



Higher-order learning

Piotr Evdokimov¹ · Umberto Garfagnini² 

Received: 5 November 2020 / Revised: 6 December 2021 / Accepted: 7 December 2021 /
Published online: 26 December 2021
© The Author(s) 2021

Abstract

We design a novel experiment to study how subjects update their beliefs about the beliefs of others. Three players receive sequential signals about an unknown state of the world. Player 1 reports her beliefs about the state; Player 2 simultaneously reports her beliefs about the beliefs of Player 1; Player 3 simultaneously reports her beliefs about the beliefs of Player 2. We say that beliefs exhibit *higher-order learning* if the beliefs of Player k about the beliefs of Player $k - 1$ become more accurate as more signals are observed. We find that some of the predicted dynamics of higher-order beliefs are reflected in the data; in particular, higher-order beliefs are updated more slowly with private than public information. However, higher-order learning fails even after a large number of signals is observed. We argue that this result is driven by base-rate neglect, heterogeneity in updating processes, and subjects' failure to correctly take learning rules of others into account.

Keywords Higher-order expectations · Learning · Theory of mind

JEL Classification C90 · D83 · D89

1 Introduction

Beliefs about the beliefs of others arise naturally in strategic settings and feature prominently in coordination problems (Keynes, 1937; Morris & Shin, 2002), global games (Carlsson & Van Damme, 1993), and financial markets (Morris & Shin, 1998). While a number of papers have investigated higher-order beliefs both theoretically and experimentally, the question of how accurate they are and how

The replication material for the study is available at <https://doi.org/10.17605/osf.io/jp5su>.

✉ Umberto Garfagnini
u.garfagnini@surrey.ac.uk

¹ Department of Applied Economics, Higher School of Economics, Moscow, Russia

² School of Economics, University of Surrey, Guildford, UK

they are updated in response to new information has received little empirical attention. This question has crucial implications for equilibrium concepts under incomplete information like Bayesian Nash equilibrium (Harsanyi, 1967; Mertens & Zamir, 1985) and cursed equilibrium (Eyster & Rabin, 2005), which assume that players know the updating rules used by others, as well as dynamic coordination problems. Cripps et al., (2008), for instance, show that coordination may be impossible in absence of common learning, which they define as the event when the true state becomes approximate common knowledge. In this paper, we use an experiment to provide a first step toward exploring how accurate higher-order beliefs are, how their accuracy changes as more information about the fundamentals is received, and whether common learning is possible.

The experiment proceeds as follows. In the beginning of the session, subjects are randomly matched into groups of three. Before any decision is made, an unknown state of the world is drawn at random and held fixed for 30 periods. It is common knowledge that the state is fixed for 30 periods, the same for every player in the group, and equally likely to take on one of two values. In each period, each group member observes a new signal about the state of the world, as in standard belief updating tasks (e.g. Holt & Smith, 2009). In the **public treatment**, all players in the same team observe the same signal in each period of the game. In the **private treatment**, Player 1, Player 2, and Player 3 observe conditionally independent signals from the same distribution, and each player can only observe her own signal.

After receiving her signal, Player 1 is incentivized to choose an action as close as possible to the realized state, Player 2 is incentivized to choose an action as close as possible to Player 1's action, and Player 3 is incentivized to choose an action as close as possible to Player 2's action. Subjects receive no feedback about the behavior of their matched partners for the duration of the experiment. After 30 periods, the state is revealed. The experiment is designed so that the action of Player 1 corresponds to her elicited first-order belief about the state, the action of Player 2 corresponds to her second-order belief about the belief of Player 1, and the action of Player 3 corresponds to her third-order belief about the belief of Player 2.

We use these treatments to address the basic question of whether players engage in higher-order reasoning, test predictions about the accuracy of higher-order beliefs, and study how their accuracy changes as more information is received.

Our first prediction is that higher-order beliefs are closer to the prior when information is private, regardless of the number of received signals. This is because when information is private, Player k for $k > 1$ must account for the fact that her information may be different from that of Player $k - 1$ by forming a belief closer to the prior, relative to what it would be if information was public. Because the only difference in the tasks of the higher-order players across the public and private treatments is information about signals received by other players, results in line with this prediction allow us to conclude that Players k for $k > 1$ engage in higher-order reasoning.

Second, higher-order beliefs are predicted to be more accurate on average with public than private signals. In the Bayesian benchmark, beliefs of Player k and Player $k - 1$ are identical in the public treatment, but not in the private treatment, where the fact that other players observe potentially different sequences of signals must be accounted for. Even if subjects are non-Bayesian, higher-order beliefs are

more difficult to update in the private treatment because the potential difference in histories must be taken into account.

In addition to testing predictions about how higher-order beliefs are updated, we study higher-order learning, i.e., the evolution of the accuracy of higher-order beliefs as more signals are received. As we elaborate below, failure to correctly predict the beliefs of others might be attenuated in the long term. In the Bayesian benchmark, this is true in the private treatment, where higher-order beliefs are inaccurate in early periods but become more accurate as both higher- and lower-order beliefs converge to the truth. Even if subjects are non-Bayesian and follow heterogeneous updating processes, higher-order learning might be observed, depending on how much heterogeneity is present in the data and how good subjects are at forecasting the beliefs of others. We use the experiment to address this question empirically.

We find that the first prediction is in line with the data, while the second is not. I.e., subjects account for the public vs. private nature of information, but higher-order beliefs are not more accurate when the information is public than when it is private. Moreover, the accuracy of higher-order beliefs does not improve over time in either treatment, even as a large number of signals is received; i.e., higher-order learning fails.

We argue that the observed failure of higher-order learning in both treatments is rooted in failures of Bayesian thinking (e.g., base rate neglect), heterogeneity in information processing, and subjects' failure to take this heterogeneity into account. Base-rate neglect has theoretically been shown to bound beliefs away from the correct state of the world (Benjamin et al., 2019). As first-order beliefs fail to converge despite accumulating evidence, heterogeneity in updating rules implies that different subjects will have very different long-run beliefs. This, in turn, implies that higher- and lower-order beliefs fail to converge.

To study what assumptions subjects make about the updating rules of others, we run additional within-subjects treatments where each subject reports a belief *both* in the role of Player 1 and Player 2. We find that the vast majority of subjects show a median difference of zero between their first- and second-order beliefs. We also use a counterfactual exercise to show that even if subjects reported the optimal beliefs *given the updating rules used by other subjects in the experiment*, higher- and lower-order beliefs would fail to converge. In other words, even a player with knowledge about the distribution of subjects' updating types in the experiment would fail to show higher-order learning.

Finally, we address the question of whether the observed failure of higher-order learning can be mitigated by additional information. To this end, we use the counterfactual exercise to simulate higher-order beliefs in an experiment where 300 as opposed to 30 signals are observed. We find that higher-order beliefs initially diverge and eventually plateau. Thus, we find little benefit of receiving more signals in the counterfactual exercise. To test this prediction, we collect data from an additional online experiment in which subjects receive 10 signals in every period, for a total of 300 signals in period 30. We find that higher-order beliefs in this treatment are not significantly different between the first 15 and the last 15 periods. These results are in line with the predictions of the counterfactual exercise. Overall, the

results suggest that higher-order learning is difficult to achieve, which in turn raises questions about the feasibility of common learning.

1.1 Related literature

This paper complements several strands of literature. Cripps et al., (2008) define common learning as the event where the true state becomes approximate common knowledge and provide conditions on the signal distributions of each player which guarantee that common learning is feasible. While other theoretical papers have followed the research agenda of Cripps et al., (2008) (e.g. Wiseman, 2012; Acemoglu et al., 2016), little is known about whether common learning occurs in practice. Because elicitation of an infinite hierarchy of beliefs poses an obstacle to any laboratory study, we restrict our attention to higher-order learning (i.e., increasing accuracy of higher-order beliefs). Higher-order learning is a necessary if not sufficient condition for common learning, so that a failure of higher-order learning in the laboratory would cast doubt on the feasibility of common learning, as well.

A number of papers have investigated how subjects update their beliefs in response to new information (e.g., Grether 1980, 1992; Holt & Smith, 2009) and found violations of Bayes' rule as well as substantial heterogeneity in updating rules. More recent contributions also find that decision makers process information about personal characteristics such as IQ asymmetrically, overweighting good news and underweighting bad news (e.g., Eil & Rao, 2011; Mobius et al., 2014; Coutts, 2018). In contrast to prior individual learning experiments, we investigate the updating of higher-order beliefs, expanding the existing literature to strategic settings.

We also contribute to the growing experimental literature on higher-order reasoning. Nagel (1995) introduces level- k thinking in the context of guessing games and finds that few subjects exceed two levels of reasoning.¹ Huck & Weizsäcker (2002) elicit subjects' beliefs about the lottery choices of other subjects. Although subjects are able to correctly predict the choice frequencies of other subjects on average, they find a significant and systematic bias toward a uniform prior. Kübler & Weizsäcker (2004) study how subjects process information generated through their predecessors' choices in a social learning framework. Using an error-rate model which allows to estimate how subjects reason about other subjects' behavior, they find that subjects underestimate the rationality of their immediate predecessors (similar to what found by Weizsäcker, 2003) and that the average subject's reasoning does not exceed two steps.²

¹ Early contributions to the level- k literature also include Stahl and Wilson (1994, 1995).

² Using a novel experimental design, Kneeland (2015) recovers the implied level of higher-order rationality from subjects' choices and finds higher levels of rationality than prior studies. Our stage game is similar in spirit to Kneeland's in that it links a subject's payoff function to that of her predecessor in a directed graph. Unlike Kneeland's, the payoff of the first player in our experiment depends on an unknown state of the world as opposed to another player's action. While Kneeland's experiment was designed to recover subjects' individual level of rationality by varying their positions in the graph, we introduce uncertainty about a state of the world and provide subjects with sequential information about the state to track the evolution of their higher-order beliefs. To the best of our knowledge, our paper is the first to investigate higher-order learning.

The difference in how subjects treat private and public information has been investigated in the experimental global games literature. Several papers find little differences in behavior across the two types of information in one-shot coordination games (Heinemann et al., 2004; Van Huyck et al., 2018), contrary to theoretical predictions. Cornand & Heinemann (2014) provide subjects with both public and private information about the underlying state of the world in a game with strategic complementarities. They argue that systematic mistakes in how subjects form higher-order beliefs can partly explain their observed deviations from equilibrium behavior. Our paper expands on this point by investigating how subjects update their higher-order beliefs in response to new information. We argue that beliefs about the beliefs of others may be persistently incorrect. Providing more information about the environment, e.g., the fundamentals of the economy in a global game setting (Angeletos et al., 2007), may prove inconsequential for the accuracy of higher-order beliefs, resulting in persistent mistakes in choice.

2 Experimental design

Our experimental design borrows sequential belief elicitation over a binary state space from the belief updating literature. The setup is intentionally simple. At the beginning of each session, the subjects are matched in teams of three. Within a team, each subject is randomly assigned to one of three roles: Player 1, Player 2, or Player 3, with exactly one subject in each role.³ The roles and teams stay fixed for the duration of the session. A session consists of three incentivized rounds, with each round unfolding as described below.⁴

Subjects are told that there are two urns, *Orange* and *Purple*, each containing 3 balls. The *Orange* urn contains 2 orange balls and 1 purple ball, while the *Purple* urn contains 1 orange ball and 2 purple balls. Before a round begins, the computer selects one of the two urns with equal probability for each three-player team. None of the subjects are told which urn is selected for their team.⁵ A round consists of 30 periods.⁶

³ In the initial treatments, each subject takes one action, as opposed to three actions, to avoid possible contamination between lower- and higher-order expectations. Nevertheless, in additional treatments, we elicit both first- and second-order expectations from the same subjects to shed light on the subjects' theory of mind.

⁴ We used three rounds in the experiment to allow subjects to become better acquainted with the task. Due to an experimenter error, the first session of the experiment was programmed for five instead of three rounds. For that session, we only analyze the first three rounds to be consistent with all the remaining sessions.

