N	475	473	459	414	338	394	380
				Market condition			
$\hat{\boldsymbol{\delta}}_{\boldsymbol{t}}$	0.643**	0.409***	0.455***	0.403***	0.387***	0.514***	0.461***
	(0.137)	(0.046)	(0.050)	(0.064)	(0.071)	(0.076)	(0.067)
$\hat{\boldsymbol{\beta}}_{\boldsymbol{L,t}}$	0.411***	0.288***	0.317***	0.304***	0.239***	0.415***	0.406***
	(0.130)	(0.075)	(0.075)	(0.081)	(0.082)	(0.094)	(0.089)
$\hat{\boldsymbol{\beta}}_{\boldsymbol{H,t}}$	0.679***	0.798***	0.813***	0.787***	0.735***	0.878*	0.889
	(0.095)	(0.046)	(0.050)	(0.061)	(0.063)	(0.059)	(0.062)
N	453	448	441	394	303	338	334
p-value $\boldsymbol{\delta_t == 1}$	0.0261	0.0000	0.0000	0.0000	0.0000	0.0001	0.0000
p-value $\boldsymbol{\beta_{L,t} == 1}$	0.0011	0.0000	0.0000	0.0000	0.0000	0.0001	0.0001
p-value $\boldsymbol{\beta_{H,t} == 1}$	0.0069	0.0013	0.0040	0.0061	0.0017	0.0655	0.1013
p-value $\boldsymbol{\beta_{L,t} == \beta_{H,t}}$	0.2252	0.0002	0.0002	0.0008	0.0005	0.0004	0.0002
				Difference			
Difference $\hat{\boldsymbol{\delta}}_{\boldsymbol{t}}$	0.067	0.188**	0.187**	0.195**	0.232*	0.193*	0.168
	(0.145)	(0.077)	(0.089)	(0.094)	(0.126)	(0.102)	(0.109)
Difference $\hat{\boldsymbol{\beta}}_{\boldsymbol{L,t}}$	0.267*	0.330***	0.280***	0.253**	0.289***	0.319**	0.236**
	(0.137)	(0.083)	(0.083)	(0.092)	(0.098)	(0.121)	(0.095)
Difference $\hat{\boldsymbol{\beta}}_{\boldsymbol{H,t}}$	0.155	0.102	0.068	0.067	0.077	0.022	0.071
	(0.117)	(0.065)	(0.076)	(0.074)	(0.076)	(0.098)	(0.087)
p-value asym.	0.624	0.027	0.039	0.102	0.067	0.082	0.135
p-value joint test	0.013	0.001	0.002	0.012	0.020	0.065	0.0614

*Note: Estimated coefficients of model (1). Observations where people updated in the wrong direction or with a prior of 0 percent of 100 percent are excluded are excluded. "Priors in [10,90]" means that the sample is restricted to observations with a prior in between 10 percent and 90 percent. Möbius subs. restricts the sample to subjects that never updated in the wrong direction (in the 4 rounds) and to subjects who updated their beliefs at least once. Möbius subs. 2 restricts Möbius subsample to observations with a prior in [20,80]. "p-value asym." is the p-value from the test on whether there is a difference in asymmetric updating (H0: $\hat{\beta}_{H,Group} - \hat{\beta}_{H,Group} = \hat{\beta}_{H,Market} - \hat{\beta}_{H,Market}$). "p-value joint test" is the p-value from a joint test that δ_t, $\beta_{L,t}$ and $\beta_{H,t}$ do not differ between the two treatment conditions (F-test). Standard errors clustered at matched-individual/group level; Standard errors in parentheses; * - $p < 0.1$; ** - $p < 0.05$; *** - $p < 0.01$.*

