

Beliefs about Others: A Striking Example of Information Neglect

Sebastian Fehrler

University of Bremen, SOCIUM
Research Center on Inequality and Social Policy
sebastian.fehrler@uni-bremen.de

Baiba Renerte

University of Konstanz
& Thurgau Institute of Economics
baiba.renerte@uni-konstanz.de

Irenaeus Wolff

University of Konstanz
& Thurgau Institute of Economics
irenaeus.wolff@uni-konstanz.de

February 16, 2023

Abstract

In many games of imperfect information, players can make Bayesian inferences about other players' types based on the information that is contained in their own type. Several behavioral theories of belief-updating even start from the assumption that players project their own type onto others also when it is not rational. We investigate such inferences in a simple environment that is a vital ingredient of numerous game-theoretic models and experiments, in which types are drawn from one out of two states of the world and participants have to guess the type of another participant. We find little evidence for irrational (over-)projection. Instead, between 50% and 70% of the participants in our experiment completely neglect the information contained in their own type and base their beliefs and choices only on the prior probabilities. Using several experimental interventions, we show that this striking neglect of information is very robust.

JEL classification: C91, D83.

Keywords: Projection, Bayesian inference, biased belief-updating.

1 Introduction

We examine how people make inferences about the types of other people based on their own type. There are many situations in which such inferences play a role. Just to name two examples that have been extensively discussed in the literature, inferences about other people’s types play a critical role in auctions with correlated valuations (Cremer and McLean, 1985, 1988; Brusco, 1998; Breitmoser, 2019) and elections with correlated political preferences (Goeree and Grober, 2007; Taylor and Yildirim, 2010; Agranov et al., 2017; Tolvanen, 2017).

For settings like the above, the theoretical prediction is that people make rational Bayesian inferences about the type distribution based on the information that is contained in their own type. Such inferences are standard in economic theories with fully rational agents. However, few behavioral economists would expect many people to update in a perfect Bayesian way. In fact, we recruited experimental economists via the ESA-discussion mailing list in May 2021 to predict our participants’ updating behavior in the focal situation of this paper (described below). On average, the 59 respondents estimated the fraction of participants who would update their beliefs in a perfect Bayesian way at 12%. In addition, the experimentalists estimated that 14% would adjust their belief towards the Bayesian benchmark but by too little.

Several behavioral theories, such as social projection, information projection, or type projection (Dawes, 1989; Madarász, 2012, 2015; Breitmoser, 2019) assume that people over-infer from their own type and thus over-project it onto others. This assumption is supported by a number of empirical studies (e.g., Messé and Sivacek, 1979; Offerman et al., 1996; Aksoy and Weesie, 2012; Blanco et al., 2014; Danz et al., 2018). This is also what our sample of experimental economists predicted to be the most prevalent behavior: on average, 34% were predicted to adjust their beliefs too strongly. And to be honest, over-inference was also what we expected to be the modal behavior before we ran the experiments.

However, the above-mentioned over-inference stands in stark contrast with findings from other studies in which inferences are not about another person’s type. These studies show that people often neglect available sources of information (e.g., Friedman, 1998; Page, 1998; Hanna et al., 2014; Esponda and Vespa, 2014; Enke and Zimmermann, 2019). Nonetheless, the mean estimate of the 59 experimental economists for the fraction of participants who do not adjust the prior probabilities at all was only 26%.¹ In light of the contrasting findings from the literature, it would be important to study a situation that is a wide-spread ingredient of both theoretic and experimental work and in which both (over-)projection and information neglect are possible. This is what we do in this paper.

In a laboratory experiment with three consecutively developed treatments and an additional control experiment, we make the players’ types very salient and thus create favorable conditions for finding rational projection and over-projection. For the purposes of this paper, we use the term “rational projection” to refer to situations in which the rational response directionally coincides with the projection

¹Note that the estimates of those whose work was closest to ours—judging by their self-reports in the survey—did not perform much better than others in the sample.

response, but does not overshoot to over-projection. The basic set-up that we use to study projection is a common ingredient in many experimental studies of games of imperfect-information (e.g., Guarnaschelli et al., 2000; Holt and Smith, 2009, 2016; Bouton et al., 2016; Fehrler and Hughes, 2018).