⁵ Thus, a new urn is independently drawn at the beginning of the next round.

⁶ While 30 repetitions is a large number for an individual updating task, our experiment focuses on the elicitation of higher-order expectations. We chose to have many periods to allow subjects to better refine their beliefs about the state and, most importantly, their beliefs about the beliefs of others. While sequentiality leaves the results open to behavioral biases documented in experiments on individual beliefs, it is also crucial for an investigation of common learning.

In the **public treatment**, the computer draws a ball with replacement from the selected urn in every period, and shows the ball to all subjects in the same team. I.e., the players receive public signals about the color of the urn. In the **private treatment**, the computer draws a ball from the selected urn in every period *separately for each subject*. The subject is shown the color of her drawn ball but not the color of the ball drawn for her matched partners. I.e., the signals received by the subjects are private and conditionally independent.

In every period, Player 1 reports her belief about the color of the urn (i.e., the state of the world), Player 2 reports her belief about the belief of Player 1, and Player 3 reports her belief about the belief of Player 2. Each subject makes one guess. The experiment is framed neutrally in that it avoids any reference to guesses or beliefs, instead explaining the task as a betting problem.

To avoid the influence of risk aversion on subjects' elicited beliefs, we employ the Binarized Scoring rule of Hossain & Okui (2013), which is incentive compatible irrespective of attitudes toward risk and relatively simple to implement.⁷ The rule is applied to elicit Player k 's beliefs about the underlying random variable of interest individually in every period of a round. The rule works as follows:

1. Player k takes an action $a_k \in [0, 1]$;⁸
2. A random variable of interest, Z , is realized;
3. The player's loss is computed according to a loss function $L(a_k, z)$, where z is the realization of Z ;
4. The computer draws a number c uniformly at random from the interval $[0, 1]$;
5. If $L(a_k, z) \leq c$, Player k receives a monetary reward of R_1 , otherwise the player receives $R_0 < R_1$.

We employ a quadratic loss function $L(a, z) = (a - z)^2$. For Player 1, Z is either 1, if the selected urn is *Orange*, or 0 otherwise. For Player $k \in \{2, 3\}$, Z corresponds to the action chosen by Player $k - 1$.

Each subject is paid on the basis of one randomly chosen period of a randomly chosen round. Paying for only one randomly chosen period breaks any intertemporal hedging across periods (and rounds), turning each period of a round into a static task.⁹ Thus, regardless of attitudes to risk, the optimal response of Player 1 is to report the probability that he/she assigns to the color of the urn being *Orange*. For Player 2 and Player 3, the optimal response corresponds to the expectation of the action chosen by the preceding player, that is, $a_k = E[a_{k-1} | \mathcal{I}_k]$, where \mathcal{I}_k is Player k 's

⁷ Danz et al., (2020) find a pull-to-the-center effect on beliefs elicited with the binarized scoring rule, which we were unaware of when we ran the experiment. While this effect could in principle affect our results, it would impact lower- and higher-order beliefs equally, improving the accuracy of higher-order beliefs and reducing the failure of higher-order learning in the data. We leave it to future research to test whether belief accuracy can be improved with the use of other scoring rules.

⁸ The action set in the experiment was a discretized version of $[0, 1]$ with $A_k = \{0, 0.01, 0.02, \dots, 0.98, 0.99, 1\}$.

⁹ Azrieli et al., (2018) show theoretically that selecting one task at random is the only incentive compatible way to pay subjects under a monotonicity assumption on subjects' preferences.

information set, $k = 2, 3$.¹⁰ Thus, Player 2's action corresponds to her expectations about Player 1's beliefs about the state of the world (that is, her *second-order expectations*), whereas Player 3's action corresponds to her expectations about Player 2's expectations about Player 1's beliefs about the state of the world (that is, her *third-order expectations*).

Note that two subjects in the role of Player k for $k > 1$ might have different beliefs about the beliefs of Player k that nevertheless have the same mean. In the analysis that follows, we measure belief accuracy using the elicited means. According to this measure, two players might have the same accuracy of higher-order beliefs according to our measure despite the fact that one of them has very precise beliefs (e.g., concentrated on 0.5) while the other has beliefs that are more diffuse (e.g., a uniform distribution over 0 and 1). We leave the extension of our paper to elicited distributions of beliefs to future research.

While Player 2's task involves only Player 1, Player 3's task involves both Player 1 and Player 2.¹¹ Subjects in the role of Player 1 are told that they will be matched with a subject in the role of Player 2 and a subject in the role of Player 3 but that the decisions made by those players will be inconsequential for her/his own performance. Moreover, Player 1 is not told what tasks Player 2 and Player 3 are given. Player 2 is explained Player 1's task but not Player 3's task. Player 3 is the only player with full information about the structure of all tasks.¹²

Subjects receive no feedback about their own performance or the performance of their matched partners for the entire duration of the experiment. At the end of each round, the correct composition of the urn is revealed to all subjects in the same team and no other information is disclosed.¹³ Lack of feedback is common in experiments measuring subjects' beliefs about other subjects' beliefs (see e.g., Stahl & Wilson, 1995; Costa-Gomes et al., 2001; Costa-Gomes & Crawford, 2006; Costa-Gomes & Weizsäcker, 2008), and we refrain from providing subjects with feedback for two main reasons. First, our experiment tries to identify a subject's mental model of other subjects and the possible effect of introspective learning rather than her response to reinforcement learning. Second, it would not be possible to implement the private treatment with period-to-period feedback. Observing a partner's action

¹⁰ A different approach could have been to elicit the distribution of Player 2 and Player 3's beliefs rather than the mean of those distributions. We did not follow this approach to keep each subject's task as simple as possible. Manski & Neri (2013) elicit both the mean and the distribution of subjects' second-order beliefs in a Hide-and-Seek game and find a general consistency between the two, that is, the mean of the distribution of second-order beliefs is consistent with its point estimate.

¹¹ Note, however, that Player 3's payoff only depends on the action of Player 2 and not Player 1.

¹² Notice that this makes social preferences theoretically irrelevant.

¹³ At the end of a round of the private treatment, a subject is also shown the cumulative number of balls drawn for each of her matched partners, but still no information about her partners' guesses. This is chosen to further emphasize to a subject the privacy of signals and that her matched partners can observe different histories than her own. While doing this provides additional information to a subject about likely histories, we believe this issue to be minor as we only give a subject a snapshot of the cumulative number of balls of each color seen by her partners at the end of the round rather than the sequence of balls drawn from the urn in each period, which we believe it would be more likely to inform decision making in subsequent rounds.

would reveal information about that partner's private signals, thus affecting a subject's choice of action in subsequent periods.

2.1 Predictions

Consider the Bayesian benchmark where players are rational, believe that others are Bayesian and rational, and believe that others believe that others are Bayesian and rational. In the public treatment, this benchmark predicts that the beliefs of Players 1–3 coincide after any possible history of observed signals. In the private treatment, uncertainty about the information of others in the private treatment creates a wedge between the average beliefs of Player k and the average beliefs of Player $k - 1$ for all $k \geq 2$, conditional on the true state of the world (see Fig. 1 and Appendix B for the proof). The logic is as follows. If Player k knew which history Player $k - 1$ has observed, her action would correspond to Player $k - 1$'s action. Private signals imply that Player $k - 1$ might have observed a different history. Player k must therefore use her own observed history to compute the distribution of signal histories observed by Player $k - 1$. Thus, uncertainty slows down the evolution of Player k 's beliefs about Player $k - 1$'s beliefs in expectation, as Player k has to give positive weight to beliefs that Player $k - 1$ is likely to hold with only small probability. Over time, the probability of such histories becomes vanishingly small and Player k 's beliefs converge to 1, which is also the limit of Player $k - 1$'s beliefs. This implies the following prediction:

Prediction 1 *Higher-order beliefs are closer to the prior with private information.*

Note that if a shift from public to private information generates a change in higher-order beliefs, we can conclude that the subjects are engaging in higher-order reasoning.¹⁴

Define belief accuracy of a subject in role $k > 1$ in a given period as one minus the absolute distance between the subject's reported belief and the reported belief of the subject's matched partner.¹⁵ Because matched partners observe identical signal histories in the public but not the private treatment, the following prediction follows:

Prediction 2 *On average, higher-order beliefs are more accurate in the public treatment, regardless of the number of signals observed.*

We define higher-order learning as *increasing accuracy of higher-order beliefs*. The beliefs of Bayesian players are predicted to be perfectly accurate in the public treatment regardless of how many signals are received. On the other hand, based on the results of previous studies, we should expect laboratory subjects to deviate from

¹⁴ While the treatments were not designed to identify levels of rationality, as in Kneeland (2015), we can use them to test the simpler hypothesis of whether subjects engage in higher-order reasoning at all.

¹⁵ Formally, belief accuracy is equal to $1 - |a_k - a_{k-1}|$, for $k = 2, 3$.

Bayesianism.¹⁶ In the presence of heterogeneity of updating processes, accuracy of higher-order beliefs depends on the extent to which deviations from Bayesian updating are forecasted and shared.

To illustrate this point, Fig. 2 shows the predicted evolution of the average distance between first- and second-order beliefs in the public treatment under varying assumptions. In all cases, the population consists of a mix of Bayesians and non-Bayesian λ -types.¹⁷ For a λ -type, the posterior belief in every period is the Bayesian belief, given the subject's prior belief and current signal, with weight λ and the prior with weight $1 - \lambda$.¹⁸

We assume that Bayesian players believe that others are Bayesian and that others believe that others are Bayesian; λ -types believe that others are λ -types, and that others believe that others are λ -types.

The dotted line, for which the average belief accuracy is closest to one, represents the predictions of a model in which the parameters λ are drawn from a uniform distributions over $\lambda = 1$, $\lambda = 0.9$, and $\lambda = 0.8$ for both first- and second-order beliefs. Because all players are close to being fully Bayesian, both first- and second-order beliefs quickly converge to the truth, the distance between them converges to zero, and belief accuracy converges to its maximal value.

The solid line represents the predictions of a model in which the parameters λ are drawn from a uniform distributions over $\lambda = 1$, $\lambda = 0.55$, and $\lambda = 0.1$. In this case, the population of players consists of Bayesian learners, slow learners, and an intermediate type, and higher-order learning is considerably slower.

Higher-order learning can be facilitated by forming a correct mental model of others' updating behavior. Thus, the dashed line represents the predicted accuracy of *optimal* second-order beliefs, if actual beliefs are drawn from a uniform distribution over $\lambda = 1$, $\lambda = 0.55$, and $\lambda = 0.1$, and the optimal higher-order beliefs are the expected lower-order beliefs given the distribution of types. Note that the dashed line is above the solid line, capturing the intuition that higher-order beliefs are more accurate if the deviations from Bayesian updating are correctly forecasted.

The assumptions underlying the predictions in Fig. 2 are ad hoc and made only for illustrative purposes; ultimately, the question of how much heterogeneity is present in the data and how well deviations from Bayesianism are forecasted is an empirical one. Answering this question is one of the goals of our experiment.

¹⁶ Holt & Smith (2009), for instance, show that belief updating processes are subject to significant heterogeneity.

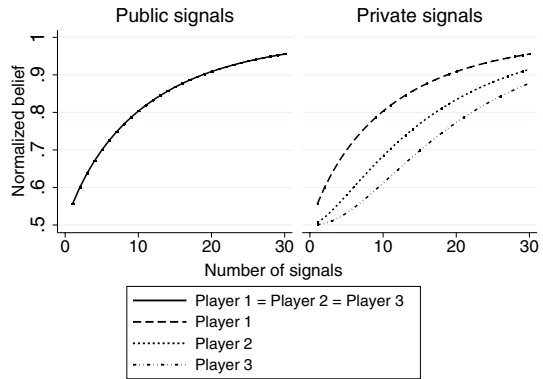
¹⁷ All of our simulations are based on a sample of 10,000 randomly drawn and matched players.