Figure A1: Distribution market price and trading volume

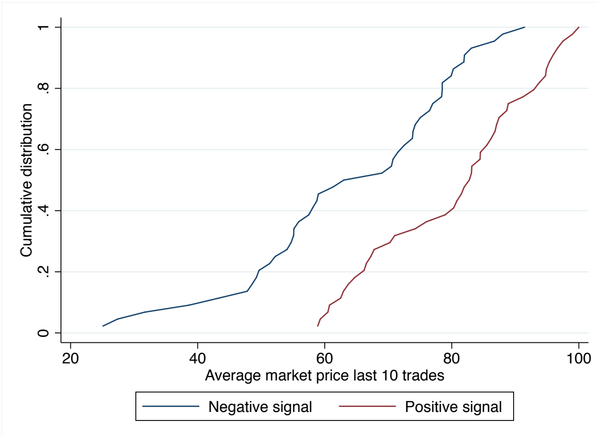

(a) Market price

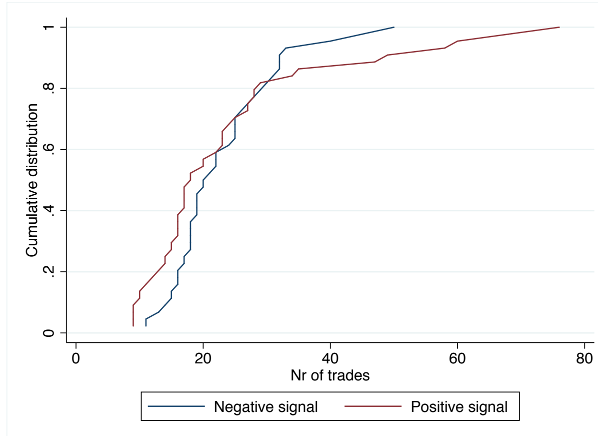

(b) Trading volume

Notes: For each market, we calculate the average price of the last 10 trades and the total number of trades. The figures show the cumulative distributions of these two variables, conditional on the signal.

Figure A2: Market prices over rounds

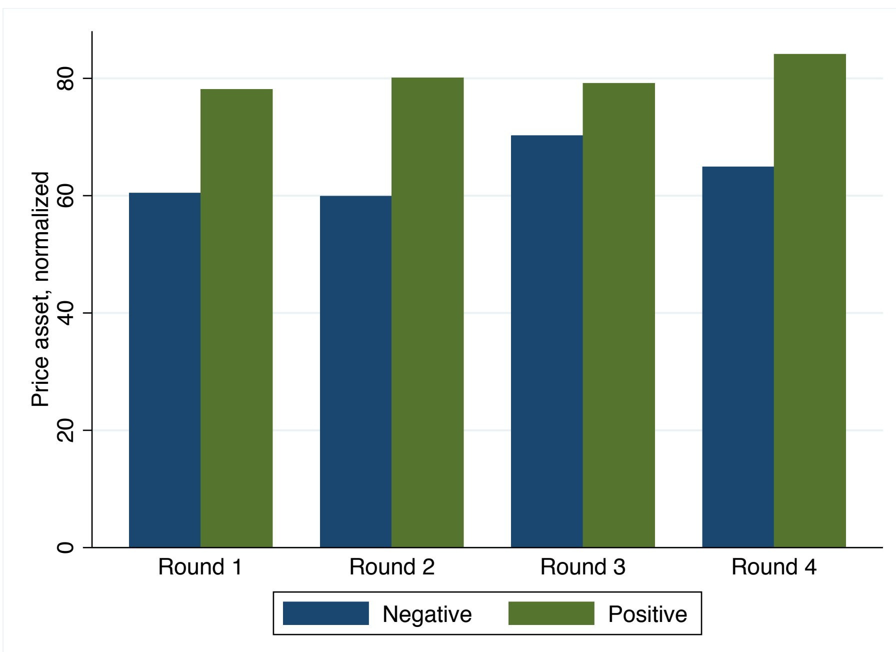

Notes: The figure gives the average asset price conditional on the signal for each of the four market rounds.