There are two possible states of the world, represented by a dark urn and a light urn. Each urn contains two balls of its own color and one ball of the other urn's color. One of the states/urns is randomly drawn. Then, player A's type (a dark ball or a light ball) is drawn from the true state of the world. After replacement of A's ball, player B's type (ball) is drawn from the same urn. Player A does not get to see B's type. Instead, player A's task is to guess B's type. Between rounds, we vary the prior probabilities of the states of the world.²

In this setup, the Bayesian response often coincides with the response of an over-projecting agent. Nevertheless, we are able to disentangle rational projection from over-projection as we elicited participants' beliefs about how likely the other player is of a certain type. However, given the lack of evidence for any type of projection in our data, there is no need to disentangle Bayesian updating from over-projection: while many subjects do react to changes in the prior, they do not react at all to the information that is contained in their own type. For example, when asked about their belief about the other player's type in the easiest setting of a 50-50 prior for either urn, a majority of around 70%(!) of subjects state the prior probability.³ Note the difference to the prediction of even the experimental economists (who should be more inclined to expect deviations from fully rational behavior than any other group of economists): they predicted, on average, that only 26% of the population would report a 50-50 probability as their posterior.⁴

We test the robustness of this result with a battery of interventions that nudge participants toward making use of this source of information. The interventions that we employ include (i) belief-elicitation tasks (with respect to both the other's type and the state of the world) between the guessing rounds to highlight the importance of belief-updating, (ii) tasks without replacement to highlight the importance of replacement, (iii) tasks with several additional draws from the selected urn to stimulate learning, (iv) the use of physical urns and balls to avoid potential confusion about computerized processes, and (v) questions about the subject's reasoning behind previous choices to increase deliberation and facilitate the identification of reasoning errors. While participants do react to the interventions themselves, none of the interventions is very effective in reducing information neglect in the subsequent type-guessing task; only asking for beliefs or making a player guess the color of her own next ball (in an additional control experiment) slightly reduces it. Hence, we conclude that the striking degree of information neglect that we observe is a robust finding.

²Our set-up is "value-neutral" and does not involve "ego-relevant" belief-updating, such as updating from feedback on own performance, to avoid biases through motivated reasoning (see Coutts, 2018).

³This is also a strong piece of evidence against stochastic choice as a possible alternative explanation for the low projection rates that we find in our main task. If participants were randomizing their choices and answers, they would not be so consistent in their belief statements.

⁴Note that we did incentivize the respondents: whoever would be closest in their predictions on this and three other experiments would earn a prize of 200 Euros.

The self-described reasoning of participants suggests that many believe that nothing has changed when the ball is put back into the urn as compared to the situation before receiving the ball. Hence, the probability that the other person sees a ball of the same color must equal the prior probability. This resembles the “no-change principle” that many people appear to be following in the more complicated Monty Hall problem (Burns, 2017)—a classic setting in which information neglect is the rule rather than the exception.⁵ Next, we discuss the related literature and state our hypotheses.

2 Related literature and hypotheses

2.1 Projection

“Projection” can be defined as follows. People who project infer from their own type that another person shares their type with a higher probability than the prior probability. This can be the rational result of Bayesian updating—when one’s type is indeed informative about the other player’s type—or it may be irrational. Our definition includes all sorts of projections, from social projections (Dawes, 1989) to information projection (Madarász, 2012), action projection (Al-Nowaihi and Dhami, 2015) and type projection (Breitmoser, 2019).

Social projection is an established concept in psychology, incorporating the assumption that people tend to project their own preferences, behavior and intentions onto others (Ross et al., 1977). A mathematical foundation for social projection was first suggested by Dawes in his rationalization of the false-consensus effect (Dawes, 1989). Today, the terms “false-consensus bias” and “social-projection bias” are often used synonymously. Closely related to social projection, Madarász (2012; 2015) studies projection of information from an informed participant to an uninformed one.

Another closely related concept is type projection. It incorporates a social-projection bias similar to the one defined above and also relates to information projection. Game-theoretically—though perhaps not psychologically—all bits of private information, preferences and informative signals can be subsumed under the term “type” and treated equivalently with the set of tools that are available for the analysis of games of imperfect information. Breitmoser (2019) sets up an elegant model of type projection in Bayesian games, where some players overestimate the probability that other people in their group share their type. He argues that people over-infer from their own type and project it onto others. He defines “type” in the context of auctions with two-sided imperfect information as a person’s signal about the object value. Type-projecting agent believes that with some probability p , their opponents have the same type as themselves, and with probability $(1 - p)$, their opponents have the types that Bayesian updating would indicate. The projecting agents then best-respond to the mixture of Bayesian and projected types.