¹⁸ This model is borrowed from Epstein et al., (2008), who axiomatized a non-Bayesian updating rule which includes Bayes' rule as a special case. Formally, fix a signal history $h_n = \{s_1, s_2, \dots, s_n\}$, and let μ_{n-1} denote a player's prior belief about the state of the world in period n , before having observed the current signal s_n . Following Epstein et al., (2008), we assume that a λ -type forms the following posterior belief in period n

$$P_1(h_n) = \lambda P_1^{Bayes}(\mu_{n-1}, s_n) + (1 - \lambda)\mu_{n-1} \quad (1)$$

that is, a subject's posterior belief is a weighted average of her one-step Bayesian posterior belief and her prior belief. $P_1^{Bayes}(\mu_{n-1}, s_n)$ is the Bayesian posterior which results from updating the prior μ_{n-1} given the current signal s_n .

Fig. 1 The predicted evolution of expected first-, second-, and third-order beliefs. The beliefs are normalized by the correct state so that the variable being plotted is B when the state is orange and $1 - B$ when the state is purple, where B is the belief of a Bayesian decision maker



2.2 Implementation

The experiment was conducted at Instituto Tecnológico Autónomo de México in Mexico City between October and December 2017 using the software z -Tree (Fischbacher, 2007). Data were collected from 120 subjects in 7 sessions for the public treatment and from 129 subjects in 8 sessions for the private treatment. A session lasted 75 minutes on average. All subjects were undergraduate students recruited from the general student population. Each subject could only participate once.

Each session started with subjects signing the consent forms, reading the instructions, and completing an incentivized quiz.¹⁹ Every subject was guaranteed a 100 Mexican pesos show-up fee (\approx US\$5.26 at the time of the experiment) in addition to the earnings from the quiz (2 Mexican pesos for each correct answer). These earnings were called the subject’s “guaranteed earnings.” Each subject was also given an initial endowment of 80 pesos which the subject had a chance to either double or lose completely according to the following procedure based on the binarized scoring rule.²⁰ The computer randomly selected one period of play for each subject. Given a subject’s loss for the period from her decision, the computer independently drew a number c that was uniformly distributed between 0 and 1, and the subject’s “additional earnings” were determined as follows:

$$\text{Additional earnings} = \begin{cases} 2 * \text{Initial endowment,} & \text{if } (a - z)^2 \leq c, \\ 0, & \text{otherwise.} \end{cases} \quad (2)$$

¹⁹ The instructions can be found in the online appendix. While the sample instructions are in English, the actual instructions were administered in Spanish. The answers to all of the quiz questions were incentivized.

²⁰ At the time of the experiment, the minimum wage in Mexico was about 70 pesos per day, which is arguably a poor reference point for students at a private research university such as ITAM. For a better one, consider that the cost of a 15km Uber ride was around 80 pesos.

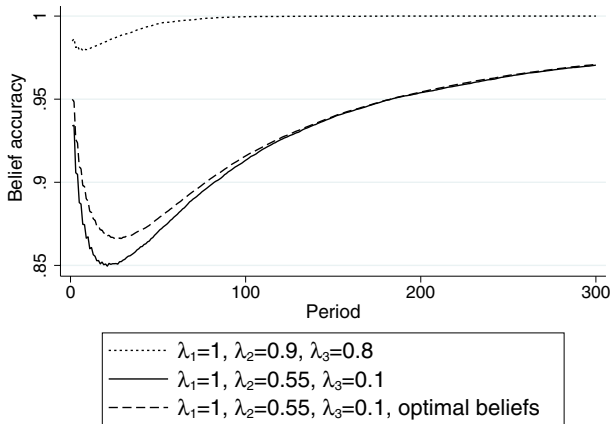


Fig. 2 The simulated evolution of expected accuracy of higher-order beliefs in the public treatment. Each line represents the average accuracy in a simulated population of players. The population in each case consists of an equal mix of three types of players, with λ_i indexing the updating rule used by type i as described in the text

The payment rule was clearly explained to the subjects in the instructions, and several examples were provided.²¹

Our presentation of the experimental results is structured as follows. Section 3 contains our main results on the effect of private vs. public information on higher-order beliefs, belief accuracy, and the failure of higher-order learning. Section 4 presents the results of additional treatments in which first- and higher-order beliefs are elicited in a within-subject design. These treatments replicate several of the main findings from Sect. 3 and shed light on subjects' theory of mind. Finally, Sect. 5 investigates possible reasons for the observed failure in higher-order learning, highlighting the impacts of base-rate neglect and heterogeneity in updating rules. We also report the results of an additional treatment in which subjects observe up to 300 signals, as opposed to 30 in the other treatments.

3 Main results

Result 1 *Higher-order beliefs are closer to the prior with private information, suggesting that players engage in higher-order reasoning.*

Figure 3 shows the average reported beliefs of subjects in all player roles and treatments. The beliefs are normalized by the correct state so that the variable being plotted is B when the state is orange and $1 - B$ when the state is purple, where B is

²¹ The software also provided an initial screen that a subject could use to experiment with the calculation of the loss based on the distance between a hypothetical guess (entered by the subject) and possible target numbers.

the subject's reported belief. For each period, the normalized beliefs are averaged across all subjects in the given treatment and player role, as well as all observed signal histories.

Consistent with the predictions, we find that higher-order expectations are closer to prior beliefs with private than public signals. This suggests that subjects understand that information of others differs from their own.²² Thus, when the normalized expectations of Players 2 and 3 are regressed against a dummy variable for the treatment with private signals, the private dummy in this regression is negative and significant ($P < 0.01$; first column of Table 1). It remains significant if we control for period number and the interaction between period number and the private treatment ($P < 0.05$; second column of Table 1).

While the belief of a Bayesian Player 3 in the private treatment is closer to 0.5 than that of Player 2, we do not find evidence of such behavior in the data: in a regression of the normalized expectations of Players 1, 2 and 3 in the private treatment against a Player 2 dummy and a Player 3 dummy, the two dummy variables are not significantly different ($P = 0.883$; third column of Table 1).²³ Thus, the effect of private information on higher-order beliefs appears to be limited. One possibility is that because Player 3 faces a more difficult information processing task than Player 2 in the private treatment, her behavior is further away from best responding than that of Player 2. This, however, is not the case: as we show below, Player 2 and Player 3 both show substantial and similar deviations from best-responding to their partners.

In line with the Bayesian benchmark, Players 1, 2 and 3 report similar beliefs on average in every period of the public treatment. This can be seen in the regression results reported in the fifth column of Table 1, normalized beliefs in the public treatment are regressed against a Player 2 dummy and a Player 3 dummy; neither dummy variable is significant, with p -values of $P = 0.479$ and $P = 0.354$, respectively. The two dummy variables are also not significantly different ($P = 0.8115$).²⁴ This result is consistent with a number of possibilities. One is that the three types of players in the public treatment follow similar updating rules and are correctly-guessing their target players' beliefs on average. Another is that higher-order beliefs are inaccurate (because some subjects are over- and some under-guessing) but appear accurate on average.²⁵

²² In a different framework, Cornand & Heinemann (2014) find that higher- and lower-order beliefs treat signals about the state of the world differently, which suggests understanding of the difference between one's own and others' private information.

²³ The fourth column controls for period number and the relevant interactions. The coefficients on the two dummies are not significantly different in that specification ($P = 0.3964$), and neither are the coefficients on the two interaction terms ($P = 0.3691$).

²⁴ The regression in the sixth column of Table 1 includes a period variable and interactions with the player dummies; only the period variable is significant in that specification ($P < 0.001$).

²⁵ Consider, for instance, the possibility that every player reports Bayesian beliefs plus or minus some ϵ term, where ϵ follows a known distribution with mean zero. The distance between Player k and Player $k - 1$'s beliefs will be positive on average despite the fact that average beliefs are equal.

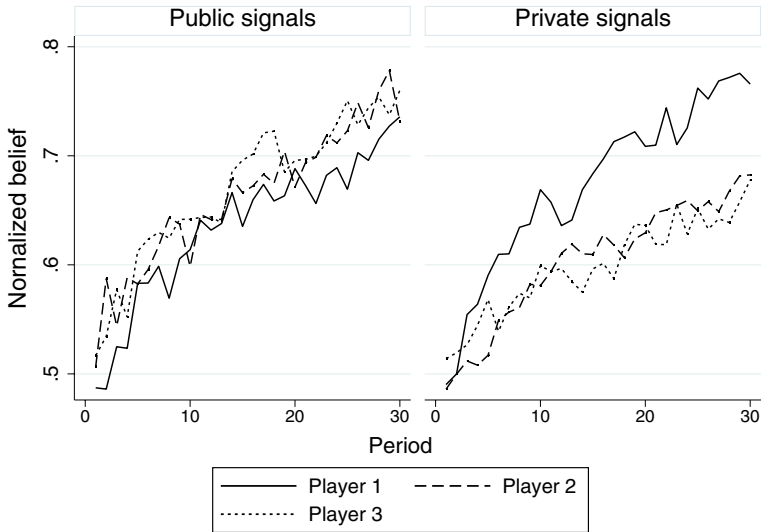


Fig. 3 The evolution of subjects' first-, second-, and third-order beliefs. The beliefs are normalized by the correct state so that the variable being plotted is B when the state is orange and $1 - B$ when the state is purple, where B is the reported belief. The normalized beliefs are averaged across all subjects and signal histories for each treatment and player role. As predicted, higher-order beliefs are closer to the prior when information is private (Result 1)

Result 2 *Higher-order beliefs are not more accurate in the public than the private treatment.*

We measure belief accuracy as $1 - |a_{it} - a_{-i,t}|$, i.e., one minus the absolute distance between a subject's beliefs and those of her matched partner. Figure 4 plots the evolution of belief accuracy over time in all relevant experimental conditions. Contrary to Prediction 2, beliefs are not more accurate with public than private information. This result is confirmed in a regression of $1 - |a_{it} - a_{-i,t}|$ against a private treatment dummy, whether it is run for Players 2 and 3 together ($P = 0.365$; first column of Table 2), Player 2 separately ($P = 0.550$; second column), or Player 3 separately ($P = 0.490$; third column).

Result 3 (Failure of higher-order learning) *Higher-order beliefs diverge from lower-order beliefs in the experiment. The period of divergence is very long; even 30 periods is not enough for convergence.*

Recall that we define higher-order learning as increasing accuracy of higher-order beliefs over time. Contrary to higher-order learning, we find no significant period trend in the public treatment ($P = 0.363$; fourth column of Table 2) and a negative period trend, suggesting decreasing belief accuracy over time, in the private treatment ($P < 0.01$).

To summarize, higher-order beliefs are as inaccurate with public information as they are when information is private and therefore more difficult to process.

Table 1 Analysis of average normalized observed expectations

	Players 2 and 3		All players, private		All players, public	
Private	-0.0651*** (0.0221)	-0.0400** (0.0174)				
Period			0.00706**** (0.000720)	0.00854**** (0.00113)		0.00710**** (0.00117)
Period × Private		-0.00162 (0.00109)				
Player 2			-0.0697* (0.0358)	-0.0330 (0.0257)	0.0264 (0.0373)	0.0277 (0.0283)
Player 3			-0.0742** (0.0343)	-0.0146 (0.0255)	0.0341 (0.0366)	0.0342 (0.0295)
Period × Player 2				-0.00237 (0.00159)		-0.0000800 (0.00166)
Period × Player 3				-0.00385** (0.00165)		-0.00000986 (0.00144)
Constant	0.666**** (0.0160)	0.557**** (0.0136)	0.673**** (0.0276)	0.541**** (0.0205)	0.636**** (0.0292)	0.526**** (0.0215)
Observations	14940	14940	11610	11610	10800	10800

Output of OLS regressions of the form $Y = \beta X + \epsilon$, where $Y = B$ when the state is orange and $Y = 1 - B$ when the state is purple, and B is the subject's reported belief. The first two columns consider Players 2 and 3 in all treatments; the third and fourth columns consider Players 1--3 in the private treatment; the last two columns consider Player 1--3 in the public treatment. Standard errors are clustered at the level of individual subjects

Subject-clustered standard errors in parentheses

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$, **** $p < 0.001$

Moreover, higher-order beliefs do not become more accurate over time in either the private treatment or the public treatment, where they are predicted to always be fully accurate by the Bayesian benchmark.