Figure A3: Development market price over time

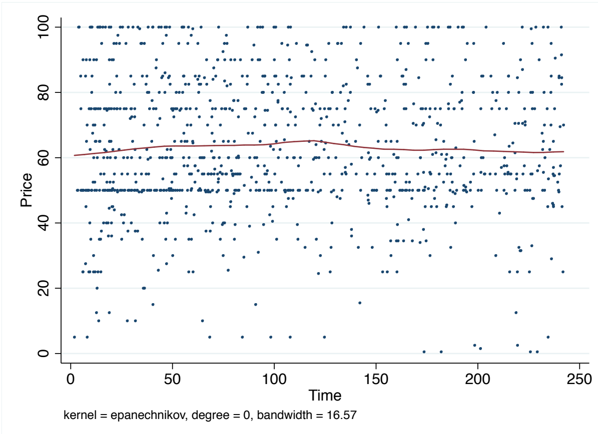

(a) Negative signal

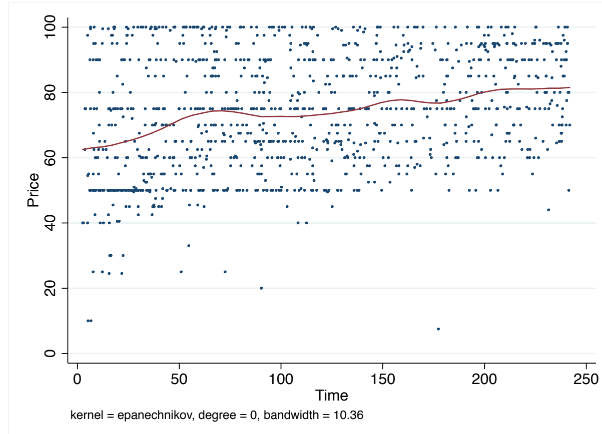

(b) Positive signal

Notes: In each market, participants traded for 4 minutes. The figure shows the correlation between asset prices and past time. We illustrate the time trend with a kernel-weighted local polynomial regression.

Figure A4: Asymmetric updating in Market condition

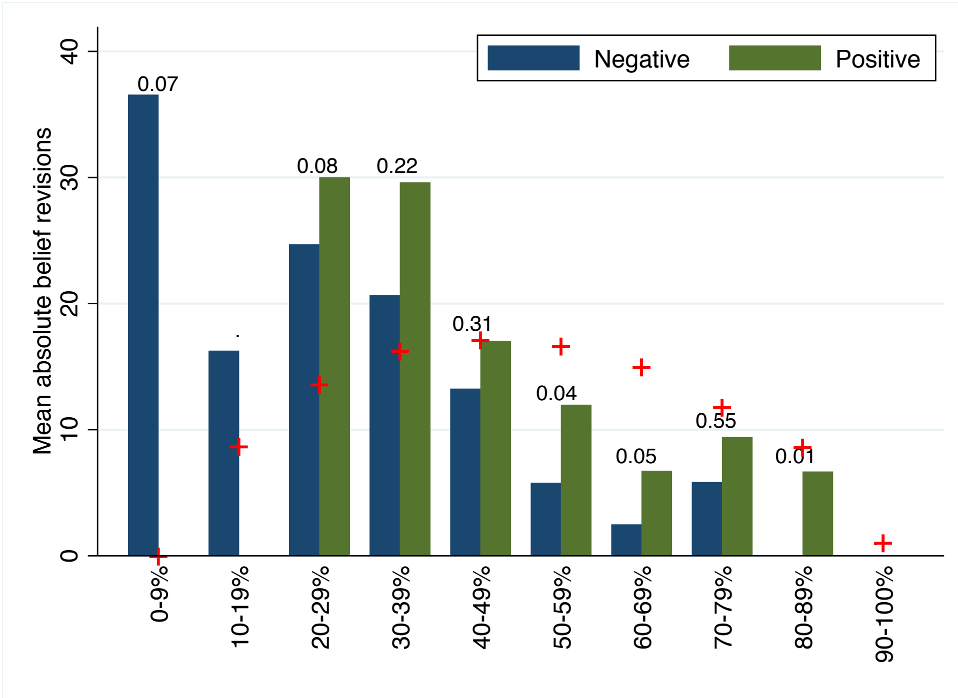

Mean absolute belief revisions by decile of prior belief in being of type equal to the signal received (following Möbius et al., 2017). The number on top of the bars indicate p-values. + indicates the rational benchmark of Bayesian updating. Observations where people updated in the wrong direction are excluded.

Appendix B: Updating in the Market condition

In this Appendix, we adapt the framework introduced in Section 2.4 to the *Market* condition. In particular, we allow the decision makers to update her belief not only in response to the signal from the urn, but also in response to the market prices. We assume that there are two possible signals from the market, a good signal and a bad signal. The decision maker makes two sequential updating decisions. First, she starts with a prior ($belief1$), and incorporates the signal from the urn, $s_{U,i}$, into her belief ($belief2$). Second, she starts from the updated belief ($belief2$) and incorporates the market signal, $s_{M,i}$, into her belief ($belief3$). We follow Möbius et al. (2017) by modeling the updating behavior with a parameterized version of Bayes rule. The first belief update is captured by:

$$logit(belief2_i) = \delta_U logit(belief1_i) + \beta_{U,L}\lambda_{U,L}\mathbf{1}(s_{U,i} = neg) + \beta_{U,H}\lambda_{U,H}\mathbf{1}(s_{U,i} = pos.) \quad (2)$$

where $\mathbf{1}(s_{U,i} = pos.)$ indicates that the signal from the urn was positive, $\lambda_{U,H} = -\lambda_{U,L} = \ln(2)$ is the log likelihood ratio of a positive signal from the urn, δ_U measures the weight placed on the prior belief and $\beta_{U,L}$ and $\beta_{U,H}$ measures how strong the decision maker react to positive and negative signals from the urn.