⁵We describe and discuss the Monty Hall problem in the next section.

There is evidence for projection both in non-strategic and strategic environments and both between individuals (Messé and Sivacek, 1979; Offerman et al., 1996; Aksoy and Weesie, 2012; Blanco et al., 2014; Danz et al., 2018) and within individuals but between different points in time (Gilbert et al., 1998; Read and Van Leeuwen, 1998; Conlin et al., 2007; Danz et al., 2015), which we interpret as suggesting that it might be a general phenomenon. In line with the literature on projection, we thus state the hypothesis that participants (over-)infer from, and thus (over-)project, their own types in our simple belief-updating tasks.

Projection Hypothesis: *Participants (over-)infer from their own types and thus (over-)project their own types onto others.*

However, a second strand of literature has revealed evidence that subjects often neglect available sources of information.

2.2 Information Neglect

Several theories predict that people will ignore one or another element of the information structure, often in situations in which non-trivial contingent reasoning is necessary.⁶ We subsume these theories under the label “information-neglect theories”.

For example, there is the work by Eyster and Rabin (2005, 2010), who argue that there are two ways in which people in strategic settings may err when inferring from information—first, an extreme form of inferential neglect (cursed equilibrium), and second, a form of inferential naivety. Type projection is conceptually related to cursed equilibrium (Eyster and Rabin, 2005). Both models are solution concepts for Bayesian games, in which agents distort the correlation between the opponents’ actual types and their strategies. The difference between cursed equilibrium and type-projection equilibrium is that type-projecting agents project their own type onto their opponents, while cursed agents act as if they face a random type that is drawn from the prior type distribution. This results from a neglect of a possible correlation between the opponents’ (potential) types and their actions. Put differently, cursed players neglect the implicit information about the type they are facing, which is revealed through the history of play, and stems from the fact that different types have different incentives to choose certain actions.⁷ In another contribution, the same authors contrast cursed “rational herding” and “naïve herding” (or inferential naivety; Eyster and Rabin, 2010). Under inferential naivety, players fail to fully attend to the strategic logic of the setting they are in. In consequence, they naïvely believe that each previous player’s action reflects solely that player’s private information.

⁶For recent general discussions of the difficulties with contingent reasoning see Miller and Sanjurjo (2019) and Martínez-Marquina et al. (2019). Note also the various terms used in the literature to describe related concepts to information neglect, e.g. conservatism (Edwards, 1968; Fischhoff and Beyth-Marom, 1983; El-Gamal and Grether, 1995) and base-rate fallacy (Meehl and Rosen, 1955; Kahneman and Tversky, 1972; Bar-Hillel, 1980).

⁷A classic example is the inference that a buyer can (neglect to) make about the value of a product (the seller’s type) from the seller’s willingness to sell it at a certain price in Akerlof’s (1970) market for lemons.

There are further information-neglect concepts. First, correlation neglect in the updating process can make beliefs excessively sensitive to well-connected information sources (Levy and Razin, 2015; Enke and Zimmermann, 2019). Second, pivotality neglect and difficulties to infer from hypothetical events cause biases that can explain naïve voting and bidding behavior (Esponda and Vespa, 2014). Third, base-rate neglect (Bar-Hillel, 1980) and other fallacies are known to cause ignorance of some relevant (for example, statistical) information in favor of using other irrelevant information. Fourth, Hanna et al. (2014) develop a theory of learning through noticing and present supporting evidence from a field experiment, in which farmers often fail to notice readily available information to optimize their production technology.

Benjamin (2019) conducts a review and meta-analysis of the literature on belief-updating. In settings that are similar to ours but that do not elicit inferences about other people's types, people tend to underweight both new information and prior information, to roughly similar degrees.⁸ However, the study provides no indication that people may neglect information altogether.