4 Within-subjects data and theory of mind

To explore in more detail what beliefs subjects form about the information processing of others, we collected data from two additional treatments, within-public and within-private, which were conducted online. These treatments are similar to the public and private treatments described above in all respects but the following. First, subjects are matched into teams of two instead of three players. Second, subjects go through one single round of 30 periods, as opposed to three rounds.²⁶ Third, and most importantly, we elicit both first- and second-order beliefs for every subjects in

²⁶ The online experiment has a single round to make it shorter, given the possibility of dropouts.

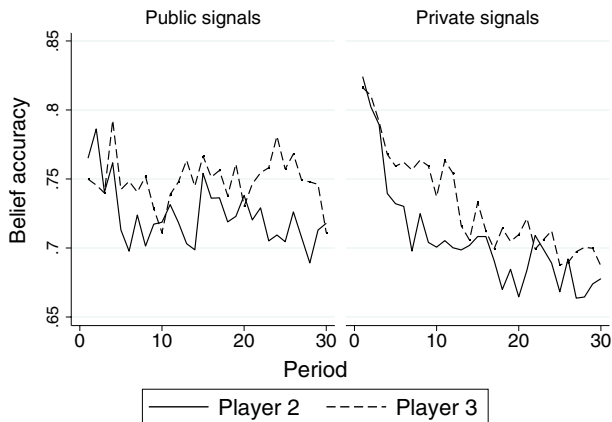


Fig. 4 The failure of higher-order learning. Accuracy of higher-order beliefs is measured by $1 - |a_{it} - a_{-it}|$. The data are plotted for different treatments, player types, and periods. Higher-order beliefs are not more accurate with public than private information and fail to become more accurate over time in either treatment

every period. This allows us to explore whether subjects assume that others process information differently than they themselves do.

The subject pool consists of U.S. workers on Amazon Mechanical Turk (MTurk), and the experiment was conducted using the software oTree (Chen et al., 2016). For the within-public treatment, we collected data from 204 subjects between June and September 2019. Data for the within-private treatment were collected from an additional 150 MTurk subjects at the request of a referee. The average hourly wage was \$14.75, which is more than three times higher than the standard MTurk task (Hara et al., 2018). Subjects took 16 minutes on average to complete the experiment. Further implementational details can be found in Appendix A.²⁷

Each subject is matched with a partner for a single round of 30 periods. In each round, each subject receives one signal about the state of the world and provides first- and second-order beliefs. A random period and belief type are drawn for payment, and the payment is determined using the binarized scoring rule, as in the between-subjects treatments. I.e., the within-public and within-private treatments are similar to their between-subjects counterparts with the difference that only one round of matching occurs and beliefs of Players 1 and 2 are elicited within-subjects.

The first two panels of Fig. 5 plot the average normalized first- and second-order beliefs in the within-public and within-private treatments. We find that first-order beliefs in the within-public and within-private treatments evolve similarly to their

²⁷ The average length of the lab sessions (75 minutes) was much longer than that of the online sessions. This is because each lab session included welcoming subjects, distributing, explaining and collecting the consent forms, distributing the instructions, actual playing time, and paying subjects at the end. Furthermore, lab subjects went through three rounds of belief elicitation, each consisting of 30 periods. Thus, the difference in session length was due both to longer playing time and logistics related to the implementation of a session in the laboratory.

Table 2 Analysis of belief accuracy

	Lab				MTurk	
	Players 2 and 3	Player 2	Player 3	Players 2 and 3	Player 2	Player 2
Private	-0.0173 (0.0190)	-0.0169 (0.0282)	-0.0177 (0.0255)	0.0307 (0.0194)	-0.0136 (0.0172)	-0.0402** (0.0189)
Period				-0.000556 (0.000610)		-0.00196*** (0.000701)
Period × Private				-0.00310*** (0.00101)		0.00172* (0.00101)
Constant	0.736**** (0.0150)	0.724**** (0.0225)	0.749**** (0.0199)	0.745**** (0.0136)	0.672**** (0.0123)	0.703**** (0.0131)
Observations	14940	7470	7470	14940	10620	10620

Accuracy of higher-order beliefs is measured by $1 - |a_{it} - a_{-i,t}|$. Output of OLS regressions of the form $Y = \beta X + \epsilon$, where $Y = 1 - |a_{it} - a_{-i,t}|$. The first four columns use data from the laboratory; the last two columns use data from the MTurk treatments discussed in Sect. 4. Standard errors are clustered at the level of individual subjects. The period trend is not significantly positive for any treatment or player type, suggesting a failure of higher-order learning

Subject-clustered standard errors in parentheses

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$, **** $p < 0.001$

counterparts in the laboratory experiment ($N=437$, $P = 0.654$).²⁸ We also find no significant effect of private information when we compare second-order beliefs in the within-public and within-private treatments ($P = 0.560$), which suggests that Result 1 does not replicate in the within-subjects data. On the other hand, comparing the average difference between first- and second-order beliefs in the within-public and within-private treatments, we find the difference to be twice as large on average in the private case, although the difference-in-differences is only marginally significant ($P < 0.1$). On average, second-order beliefs are shaded toward the prior more in the private than in the public treatment, although the size of the shading is only 0.034 in within-private and 0.014 in within-public.²⁹

To further explore the effect of private information, we analyze the gap between first- and second-order beliefs at the subject level. To this end, we compute the mean and median difference between first- and second-order beliefs for each subject. If the effect of private information is correctly taken into account, subjects should report equal first- and second-order beliefs more often in within-public than within-private. To test this prediction, we create a dummy variable equal to one if the mean difference between first- and second-order beliefs is negative (first column of Table 3), positive (second column), and zero (third column) and regress these dummy

²⁸ We obtain this p -value from a regression of normalized first-order beliefs against an MTurk dummy variable using the data from the public, private, within-public, and within-private treatments. For the data from the between-subjects treatments, only the first round (the first 30 periods) is used.

²⁹ Both of these values are significantly bigger than zero with largest $p < 0.05$.

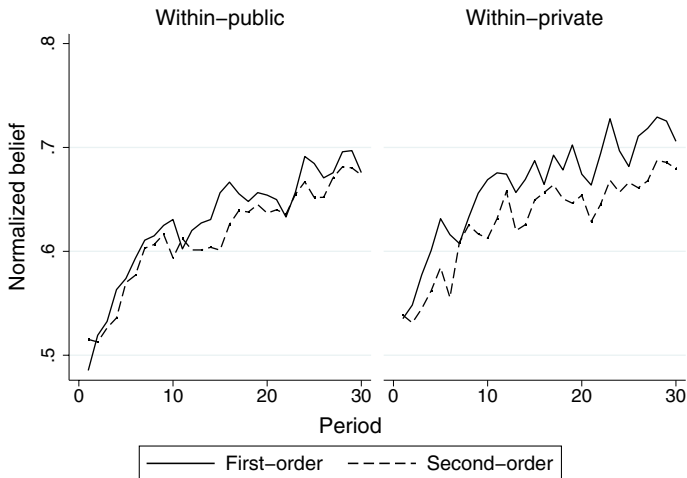


Fig. 5 The evolution of first- and second-order beliefs in the MTurk treatments. The beliefs are normalized by the correct state so that the variable being plotted is B when the state is orange and $1 - B$ when the state is purple, where B is the reported belief

variables against the treatment dummies. We also repeat this exercise for the median in the last three columns of Table 3.

While we find no significant effects on mean differences, private information causes a shift in the medians. Relative to the within-public baseline, we find that the proportion of subjects reporting a positive median difference between first- and second-order beliefs in within-private increases by 13.6% ($P < 0.01$), while that reporting a zero median difference decreases by 14.3% ($P < 0.01$). Thus, a significant fraction of subjects report higher-order beliefs closer to the prior in the presence of private information.

We conclude that the effect of private information is weaker but nevertheless significant in the within-subjects treatments. One possibility is that the within-subjects nature of the design lessened the impact of private information due to bounded rationality. I.e., the subjects might have found it difficult to reason about private information in a setting where they had more tasks (reporting beliefs for both player roles). Another is that the effect would be stronger with learning (i.e., more rounds of matching). We leave these questions open for future research.

Results 2 and 3 replicate in the within-subjects data (last two columns of Table 2). The first two panels of Fig. 6 plot the evolution of belief accuracy, $1 - |a_{it} - a_{-i,t}|$, over time in the within-public and within-private treatments. The figure suggests that higher-order beliefs are not more accurate with public than private information ($P = 0.431$, fifth column of Table 2). The period effect on belief accuracy is negative in the within-public treatment ($P < 0.01$) and insignificant in the within-private treatment ($P = 0.739$, last column of Table 2). Overall, higher-order beliefs do not become more accurate over time.

We can use subject-level differences between first- and second-order beliefs to infer what assumptions subjects make about the reasoning of others. A substantial

Table 3 The effect of private information on the median and mean of the difference between first- and second-order beliefs

	(1)	(2)	(3)	(4)	(5)	(6)
	Mean			Median		
	Neg.	Pos.	Zero	Neg.	Pos.	Zero
Private	-0.0659 (0.0508)	0.0843 (0.0527)	-0.0184 (0.0271)	0.00627 (0.0386)	0.136*** (0.0493)	-0.143*** (0.0532)
Constant	0.373**** (0.0339)	0.549**** (0.0349)	0.0784**** (0.0189)	0.147**** (0.0249)	0.230**** (0.0296)	0.623**** (0.0340)
Observations	354	354	354	354	354	354

Output of OLS regressions, where each observation is a subject in the experiment. The fraction of subjects with a median positive difference between first- and second-order beliefs is greater in the within-private than the within-public treatment

Standard errors in parentheses

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$, **** $p < 0.001$

fraction of subjects—62.3% in the within-public treatment and 48% in the within-private treatment—report a median difference of zero (Table 3). In the private treatment, this suggests that some subjects assume the private information of others to be the same as their own. Projection of private information unto others has recently been experimentally investigated by Danz et al., (2019). To the extent that such behavior deviates from Bayesian use of objective information, it precludes higher-order learning.

A subject in the public treatment might put equal probabilities on her matched partner over- and under-updating relative to her own belief, which would predict equal first- and second-order beliefs despite the fact that the subject believes her partner to be less Bayesian than she herself is.³⁰ Nevertheless, assuming that one's partner has equal beliefs on average might exacerbate belief inaccuracy in the presence of heterogeneity, as argued in Sect. 2.1 (Fig. 2). In Sect. 5, we model subjects' deviations from Bayesian thinking and explore their influence on higher-order belief accuracy in more detail.

5 Base-rate neglect and long-run behavior

As discussed in Sect. 2.1, higher-order learning might fail in this case if deviations from Bayesian updating are not anticipated or shared. We now argue that both of these issues are present in the data. First, there exists substantial heterogeneity in updating types (i.e., deviations from Bayesian updating are not shared). Second, if subjects correctly took this heterogeneity into account, their higher-order beliefs would have been more accurate (i.e., deviations from Bayesian updating are to some extent not anticipated). Nevertheless, higher- and lower-order beliefs would fail to

³⁰ This argument was suggested by a referee.