The second update is captured by:

$$logit(belief3_i) = \delta_M logit(belief2_i) + \beta_{M,L}\lambda_{M,L}\mathbf{1}(s_{M,i} = neg.) + \beta_{M,H}\lambda_{M,H}\mathbf{1}(s_{M,i} = pos.) \quad (3)$$

where $\mathbf{1}(s_{M,i} = pos.)$ indicates that the market signal was positive, δ_M measures the weight placed on the prior belief, $\beta_{M,L}$ and $\beta_{M,H}$ measures how much weight the decision maker puts on the market signal and $\lambda_{M,H}$ is the log likelihood ratio of a positive market signal given the signal from the urn:

$$\lambda_{M,L} = \ln\frac{\Pr(s_{M,i} = neg.|More\ points, s_{U,i})}{\Pr(s_{M,i} = neg.|Less\ points, s_{U,i})}$$

$$\lambda_{M,H} = \ln\frac{\Pr(s_{M,i} = pos.|More\ points, s_{U,i})}{\Pr(s_{M,i} = pos.|Less\ points, s_{U,i})}.$$

Note that $\lambda_{M,L}$ is zero if the market signal is perceived as non-informative conditional on the signal from the urn. If $\beta_{M,L}\lambda_{M,L} \neq 0$ ($\beta_{M,H}\lambda_{M,H} \neq 0$), the subject perceives the negative (positive) market signals as informative, and incorporates them into her belief.

Combining equations (2) and (3) results in:

$$logit(belief3_i) = \delta_M \delta_U logit(belief1_i) + \delta_M \beta_{U,L} \lambda_{U,L} \mathbf{1}(s_{U,i} = neg.) + \delta_M \beta_{U,H} \lambda_{U,H} \mathbf{1}(s_{U,i} = pos.) + \beta_{M,L} \lambda_{M,L} \mathbf{1}(s_{M,i} = neg.) + \beta_{M,H} \lambda_{M,H} \mathbf{1}(s_{M,i} = pos.) \qquad (4)$$

In the following, we want to test whether subjects react to a positive and negative market signals, that is, $\beta_{M,L}\lambda_{M,L} \neq 0$ and $\beta_{M,H}\lambda_{M,H} \neq 0$. How should we estimate the parameters of model (4)? First, we need a measure for a good and a bad market signal. We use two different measures:

- **Measure 1**: bad market signal = the normalized average market price[27] is lower than a subject's prior belief ($belief1_i$); good market signal = the normalized average market price is higher than a subject's prior belief.
- **Measure 2**: bad market signal = the normalized average market price is lower than a subject's estimated $belief2_i$ (using the parameters estimated in the *Collective* condition); good market signal = the normalized average market price is higher than a subject's estimated $belief2_i$.

A second challenge is that we can not identify the parameters of interested by only using data from the *Market* condition. We solve this issue by using data from the *Collective* condition to estimate the parameters of equation (2), δ_U, $\beta_{U,L}$ and $\beta_{U,H}$ (see Table 4). Using these estimates allows us to identify δ_M, $\beta_{M,L}\lambda_{M,L}$ and $\beta_{M,H}\lambda_{M,H}$.[28] We bootstrap estimation of equation (2) (with the *Collective* condition data) and estimation of equation (3) (with *Market* condition data) together to calculate standard errors. Table B1 gives the estimates.

We find some evidence that subjects respond to positive market signals by increasing their confidence. However, subjects ignore bad market signals. Subjects thus seem to respond asymmetrically to market prices in a way that helps them to stay (over-)confident. Another interesting finding is that base-rate neglect also applies to the second update, that is, $\delta_M < 1$. This suggests that a potential explanation for the difference in δ_t between the *Market* and the

[27] The normalized market price is the price to which a risk-neutral trader would buy the asset if she believed that the probability of winning the competition was higher than the price and would sell the asset if she believed that it was lower than the price.