Finally, the most studied situation in which information is often neglected is the Monty Hall problem (Friedman, 1998; Page, 1998). In this puzzle, a game-show participant faces three doors. Behind two of the doors, there is a worthless prize (a "goat"), while behind the third, there is a highly valuable prize (a "car"). The most popular protocol now proceeds as follows: the participant chooses one of the three doors, which remains closed. The game-show host (famously, Monty Hall in "Let's make a deal") opens one of the remaining two doors but never opens the door hiding the car (uniformly randomizing which door to open when the participant's initially chosen door hides the car). Finally, the participant is asked whether she wants to stick to her original choice of doors or to switch to the other unopened door that she had not chosen at the outset. Under this set of rules (in particular, that the host always opens a door, that the opened door never hides a car, and that the host randomizes between goat-hiding doors), it should be clear that switching yields a probability of two thirds of getting the car.⁹ However, most people who are confronted with the puzzle prefer to stick to their initial choice, even after many repetitions (e.g., Friedman, 1998).

In the Monty Hall problem, most people seem to be "cursed". They neglect the information contained in the host's choice of which alternative door to open: if this choice was random (and not restricted to doors hiding goats), then switching and non-switching would lead to the same probabilities. Note though that cursedness alone would not be enough to explain choices in the Monty Hall problem: if people saw the two remaining options as equivalent, they should not favor sticking to their initial choice.¹⁰ After Friedman (1998) introduced the problem into economics, a number of researchers have

⁸Benjamin follows Grether's (1980) framework, in which the posterior odds of the urns given the observed signal are equal to the ratio of the signal's probabilities under the two types of urns, elevated to the power of " c ", times the prior odds, elevated to the power of " d ". Bayesian updating would imply $c = d = 1$, but Benjamin estimates the values to be $c = 0.38$ and $d = 0.43$ in incentivized studies of "simultaneous-sample updating problems" (Tables 2 and 4).

⁹See, e.g., Selvin (1975b,a), Nalebuff (1987), or Gillman (1992).

¹⁰See, for example, Friedman (1998) for a discussion of additional forces that may combine with cursedness to yield the observed non-switching behavior.

studied variants of the Monty Hall problem, looking at whether individual biases in the task would be reflected in market prices of assets on the possible doors (Kluger and Wyatt, 2004; Siddiqi, 2009), at whether communication and competition eliminate the bias or at how intelligence and communication relate to learning the optimal strategy (Palacios-Huerta, 2003).¹¹ Tor and Bazerman (2003) attribute non-switching (as well as two other empirical phenomena) to missing attention and a resulting failure “to consider all of the information needed”.

This review of information-neglect studies and theories is certainly not exhaustive but sufficient to raise doubts about whether people are always able to make use of the information contained in their own type.¹² Hence, people may not project their type onto others, even when it would be rational. We thus formulate the following alternative hypothesis to projection:

Information-Neglect Hypothesis: *Participants neglect the information that their types contain.*

Several studies discussed in this subsection use nudges to encourage subjects to use available information. We discuss these nudges and how they relate to our own in the following.

2.3 Nudging people to use the available information

Enke and Zimmermann (2019) show that repeatedly pointing the subjects to the correlation structure, by juxtaposing a set-up with correlation and a set-up without correlation, effectively reduces correlation neglect, causing subjects to take this element of the information structure into account. In our set-up, we contrast the basic set-up with replacement with a set-up without replacement to point the subjects to the fact that the second ball is always drawn from the same urn.

Esponda and Vespa (2014) try several interventions to make subjects condition their behavior on the hypothetical information contained in the event of being pivotal. These include feedback, sequential rather than simultaneous play and several verbal hints. However, when they switch back to the main task in the experiment, error rates jump back to the initial, high level. Taking a similar approach, we also implement several interventions to improve participants’ reasoning.

Hanna et al. (2014) summarize the relevant information for the farmers who are then better able to notice them and take them into account. In our experiment, we go to great lengths to make the set-up as clear and easy to understand as possible. We try two different decision screens and even use physical urns and balls to demonstrate the information structure at the beginning of the session. We thus can rule out people not noticing or understanding what happens when as an explanation for information neglect in our experiment.

¹¹“Competition” here includes detailed empirical feedback on the competitors’ strategies and performance (see Slembeck and Tyran, 2004).

¹²Further recent studies on behavioral updating, which are less closely related, include the works of Ambuehl and Li (2018), Coutts (2018), and Martínez-Marquina et al. (2019), who look at various mistakes people make in Bayesian updating but not specifically at information neglect.