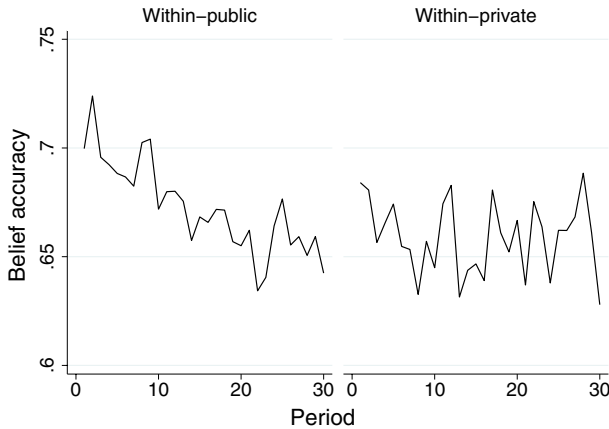


Fig. 6 Failure of higher-order learning in the MTurk treatments. Accuracy of higher-order beliefs is measured by $1 - |a_{it} - a_{-it}|$. Higher-order beliefs in the MTurk treatments are not more accurate with public than private information and fail to become more accurate over time

converge even if subjects were able to forecast the beliefs of others optimally. As we show below, this a consequence of base-rate neglect: subjects’ updating rules are such that neither higher- nor lower-order beliefs converge to the truth even after a large number of signals, making belief inaccuracies persistent.

Consider the case of public signals. We model deviations from Bayesianism following the approach in Grether (1980). Note that Bayes’ rule implies that:

$$\frac{\mu_n}{1 - \mu_n} = \frac{\mu_{n-1}}{1 - \mu_{n-1}} \underbrace{\frac{\text{Prob}(\text{Current ball}|\text{Urn} = \text{Orange})}{\text{Prob}(\text{Current ball}|\text{Urn} = \text{Purple})}}_{LR_n}, \tag{3}$$

where μ_n is the subject’s (first-order) posterior belief, μ_{n-1} is her prior belief,³¹ and LR_n is the likelihood ratio following the observation of the current ball, with $LR_n \in \{LR_{orange} = 2, LR_{purple} = \frac{1}{2}\}$. The following model can be estimated to capture the extent to which subjects deviate from correctly taking into account prior and new information:

$$\ln\left(\frac{\mu_n}{1 - \mu_n}\right) = \beta_0 + \beta_{prior} \ln\left(\frac{\mu_{n-1}}{1 - \mu_{n-1}}\right) + \beta_{LR} \ln(LR_n) + \epsilon. \tag{4}$$

For simplicity, we focus only on the public between-subjects treatment.³² Following Holt and Smith,(2009), we recode 0 guesses as 0.01 and 1 guesses as 0.99 to ensure that equation (4) is well-defined.

³¹ The prior belief in each period is defined as the reported first-order belief from the previous period. In period 1, the prior belief is exogenously given and set at 0.5.

³² The model is not valid for the private treatment. In the within-subjects treatments, two updating rules would need to be estimated for every subject.

Figure 7 shows the histograms of the β_{Prior} and β_{LR} coefficients for subjects in the public treatment. The figure suggests that a substantial degree of heterogeneity is present in the data. Moreover, the distributions of coefficients do not vary significantly across player roles, with the exception of the difference in β_{Prior} between Players 1 and 2 and the difference in β_{LR} between Players 1 and 3, both of which are marginally statistically significant according to a Kolmogorov-Smirnov test ($P < 0.1$).

The estimated distributions of β_{Prior} and β_{LR} allow us to perform the following counterfactual exercise. We form 5000 simulated groups of Player 1, Player 2, and Player 3. For each player in each group, we randomly draw a vector $(\beta_{Prior}, \beta_{LR})$ from the empirical distribution of parameters corresponding to her player type. We then randomly draw 300 signals for each 3-player team. For each player in each group and following each signal, we generate posterior beliefs recursively using the following model, which can easily be obtained from (4):

$$E[\mu_n | \mu_{n-1}, LR_n] = \frac{\mu_{n-1}^{\beta_{Prior}}}{\mu_{n-1}^{\beta_{Prior}} + e^{-\beta_0} (1 - \mu_{n-1})^{\beta_{Prior}} LR_n^{-\beta_{LR}}}. \quad (5)$$

For each player in the role of Player 2 or Player 3, we compute belief accuracy in each period based on the simulated beliefs of the player and the player's matched partner. We then average the distances across all players in a given role for a period-specific prediction of belief accuracy.

The predicted accuracy of higher-order beliefs is reported in Fig. 8. Focusing on the first 30 periods, the simulations provide a reasonable match for the data. While the observed distances between higher- and lower-order beliefs are noisier than the simulated ones, average belief accuracy is 0.72 for Player 2 and 0.75 for Player 3 in the data; the average simulated belief accuracies in the first 30 periods are 0.76 and 0.78 for Players 2 and 3, respectively.

Second, given the updating rules used by the subjects, higher- and lower-order beliefs fail to converge even after 300 periods. Instead, belief accuracy decreases initially and remains flat as more signals are received. Thus, the updating rules used by the players generate a bound for the accuracy of higher-order beliefs. This implies the following prediction:

Prediction 4 *The accuracy of higher-order beliefs does not improve any more after 300 than 30 signals.*

We also simulate optimal beliefs, i.e., the beliefs that would be reported by a sophisticated player that took the empirical distribution of updating coefficients of the target player into account. In order to do this, for every simulated group, every Player k , $k > 1$ in that group, and every realized public history of signals, we first compute the belief corresponding to each possible updating type of Player $k - 1$ (using that player's empirical distribution of updating coefficients), and then average out those posterior beliefs. The average belief corresponds to Player k 's optimal belief given her observed signal history.

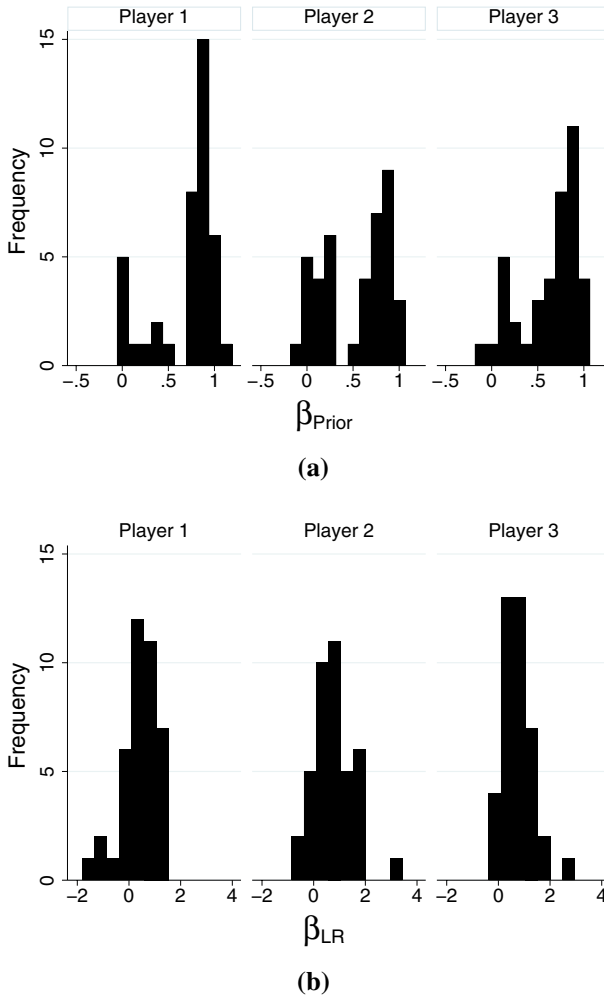
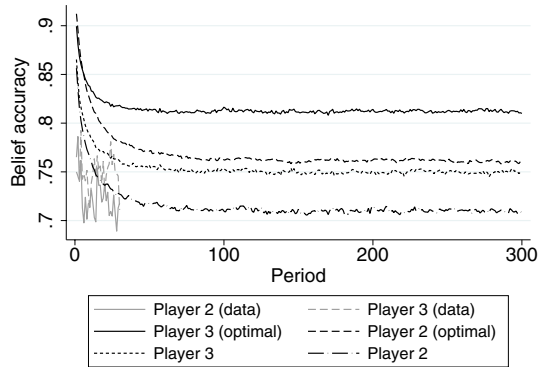


Fig. 7 Histograms of updating parameters estimated at the level of individual subjects in the public treatment ($N = 120$). Following Grether, (1980), the parameters are estimated from the following model: $\ln\left(\frac{\mu_n}{1-\mu_n}\right) = \beta_0 + \beta_{Prior} \ln\left(\frac{\mu_{n-1}}{1-\mu_{n-1}}\right) + \beta_{LR} \ln(LR_n) + \epsilon$

The average accuracy of optimal beliefs is reported in Fig. 8. We find that optimal beliefs are 19% more accurate than observed beliefs for Player 2 and 25% more accurate for Player 3. I.e., taking the distribution of updating types into account confers a benefit. This benefit, however, is limited. Moreover, even if players formed beliefs optimally by taking the distribution of updating types into account, higher- and lower-order beliefs would still diverge. Convergence would not take place even after a large number of signals.

Why do optimal higher-order beliefs fail to correctly predict lower-order beliefs? Our analysis of updating rules shows a pervasive amount of base-rate neglect, that

Fig. 8 Simulated long-run belief accuracy and data in the public treatment. Given subjects' updating rules estimated from the public treatment, higher- and lower-order beliefs fail to converge even after 300 periods in all of the simulations




is, the tendency to underuse one's own previous information.³³ This is reflected in the coefficient β_{prior} being less than 1. While the average subject also manifests an under-inference to new information, that is, $\beta_{LR} < 1$, suppose that β_{LR} were equal to 1 for simplicity. Would beliefs converge to the correct state of the world for an agent who exhibits base-rate neglect? To illustrate, Fig. 9 simulates the beliefs of an agent with mild base-rate neglect ($\beta_{prior} = 0.9$) and no under- or over-inference from new information ($\beta_{LR} = 1$) over 5000 randomly drawn histories and averages out beliefs by period. The simulation shows that even a mild base-rate neglect will lead to long-run beliefs failing to converge and exhibiting non-negligible uncertainty about the correct state of the world.

This observation is not a coincidence. In a recent paper, Benjamin et al., (2019) show theoretically that base-rate neglect has a moderating effect on beliefs, relative to the Bayesian benchmark, and that beliefs fail to converge to the correct state even after observing a large amount of information. Thus, the behavior highlighted in Fig. 9 is a long-run implication of base-rate neglect. In the presence of base-rate neglect and heterogeneity in belief updating, beliefs of players in different roles will converge to different limiting beliefs, if they converge at all. Thus, increasing the amount of information is predicted to generate a failure of higher-order learning even if higher-order beliefs are formed optimally, that is, taking the distribution of updating types into account. This highlights that failure of higher-order learning is generated by the type of heterogeneity present, and not the presence of heterogeneity per se. For example, suppose that agents exhibited no base-rate neglect but under-inferred information contained in new signals. Even with heterogeneity in the β_{LR} parameters across agents, individual beliefs would converge to the correct state, albeit at slower rates than the Bayesian benchmark. Thus, as long as agents believed that others' were heterogeneous only in the β_{LR} parameter, higher-order beliefs would become more accurate over time.

³³ This is a well-known fact in the belief updating literature. For example, see Benjamin (2019, Sect. 4.3, Stylized Fact 9).