[28] We can not differentiate between β_M and λ_M. However, this is not necessary to test whether subjects react to a positive and negative market signals, that is, $\beta_{M,L}\lambda_{M,L} \neq 0$ and $\beta_{M,H}\lambda_{M,H} \neq 0$.

Collective condition (see Table 4) is that subjects update twice in the *Market* condition ($\delta_t = \delta_M \delta_U$) and therefore put less weight on the initial prior. Indeed, multiplying $\hat{\delta}_U = 0.642$ (as estimated in the *Collective* condition data, see Table 4) with $\hat{\delta}_M = 0.747$ (as estimated in the *Market* condition data, see Table B1) results in $\hat{\delta}_t = 0.480$ for the sample with a prior belief in [20,80]. This is similar to the estimate we find when we estimate the parameters of model (1) in the *Market* condition (see Table 4, $\delta_t = 0.455$).

Table B1: Updating in response to the market signal

	Prior in (0,100)	**Prior in [10,90]**	**Prior in [20,80]**	**Prior in [30,70]**	**Prior in [40,60]**	**Sub-sample 1**	**Sub-sample 2**	**Sub-sample 3**
				Measure 1				
$\hat{\boldsymbol{\delta}}_{\boldsymbol{M}}$	0.825***	0.708***	0.747***	0.740***	0.707***	0.907***	0.771***	0.785***
	(0.124)	(0.062)	(0.065)	(0.073)	(0.095)	(0.137)	(0.085)	(0.084)
$\hat{\boldsymbol{\beta}}_{\boldsymbol{U,L}}\hat{\boldsymbol{\lambda}}_{\boldsymbol{U,L}}$	-0.057	0.000	-0.017	-0.044	0.053	-0.005	0.016	-0.023
	(0.088)	(0.081)	(0.085)	(0.079)	(0.068)	(0.099)	(0.096)	(0.101)
$\hat{\boldsymbol{\beta}}_{\boldsymbol{M,H}}\hat{\boldsymbol{\lambda}}_{\boldsymbol{M,H}}$	0.113	0.141***	0.129***	0.118**	0.108**	0.065	0.141**	0.106**
	(0.078)	(0.042)	(0.042)	(0.048)	(0.044)	(0.082)	(0.055)	(0.047)
N	453	448	441	394	303	377	338	334
				Measure 2				
$\hat{\boldsymbol{\delta}}_{\boldsymbol{M}}$	0.841***	0.723***	0.758***	0.770***	0.716***	0.917***	0.783***	0.801***
	(0.126)	(0.060)	(0.064)	(0.072)	(0.092)	(0.138)	(0.087)	(0.081)
$\hat{\boldsymbol{\beta}}_{\boldsymbol{U,L}}\hat{\boldsymbol{\lambda}}_{\boldsymbol{U,L}}$	-0.058	0.024	0.036	-0.016	0.112	-0.014	0.051	0.017
	(0.134)	(0.120)	(0.106)	(0.102)	(0.076)	(0.139)	(0.137)	(0.128)
$\hat{\boldsymbol{\beta}}_{\boldsymbol{M,H}}\hat{\boldsymbol{\lambda}}_{\boldsymbol{M,H}}$	0.094	0.119***	0.101**	0.098**	0.095**	0.060	0.117**	0.080*
	(0.068)	(0.043)	(0.045)	(0.047)	(0.044)	(0.071)	(0.049)	(0.044)
N	453	448	441	394	303	377	338	334

*Note: Estimated coefficients of model (1). Observations where people updated in the wrong direction or with a prior of 0 percent of 100 percent are excluded are excluded. "Prior in [10,90]" means that the sample is restricted to observations with a prior in between 10 percent and 90 percent. Subsample 1 restricts the sample to subjects that never updated in the wrong direction (in the 4 rounds). Subsample 2 restricts Subsample 1 to subjects who updated their beliefs at least once (the subsample explored in Möbius et al. (2017)). Subsample 3 restricts Subsample 2 to observations with a prior in [20,80]. Bootstrapped standard errors (over the two-stage estimation), clustered at matched-individual/group level; Standard errors in parentheses; * - p < 0.1; ** - p < 0.05; *** - p < 0.01.*