Fig. 9 A simulated path of normalized first-order beliefs for an agent with base-rate neglect. The beliefs are normalized by the correct state so that the variable being plotted is B when the state is orange and $1 - B$ when the state is purple, where B is the simulated belief. First-order beliefs fail to converge even after 300 periods



5.1 The long treatment

We ran an additional treatment, within-long, to test Prediction 4. In this treatment, which was otherwise identical to the within-public treatment, each subject received 10 signals about the state of the world in each period. Data from 154 MTurk subjects are reported in the second panel of Fig. 10.³⁴

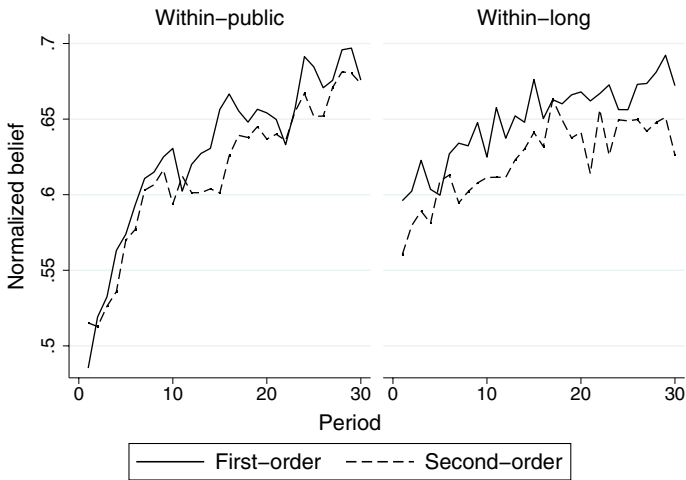
First-order beliefs in the within-long treatment are reported in the top panel of Fig. 10, with those in the within-public treatment also shown for comparison. In the first period, first-order beliefs in the within-long treatment are higher by 11 percentage points than those in the within-public treatment ($P < 0.01$). On the other hand, first-order beliefs are not significantly different across these two treatments in period 30 ($P = 0.912$). Overall, the pattern reported in Fig. 10 suggests that subjects take advantage of additional signals in the early rounds, but that there exists an upper bound on how much the average subject can learn about the state.³⁵

After 30 signals, first-order beliefs are less accurate in the within-long treatment than those in the within-public treatment. It is possible that this is driven by how the within-long treatment was implemented. As mentioned above, we presented subjects with 10 signals about the state of the world in each period of the within-long treatment.³⁶ Providing subjects with 10 signals at a time, as opposed to a sequence of 10 signals, might lead to underinference as discussed by Benjamin (2019, Sec. 4.2, Stylized Fact 2). On the other hand, the simulations, which do not assume bundling, point to an upper bound on first-order belief accuracy. Thus, the initial

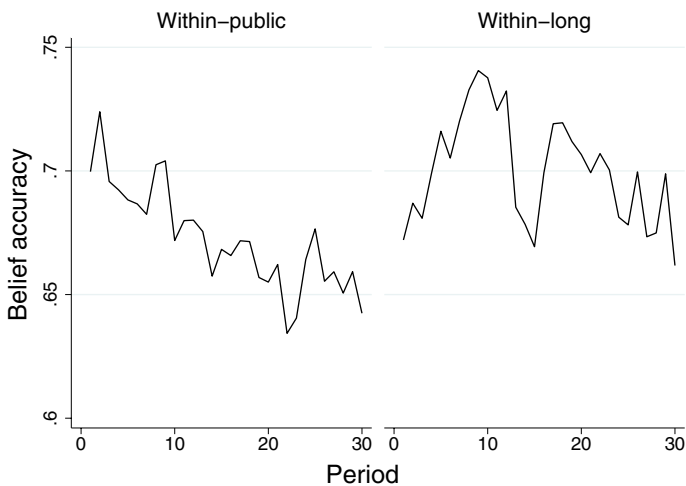
³⁴ We refer to Appendix A for details about the implementation of this treatment.

³⁵ In a recent paper, Esponda et al., (2020) investigate the long-term effect of base-rate neglect by having subjects repeat an updating task involving a binary state and a binary signal 200 times. After each task, subjects are told the realized state. Posterior beliefs are far from the Bayesian benchmark despite the large number of repetitions. While our design and scope are arguably different, Esponda et al's finding also points to a bound on learning.

³⁶ This was done in order to keep the number of elicitations the same across treatments, as well as to avoid decision fatigue.



(a)



(b)

Fig. 10 The evolution of first- and second-order beliefs **a** and failure of higher-order learning **b** in the within-long treatment. Data from the within-public treatment are shown for comparison. In **a**, the beliefs are normalized by the correct state. In **b**, the accuracy of second-order beliefs is measured by $1 - |a_{it} - a_{-it}|$

underinference in the within-long treatment might be driven by bundling, but the overall upper bound on belief accuracy is consistent with base-rate neglect.

The average belief accuracy in the within-long treatment is plotted in the bottom panel of Fig. 10, with that in the within-public treatment again shown for comparison. Overall, we find that beliefs in the within-long treatment are more accurate

than those in the within-public treatment, although the effect is only marginally significant ($P < 0.1$ in a regression of the accuracy measure on a within-long dummy, using only the data from the within-public and within-long treatments). The magnitude of belief inaccuracy in the within-long treatment remains large. Moreover, beliefs in the last 15 periods of the within-long treatment are not more accurate than those in the first 15 periods ($P = 0.361$ in a regression of the accuracy measure on a dummy variable for being in the last 15 rounds using only the data from the within-long treatment). This result is in line with Prediction 4.

The observed accuracy of second-order beliefs in the within-long treatment, together with the out of sample predictions described above, are shown together in Fig. 11. Overall, the data track the simulated predictions well. Taken together, the results in the within-long treatment suggest little benefit in terms of higher-order belief accuracy from receiving 300 as opposed to 30 public signals.

6 Conclusion

This paper presents the first experiment on how higher-order beliefs are updated in response to new information about the fundamentals. We find that subjects engage in higher-order thinking and shade their beliefs toward the prior when they receive private as opposed to public signals. On the other hand, we find that beliefs are not more accurate with public signals, contrary to the Bayesian prediction. Moreover, we find that beliefs do not become more accurate over time with either public or private signals, suggesting a failure of higher-order learning. We attribute this failure to base-rate neglect, heterogeneity in updating rules, and subjects' failure to correctly model how other players deviate from Bayesian reasoning.

Failure of higher-order learning has implications for macroeconomic models. For instance, in a Calvo model with incomplete information about nominal shocks, Angeletos and La'O (2009) show that knowledge about the evolution of first-order beliefs is insufficient to quantify the rate of price adjustment without taking into account the evolution of higher-order beliefs. In turn, higher-order beliefs affect firms' forecasts of other firms' equilibrium actions, which determine their own pricing choices. Our analysis shows that firms might have persistently incorrect beliefs about other firms' beliefs about the size of a nominal shock. More importantly, sluggish price adjustments could persist even when firms in the economy observe only publicly available information.

Our investigation focuses attention on introspective learning where subjects do not receive feedback about the beliefs of others. This design choice guarantees that subjects' higher-order beliefs are not simply the result of adaptation to the behavior of the matched partner. However, it also removes an important source of information which is often available in practice. An interesting extension would be to explore the evolution of higher-order beliefs in the presence of feedback about the average beliefs of a group of subjects to see whether the failure of higher-order learning that we observed could be reduced or even resolved. Noisy feedback, on the other hand, might generate a failure of higher-order learning similar to that we identify.

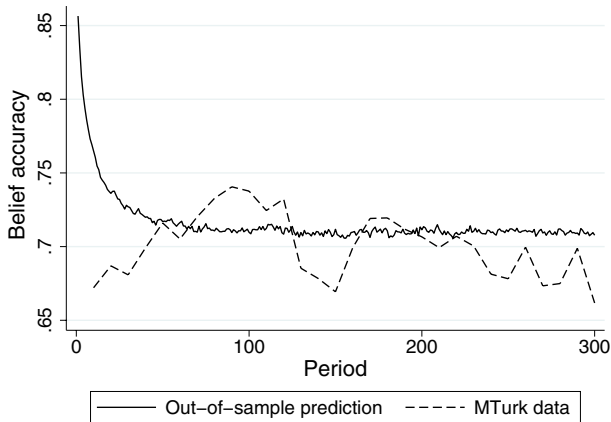


Fig. 11 Out of sample predicted accuracy of second-order beliefs together with the MTurk data in the within-long treatment. For the out of sample predictions, laboratory data from Player 2 in the public treatment is used. Belief accuracy is measured by $1 - |a_{it} - a_{-i,t}|$

Appendix

A Implementation of the online treatments

For the within-public, within-private, and within-long treatments, the subject pool consisted of U.S. workers on Amazon Mechanical Turk (MTurk).³⁷ The data were collected using the software oTree (Chen et al., 2016). Each subject was only allowed to participate in the experiment once, and the experiment was implemented through the assignment of MTurk qualifications to subjects that accepted the HIT (human intelligence task).

The data for the within-public treatment were collected separately as part of a different project. At the beginning of this treatment, subjects were administered a longer version of the Cognitive Reflection test (Frederick, 2005). The data for the within-long and within-private treatments were collected at the request of two referees and did not include any pre-experiment test.³⁸

In all the online treatments, each subject made a total of 60 decisions (30 guesses about the state and 30 guesses about the partner's guess). The experiments' framing was identical to that of the laboratory experiment. Each subject was paid on the basis of one randomly chosen decision out of 60 according to the Binarized Scoring rule with a bonus of \$3.00 or nothing. The payment rule was clearly explained to the subjects in the instructions and several examples were provided, including a table

³⁷ The data for the within-public treatment were collected by Umberto Garfagnini and funded by the University of Surrey. The data for the other two online treatments were collected by Piotr Evdokimov and funded by the Higher School of Economics.

³⁸ These two more recent treatments also included a completion bonus of \$1 to increase the attractiveness of the HIT for potential participants. As this was a flat payment, it did not affect incentives which were kept exactly the same across all the online treatments.

indicating the loss corresponding to one's choice and the different value of the target (the state of the world for the first decision, and the partner's choice otherwise). Each decision screen summarized the most important information from the instructions, including a table of potential losses. The instructions and screenshots can be found in the online Appendix.

Subjects were shown the instructions for the experiment on their screens. After a subject finished reading the instructions, the subject waited to be matched. Given the possibility of no match occurring, subjects were asked to wait for at least up to five minutes for a match (no more than five minutes for within-long and within-private), after which they were allowed to quit (were terminated from) the study and collected their earnings up to that point plus a bonus to compensate the dismissed subjects for their time.³⁹ If two subjects were matched, the experiment automatically began. Each of the two per-period choices was presented on a separate screen and a subject had two minutes (one minute in the within-long and within-private treatments) to submit each choice.⁴⁰

B The dynamics of higher-order beliefs with private signals

Suppose that learning occurs through private signals. Let $\Theta = \{\underline{\theta}, \bar{\theta}\}$ and $p = \text{Prob}(\theta = \bar{\theta})$ denote the common prior belief that the state of the world is $\theta = \bar{\theta}$ in period 0. Consider a binary signal technology with $s \in \{\underline{s}, \bar{s}\}$ and

$$\text{Prob}(s = \bar{s}|\bar{\theta}) = \bar{q}, \quad \text{and} \quad \text{Prob}(s = \bar{s}|\underline{\theta}) = \underline{q}. \quad (6)$$

Since we are assuming that players are Bayesian expected utility maximizers and the signal technology is binary, a sufficient statistic for a history of length n observed by a player is the number of signals of type \bar{s} out of n total signals observed, which we denote by \bar{n} .

Let $p_i^n(\bar{\theta}|\bar{n})$ denote Player i 's posterior belief that the state is $\bar{\theta}$ after having observed \bar{n} signals of type \bar{s} . By Bayes' rule, this probability equals

$$p_i^n(\bar{\theta}|\bar{n}) = \frac{\bar{q}^{\bar{n}}(1 - \bar{q})^{n - \bar{n}}p}{\bar{q}^{\bar{n}}(1 - \bar{q})^{n - \bar{n}}p + \underline{q}^{\bar{n}}(1 - \underline{q})^{n - \bar{n}}(1 - p)}. \quad (7)$$

Players 1, 2 and 3 form the same belief about the state of the world after having observed the same set of signals. Furthermore, if Player $i - 1$ observed a history (n, \bar{n}') , and Player i knew which history Player $i - 1$ observed, then Player i would assign probability one to Player $i - 1$ assigning probability $p_{i-1}^n(\bar{\theta}|\bar{n}')$ to state $\bar{\theta}$.

³⁹ The slight differences between the within-public and the other two treatments is due to software updates to the oTree platform which occurred after we ran the within-public treatment in 2019.

⁴⁰ The time constraint was used to catch dropouts and reassure participants that they would be paid regardless of whether their partner dropped out. While the within-public treatment allowed up to 2 minutes to make each decision, we reduced the time limit to one minute per decision after observing that the vast majority of decisions in the within-public treatment took less than 60 seconds. This is also the case in the lab treatments where subjects faced no time constraints.

However, as learning occurs through private signals, players in higher-order roles need to form beliefs about the histories that players in lower-order roles might have observed. So, suppose that Player i observes history (n, \bar{n}) , she will use this information to update her beliefs about the state of the world which also informs her about the probability with which Player $i - 1$ might have observed any feasible history. For any $\bar{n}' \in \{0, 1, \dots, n\}$, Player i assigns probability $R_i^n(\bar{n}'|\bar{n})$ to Player $i - 1$ having observed \bar{n}' signals of type \bar{s} in n draws when she observed \bar{n} such signals in n draws, this probability is given by

$$R_i^n(\bar{n}'|\bar{n}) = p_i^n(\bar{\theta}|\bar{n})r_n(\bar{n}'|\bar{\theta}) + (1 - p_i^n(\underline{\theta}|\bar{n}))r_n(\bar{n}'|\underline{\theta}), \tag{8}$$

where, from the binomial distribution,

$$r_n(\bar{n}'|\theta) = \begin{cases} \binom{n}{\bar{n}'} \bar{q}^{\bar{n}'} (1 - \bar{q})^{n - \bar{n}'}, & \text{if } \theta = \bar{\theta}, \\ \binom{n}{\bar{n}'} \underline{q}^{\bar{n}'} (1 - \underline{q})^{n - \bar{n}'}, & \text{otherwise,} \end{cases} \tag{9}$$

denotes the probability of observing \bar{n}' signals of type \bar{s} out of n signals, conditional on state θ .

Given the incentives provided by the Binarized scoring rule in an arbitrary period n , Player 2 reports her expectation about the first-order beliefs held by Player 1 which is given by

$$a_2^n(\bar{\theta}|\bar{n}) = E_n[a_1^n(\bar{\theta}|\bar{n}')|\bar{n}] = E_n[p_1^n(\bar{\theta}|\bar{n}')|\bar{n}] = \sum_{\bar{n}'=0}^n p_1^n(\bar{\theta}|\bar{n}')R_i^n(\bar{n}'|\bar{n}), \quad \forall \bar{n} \in \{0, 1, \dots, n\}. \tag{10}$$

We abuse notation slightly and write $a_2^n(\bar{\theta}|\bar{n})$ instead of $a_2^n(\bar{n})$ to make explicit the fact that the beliefs are referenced to the state of the world being equal to $\bar{\theta}$.

Similarly, Player 3 reports her expectation about the report made by Player 2 which is given by

$$a_3^n(\bar{\theta}|\bar{n}) = E_n[a_2^n(\bar{\theta}|\bar{n}')|\bar{n}] = \sum_{\bar{n}'=0}^n a_2^n(\bar{\theta}|\bar{n}')R_i^n(\bar{n}'|\bar{n}), \quad \forall \bar{n} \in \{0, 1, \dots, n\}. \tag{11}$$

Next, we prove the following result:

Proposition B1 *Suppose that the state of the world is $\bar{\theta}$ and $1 > \bar{q} > \frac{1}{2} > \underline{q} > 0$. Then, the following properties hold true:*

- a. $E_n[a_1^n(\bar{\theta}|\bar{n})|\bar{\theta}] > E_n[a_2^n(\bar{\theta}|\bar{n})|\bar{\theta}] > E_n[a_3^n(\bar{\theta}|\bar{n})|\bar{\theta}]$, for any $n > 0$.
- b. For any $i \in \{1, 2, 3\}$, $\lim_{n \rightarrow \infty} E_n[a_i^n(\bar{\theta}|\bar{n})|\bar{\theta}] = 1$.

We split the proof into several steps. We first show two auxiliary results.

Lemma B1 For any $n > 0$, $E_n [a_1^n(\bar{\theta}|\bar{n})|\bar{\theta}] < E_{n+1} [a_1^{n+1}(\bar{\theta}|\bar{n})|\bar{\theta}]$.

Proof Let

$$A(\bar{n}, n) = \frac{\bar{q}^{2\bar{n}}(1 - \bar{q})^{2(n-\bar{n})}p}{\bar{q}^{\bar{n}}(1 - \bar{q})^{n-\bar{n}}p + \underline{q}^{\bar{n}}(1 - \underline{q})^{n-\bar{n}}(1 - p)}, \tag{12}$$

and notice that

$$\begin{aligned} \frac{\partial A(\bar{n}, n)}{\partial n} &= p \ln \left(\frac{\bar{q}}{1 - \bar{q}} \right) \bar{q}^{2\bar{n}}(1 - \bar{q})^{2(n-\bar{n})} \\ &\quad \times \frac{\bar{q}^{\bar{n}}(1 - \bar{q})^{n-\bar{n}}p + \left[1 + \ln \left(\frac{\bar{q}}{1 - \bar{q}} \right) - \ln \left(\frac{q}{1 - q} \right) \right] \underline{q}^{\bar{n}}(1 - q)^{n-\bar{n}}(1 - p)}{\left[\bar{q}^{\bar{n}}(1 - \bar{q})^{n-\bar{n}}p + \underline{q}^{\bar{n}}(1 - \underline{q})^{n-\bar{n}}(1 - p) \right]^2}, \end{aligned} \tag{13}$$

which is strictly positive and well-defined because $1 > \bar{q} > \frac{1}{2} > \underline{q} > 0$. Then,

$$\frac{\partial}{\partial n} \left[\binom{n}{\bar{n}} A(\bar{n}, n) \right] = A(\bar{n}, n) \binom{n}{\bar{n}} \sum_{k=0}^{\bar{n}-1} \frac{1}{n - k} + \binom{n}{\bar{n}} \frac{\partial A(\bar{n}, n)}{\partial n} > 0. \tag{14}$$

This implies that

$$\begin{aligned} E_n [a_1^n(\bar{\theta}|\bar{n})|\bar{\theta}] &= \sum_{\bar{n}=0}^n \binom{n}{\bar{n}} A(\bar{n}, n) \\ &\leq \sum_{\bar{n}=0}^n \binom{n+1}{\bar{n}} A(\bar{n}, n+1) \\ &\leq \sum_{\bar{n}=0}^n \binom{n+1}{\bar{n}} A(\bar{n}, n+1) + \binom{n+1}{n+1} A(n+1, n+1) \\ &= \sum_{\bar{n}=0}^{n+1} \binom{n+1}{\bar{n}} A(\bar{n}, n+1) \\ &= E_{n+1} [a_1^{n+1}(\bar{\theta}|\bar{n})|\bar{\theta}], \end{aligned}$$

which completes the proof. □

Lemma B2 $E_n [a_1^n(\bar{\theta}|\bar{n})|\bar{\theta}] > E_n [a_1^n(\bar{\theta}|\bar{n})|\underline{\theta}]$, for any $n > 0$.

Proof. Note that, for any $n > 0$,

$$\begin{aligned}
 E_n[a_1^n(\bar{\theta}|\bar{n})|\underline{\theta}] &= E_n[p_1^n(\bar{\theta}|\bar{n})|\underline{\theta}] \\
 &= E_n[1 - p_1^n(\bar{\theta}|\bar{n})|\underline{\theta}] \\
 &= 1 - E_n[p_1^n(\bar{\theta}|\bar{n})|\underline{\theta}] \\
 &\geq 1 - E_{n+1}[p_1^{n+1}(\bar{\theta}|\bar{n})|\underline{\theta}] \\
 &= E_{n+1}[a_1^{n+1}(\bar{\theta}|\bar{n})|\underline{\theta}]
 \end{aligned}
 \tag{15}$$

where the inequality follows from Lemma B1. Since

$$E_1[a_1^1(\bar{\theta}|\bar{n})|\bar{\theta}] - E_1[a_1^1(\bar{\theta}|\bar{n})|\underline{\theta}] = \frac{q(1 - \bar{q})p}{(1 - \bar{q})p + (1 - \underline{q})(1 - p)} + \frac{\bar{q}(1 - \underline{q})p}{\bar{q}p + \underline{q}(1 - p)} > 0,
 \tag{16}$$

the result follows from the opposite monotonicity of the two sequences $\{E_n[a_1^n(\bar{\theta}|\bar{n})|\bar{\theta}]\}_n$ and $\{E_n[a_1^n(\bar{\theta}|\bar{n})|\underline{\theta}]\}_n$. \square

We now prove part a) of the proposition.

Lemma B3 For any $n \geq 1$ and $i \geq 2$, $E_n[a_{i-1}^n(\bar{\theta}|\bar{n})|\bar{\theta}] > E_n[a_i^n(\bar{\theta}|\bar{n})|\bar{\theta}]$.

Proof Notice that

$$\begin{aligned}
 E_n[a_i^n(\bar{\theta}|\bar{n})|\bar{\theta}] &= \sum_{\bar{n}=0}^n a_i^n(\bar{\theta}|\bar{n})r_n(\bar{n}|\bar{\theta}) \\
 &= \sum_{\bar{n}=0}^n \left[\sum_{\bar{n}'=0}^n a_{i-1}^n(\bar{\theta}|\bar{n}')R_i^n(\bar{n}'|\bar{n}) \right] r_n(\bar{n}|\bar{\theta}) \\
 &= \sum_{\bar{n}'=0}^n a_{i-1}^n(\bar{\theta}|\bar{n}') \left[\sum_{\bar{n}=0}^n R_i^n(\bar{n}'|\bar{n})r_n(\bar{n}|\bar{\theta}) \right] \\
 &= \sum_{\bar{n}'=0}^n a_{i-1}^n(\bar{\theta}|\bar{n}') [(r_n(\bar{n}'|\bar{\theta}) - r_n(\bar{n}'|\underline{\theta}))E_n[p_i^n(\bar{\theta}|\bar{n})|\bar{\theta}] + r_n(\bar{n}'|\underline{\theta})] \\
 &= E_n[a_{i-1}^n(\bar{\theta}|\bar{n})|\underline{\theta}] \\
 &\quad + E_n[p_i^n(\bar{\theta}|\bar{n})|\bar{\theta}] \{E_n[a_{i-1}^n(\bar{\theta}|\bar{n})|\bar{\theta}] - E_n[a_{i-1}^n(\bar{\theta}|\bar{n})|\underline{\theta}]\}.
 \end{aligned}
 \tag{17}$$

Then,

$$\begin{aligned}
 E_n[a_{i-1}^n(\bar{\theta}|\bar{n})|\bar{\theta}] - E_n[a_i^n(\bar{\theta}|\bar{n})|\bar{\theta}] &= \{1 - E_n[p_i^n(\bar{\theta}|\bar{n})|\bar{\theta}]\} \\
 &\quad \times \{E_n[a_{i-1}^n(\bar{\theta}|\bar{n})|\bar{\theta}] - E_n[a_{i-1}^n(\bar{\theta}|\bar{n})|\underline{\theta}]\},
 \end{aligned}
 \tag{18}$$

which is positive if and only if $E_n[a_{i-1}^n(\bar{\theta}|\bar{n})|\bar{\theta}] > E_n[a_{i-1}^n(\bar{\theta}|\bar{n})|\underline{\theta}]$.⁴¹ We know from Lemma B2 that this inequality holds for $i = 1$. So, suppose that it holds for $i = k - 1$, we will show that it also holds for $i = k$. By equivalently expressing equation (17) in terms of $\underline{\theta}$ instead of $\bar{\theta}$ gives

$$E_n[a_k^n(\underline{\theta}|\bar{n})|\underline{\theta}] = E_n[a_{k-1}^n(\underline{\theta}|\bar{n})|\bar{\theta}] + E_n[p_k^n(\underline{\theta}|\bar{n})|\underline{\theta}]E_n[a_{k-1}^n(\underline{\theta}|\bar{n})|\underline{\theta}] - E_n[p_k^n(\underline{\theta}|\bar{n})|\underline{\theta}]E_n[a_{k-1}^n(\underline{\theta}|\bar{n})|\bar{\theta}]. \tag{19}$$