Appendix C: Additional figures

Figure C1: Task interface Individual condition

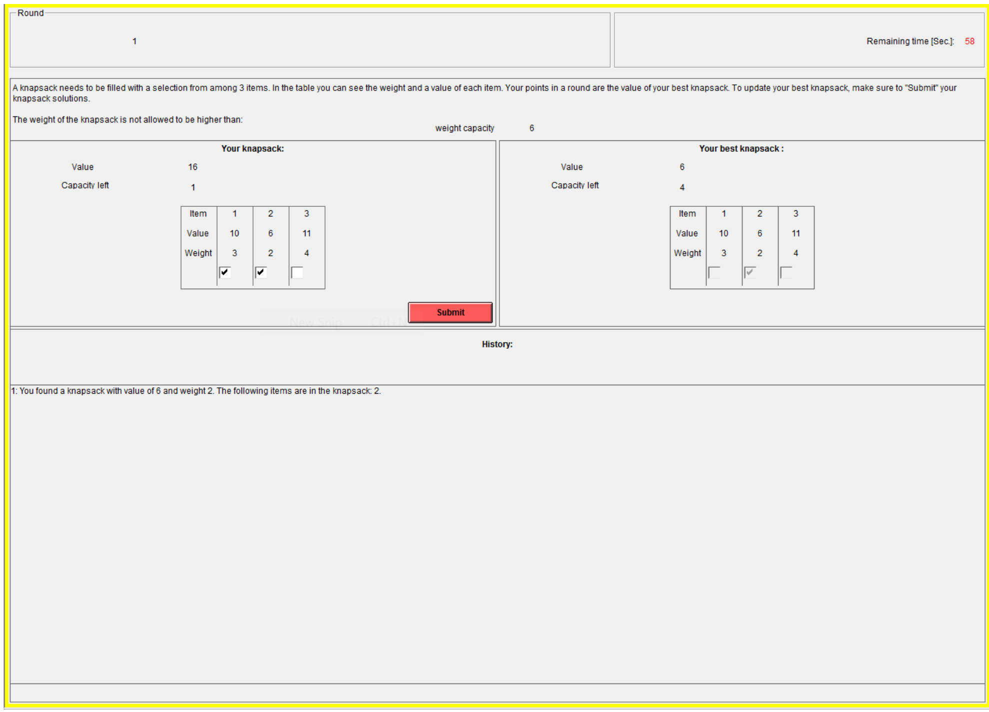

Figure C2: Task interface Collective and Market condition

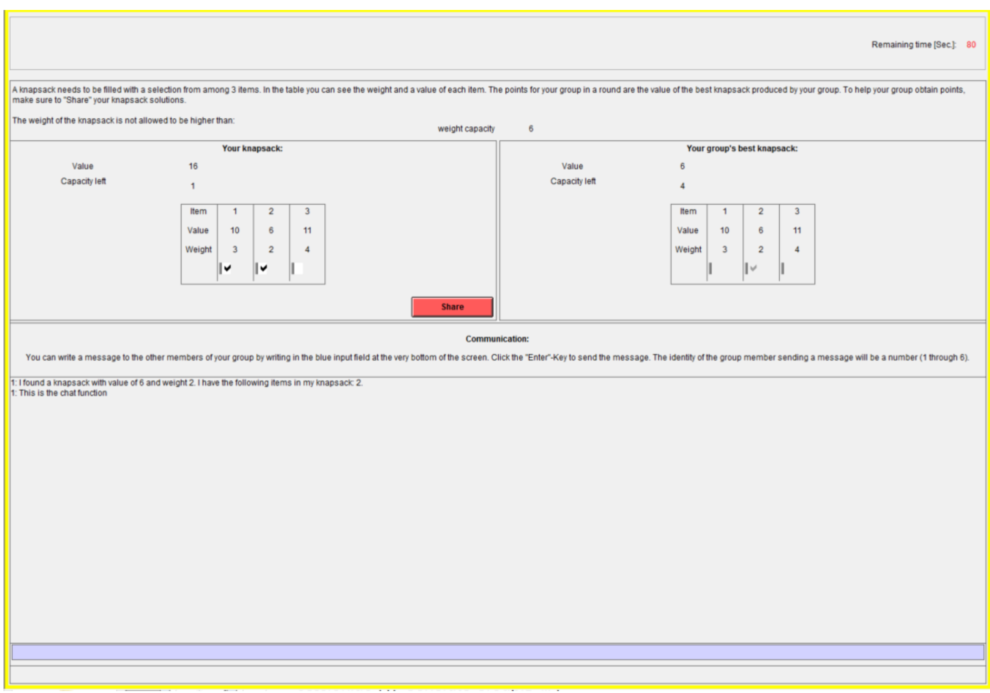

Figure C3: Market interface