Next, given that the state space is binary, it holds that

$$\begin{aligned} E_n[a_k^n(\bar{\theta}|\bar{n})|\underline{\theta}] &= E_n[E_n[a_{k-1}^n(\bar{\theta}|\bar{n}')|\bar{n}]|\underline{\theta}] \\ &= E_n[1 - E_n[a_{k-1}^n(\underline{\theta}|\bar{n})'|\bar{n}]|\underline{\theta}] \\ &= 1 - E_n[E_n[a_{k-1}^n(\underline{\theta}|\bar{n})'|\bar{n}]|\underline{\theta}] \\ &= 1 - E_n[a_k^n(\underline{\theta}|\bar{n})|\underline{\theta}]. \end{aligned} \tag{20}$$

Combining (17) and (20), and rearranging using (20), we obtain

$$\begin{aligned} E_n[a_k^n(\bar{\theta}|\bar{n})|\bar{\theta}] - E_n[a_k^n(\bar{\theta}|\bar{n})|\underline{\theta}] &= \{1 + E_n[p_k^n(\bar{\theta}|\bar{n})|\bar{\theta}] - E_n[p_k^n(\underline{\theta}|\bar{n})|\underline{\theta}]\} \\ &\quad \times \{E_n[a_{k-1}^n(\bar{\theta}|\bar{n})|\bar{\theta}] - E_n[a_{k-1}^n(\bar{\theta}|\bar{n})|\underline{\theta}]\} \\ &= \{1 + E_n[p_1^n(\bar{\theta}|\bar{n})|\bar{\theta}] - E_n[p_1^n(\bar{\theta}|\bar{n})|\underline{\theta}]\} \\ &\quad \times \{E_n[a_{k-1}^n(\bar{\theta}|\bar{n})|\bar{\theta}] - E_n[a_{k-1}^n(\bar{\theta}|\bar{n})|\underline{\theta}]\} \\ &= \{1 + E_n[a_1^n(\bar{\theta}|\bar{n})|\bar{\theta}] - E_n[a_1^n(\bar{\theta}|\bar{n})|\underline{\theta}]\} \\ &\quad \times \{E_n[a_{k-1}^n(\bar{\theta}|\bar{n})|\bar{\theta}] - E_n[a_{k-1}^n(\bar{\theta}|\bar{n})|\underline{\theta}]\} \end{aligned} \tag{21}$$

where the second equality follows from the observation that all players share the same first-order beliefs given the same history. Since $E_n[a_1^n(\bar{\theta}|\bar{n})|\bar{\theta}] - E_n[a_1^n(\bar{\theta}|\bar{n})|\underline{\theta}] > 0$ by Lemma B2, and $E_n[a_{k-1}^n(\bar{\theta}|\bar{n})|\bar{\theta}] - E_n[a_{k-1}^n(\bar{\theta}|\bar{n})|\underline{\theta}] > 0$ by the induction hypothesis, it follows that $E_n[a_k^n(\bar{\theta}|\bar{n})|\bar{\theta}] - E_n[a_k^n(\bar{\theta}|\bar{n})|\underline{\theta}] > 0$. Finally, this implies that $E_n[a_k^n(\bar{\theta}|\bar{n})|\bar{\theta}] > E_n[a_{k+1}^n(\bar{\theta}|\bar{n})|\bar{\theta}]$, for any $k \geq 1$. \square

Next, we show part *b*) of the proposition. By Lemma B1, the sequence $\{E_n[a_1^n(\bar{\theta}|\bar{n})|\bar{\theta}]\}_{n=1}^\infty$ is monotonically increasing. As the sequence is also uniformly bounded above by 1, it converges. That $\lim_{n \rightarrow \infty} E_n[a_1^n(\bar{\theta}|\bar{n})|\bar{\theta}] = 1$ follows from Markov's inequality and a standard argument on the martingale property of Bayesian updating and is thus omitted. Similarly, $\lim_{n \rightarrow \infty} E_n[a_1^n(\bar{\theta}|\bar{n})|\underline{\theta}] = 0$. Therefore, from (17) with $i = 2$, we can conclude that $\lim_{n \rightarrow \infty} E_n[a_2^n(\bar{\theta}|\bar{n})|\bar{\theta}] = 1$. Using this fact and (17) again, we can show that the claim also holds for $i = 3$. This completes the proof.

⁴¹ $E_n[p_i^n(\bar{\theta}|\bar{n})|\bar{\theta}] \in (0, 1)$.

Supplementary Information The online version contains supplementary material available at <https://doi.org/10.1007/s10683-021-09743-6>.

Acknowledgement We are grateful to Ryan Oprea and seminar participants at the ESA World Meeting (Berlin) and the University of Southampton, the Editor, and two anonymous referees for very helpful comments. Financial support from the Asociación Mexicana de Cultura A.C., the Higher School of Economics, and the University of Surrey is acknowledged.

Open Access This article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence, and indicate if changes were made. The images or other third party material in this article are included in the article's Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the article's Creative Commons licence and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this licence, visit <http://creativecommons.org/licenses/by/4.0/>.

References

- Acemoglu, D., Chernozhukov, V., & Wildiz, M. (2016). Fragility of asymptotic agreement under Bayesian learning. *Theoretical Economics*, *11*, 187–225.
- Angeletos, G.-M., Hellwig, C., & Pavan, A. (2007). Dynamic global games of regime change: Learning, multiplicity, and the timing of attacks. *Econometrica*, *75*, 711–756.
- Angeletos, G.-M., & La'O, J. (2009). Incomplete information, higher-order beliefs and price inertia. *Journal of Monetary Economics*, *56*, S19–S37.
- Azrieli, Y., Chambers, C. P., & Healy, P. J. (2018). Incentives in experiments: A theoretical analysis. *Journal of Political Economy*, *126*, 1472–1503.
- Benjamin, D., Bodoh-Creed, A., & Rabin, M. (2019). Base-rate neglect: Foundations and implications, Technical report
- Benjamin, D. J. (2019). Errors in probabilistic reasoning and judgment biases. *Handbook of Behavioral Economics: Applications and Foundations*, *1*(2), 69–186.
- Carlsson, H., & Van Damme, E. (1993). Global games and equilibrium selection. *Econometrica: Journal of the Econometric Society*, *61*, 989–1018.
- Chen, D. L., Schonger, M., & Wickens, C. (2016). oTree—An open-source platform for laboratory, online, and field experiments. *Journal of Behavioral and Experimental Finance*, *9*, 88–97.
- Cornand, C., & Heinemann, F. (2014). Measuring agent's reaction to private and public information in games with strategic complementarities. *Experimental Economics*, *17*, 61–77.
- Costa-Gomes, M., Crawford, V. P., & Broseta, B. (2001). Cognition and behavior in normal-form games: An experimental study. *Econometrica*, *69*, 1193–1235.
- Costa-Gomes, M. A., & Crawford, V. P. (2006). Cognition and behavior in two-person guessing games: An experimental study. *American Economic Review*, *96*, 1737–1768.
- Costa-Gomes, M. A., & Weizsäcker, G. (2008). Stated beliefs and play in normal-form games. *The Review of Economic Studies*, *75*, 729–762.
- Coutts, A. (2018). Good news and bad news are still news: Experimental evidence on belief updating. *Experimental Economics*, *22*, 1–27.
- Cripps, M. W., Ely, J. C., Mailath, G. J., & Samuelson, L. (2008). Common learning. *Econometrica*, *76*, 909–933.
- Danz, D., Madarász, K., & Wang, S. (2019). The biases of others: Anticipating informational projection in an agency setting, Technical report
- Danz, D., Vesterlund, L., & Wilson, A. J. (2020). Belief elicitation: Limiting truth telling with information on incentives. National Bureau of Economic Research: Technical report
- Eil, D., & Rao, J. M. (2011). The good news-bad news effect: Asymmetric processing of objective information about yourself. *American Economic Journal: Microeconomics*, *3*, 114–38.

- Epstein, L. G., Noor, J., & Sandroni, A. (2008). Non-Bayesian updating: A theoretical framework. *Theoretical Economics*, *3*, 193–229.
- Esponda, I., Vespa, E., & Yuksel, S. (2020). Mental models and learning: The case of base-rate neglect. Tech. rep.
- Eyster, E., & Rabin, M. (2005). Cursed equilibrium. *Econometrica*, *73*, 1623–1672.
- Fischbacher, U. (2007). z-Tree: Zurich toolbox for ready-made economic experiments. *Experimental Economics*, *10*, 171–178.
- Frederick, S. (2005). Cognitive reflection and decision making. *Journal of Economic perspectives*, *19*, 25–42.
- Grether, D. M. (1980). Bayes rule as a descriptive model: The representativeness heuristic. *The Quarterly Journal of Economics*, *95*, 537–557.
- Grether, D. M. (1992). Testing Bayes rule and the representativeness heuristic: Some experimental evidence. *Journal of Economic Behavior and Organization*, *17*, 31–57.
- Hara, K., Adams, A., Milland, K., Savage, S., Callison-Burch, C., & Bigham, J. P. (2018). A data-driven analysis of workers' earnings on amazon mechanical turk, in *Proceedings of the 2018 CHI Conference on Human Factors in Computing Systems*, ACM, p. 449.
- Harsanyi, J. C. (1967). Games with incomplete information played by "Bayesian" players, I–III Part I. *The basic model*, *Management Science*, *14*, 159–182.
- Heinemann, F., Nagel, R., & Ockenfels, P. (2004). The theory of global games on test: Experimental analysis of coordination games with public and private information. *Econometrica*, *72*, 1583–1599.
- Holt, C. A., & Smith, A. M. (2009). An update on Bayesian updating. *Journal of Economic Behavior and Organization*, *69*, 125–134.
- Hossain, T., & Okui, R. (2013). The binarized scoring rule. *Review of Economic Studies*, *80*, 984–1001.
- Huck, S., & Weizsäcker, G. (2002). Do players correctly estimate what others do? Evidence of conservatism in beliefs. *Journal of Economic Behavior and Organization*, *47*, 71–85.
- Keynes, J. M. (1937). The general theory of employment. *The Quarterly Journal of Economics*, *51*, 209–223.
- Kneeland, T. (2015). Identifying higher-order rationality. *Econometrica*, *83*, 2065–2079.
- Kübler, D., & Weizsäcker, G. (2004). Limited depth of reasoning and failure of cascade formation in the laboratory. *The Review of Economic Studies*, *71*, 425–441.
- Manski, C. F., & Neri, C. (2013). First-and second-order subjective expectations in strategic decision-making: Experimental evidence. *Games and Economic Behavior*, *81*, 232–254.
- Mertens, J.-F., & Zamir, S. (1985). Formulation of Bayesian analysis for games with incomplete information. *International Journal of Game Theory*, *14*, 1–29.
- Mobius, M. M., Niederle, M., Niehaus, P., & Rosenblat, T. S. (2014). Managing self-confidence: Theory and experimental evidence. National Bureau of Economic Research: Tech. rep.
- Morris, S., & Shin, H. S. (1998). Unique equilibrium in a model of self-fulfilling currency attacks. *American Economic Review*, *88*, 587–597.
- Morris, S., & Shin, H. S. (2002). Social value of public information. *American Economic Review*, *92*, 1521–1534.
- Nagel, R. (1995). Unraveling in guessing games: An experimental study. *The American Economic Review*, *85*, 1313–1326.
- Stahl, D. O., & Wilson, P. W. (1994). Experimental evidence on player's models of other players. *Journal of Economic Behavior and Organization*, *25*, 309–327.
- Stahl, D. O., & Wilson, P. W. (1995). On player's models of other players: Theory and experimental evidence. *Games and Economic Behavior*, *10*, 218–254.
- Van Huyck, J., Viriyavipart, A., & Brown, A. L. (2018). When less information is good enough: Experiments with global stag hunt games. *Experimental Economics*, *21*, 527–548.
- Weizsäcker, G. (2003). Ignoring the rationality of others: Evidence from experimental normal-form games. *Games and Economic Behavior*, *44*, 145–171.
- Wiseman, T. (2012). A partial folk theorem for games with private learning. *Theoretical Economics*, *7*, 217–239.