For the Monty-Hall set-up, Page (1998) shows that increasing the number of doors from three to at least ten increases the percentage of people who switch, but that switching in the many-doors problems does not mean participants would switch in a subsequent or simultaneously administered three-doors task. Instead of increasing the number of urns, we implement an intervention in which several balls are drawn from the same urn, which makes it easier to understand that the color of a drawn ball is an informative signal of the urn’s color.

2.4 Social learning

In the standard social-learning literature (*e.g.*, Weizsäcker, 2010) people have to guess the state of the world after observing a signal as well as the sequence of guesses of all earlier players. While in our paper, people have to guess the next draw from the urn, we ask about the probabilities of the states of the world, too. In that sense, participants in our experiment face a decision similar to that of the first player in the sequence of a standard social-learning experiment.

In De Filippis et al. (2022) and Angrisani et al. (2021), participants report two probabilities for the “good” state of the world, one after observing all earlier players’ reports and the second after additionally seeing their own private signal. While the authors focus on the update between the first and second report, our task could be likened to the first report of the second player, assuming that the first player would ‘reveal their signal’. Our paper shows that this cannot be taken for granted, as in our paper, up to 30% of the players do not ‘reveal their signal’, in the sense that they bet on the ‘opposite’ color. Note that this perspective does not consider people’s actual beliefs about others’ strategies. Rather, it shows what a ‘rational’ player’s belief about the representative player’s strategy should be.

Duffy et al. (2021) have their participants “choose between receiving a private, noisy signal about the current state or observing the past guesses of other subjects in the prior period,” to learn about whether non-optimal behavioral strategies of information choice. Our experiment informs about the value people might attach to receiving the noisy signal. Further information in this regard comes from the first players in three treatments of Kübler and Weizsäcker (2004). These players have to decide whether they want to acquire a signal—at a “high” cost, a “low” cost, or no cost at all, depending on the treatment—before betting on the state of the world. While the equilibrium prediction has players acquire the signal in all three treatments, only half, two thirds, and nine out of ten do so, respectively. While the signal take-up in the cost-free treatment is higher than what our results would suggest, our results nonetheless shed new light on why so many people are not willing to pay even a small fee for the signal in their other treatments.

3 Set-up

The main task in our experiment is a ‘Guessing Task’, variations of which can often be found in experimental studies of imperfect-information settings. In this task, participants have to guess the color of

another participant’s ball after being informed about their own type (receiving a colored ball). To this task, we add different parts in the different treatments to draw our conclusions about projection. In Treatment 1, the Guessing Task is followed by a Belief-Elicitation Task. In the Belief-Elicitation Task, participants are asked about the likelihood that the other participant’s ball of each of the two possible colors and about the probabilities of the different states of the world. We add this Belief-Elicitation Task as a diagnostic tool to understand the participants’ reasoning behind their Guessing-Task choices. In particular, the Belief-Elicitation Task would have allowed us to differentiate Bayesian updating from over-projection.

After finding evidence for information neglect in Treatment 1, we conduct two subsequent treatments as robustness checks. Also, in Treatments 2 and 3, the basic task remains the Guessing Task. In all treatments, the session starts with the Guessing Task. Then, in alternating fashion, participants face three different intervention tasks and the Guessing Task again. Our choice of intervention tasks makes participants more aware that they can learn something from the color of their own ball. Before Treatment 3, we demonstrate the whole process behind the Guessing Task in front of the participants: how an urn is randomly selected and how the balls are drawn with replacement. We demonstrate the process to ensure that participants’ failure to update is not due to the task being too abstract or poorly understood.

3.1 Experimental design

3.1.1 Guessing Task

Nature draws one of the two states of the world, or one of eight urns, as depicted in Figure 1. The states of the world can either be Dark (D) as on the left side or Light (L) as on the right side of Figure 1, with three balls in each urn. In D -urns, a fraction of $2/3$ of the balls are dark (d) and a fraction of $1/3$ is light (l). In L -urns, a fraction of $2/3$ of the balls are light (l) and a fraction of $1/3$ is dark (d). The possible priors are: $Pr(D) = 4/8$ or $Pr(D) = 5/8$, or $Pr(D) = 7/8$.¹³

First, the participants are randomly divided into pairs, and for each pair an urn is drawn. Nature’s draw of the urn is not revealed; thus, it is unknown in the remainder of the particular situation whether the balls are drawn from a D - or an L -urn. Then participant i receives a private signal s_i in the form of a ball drawn from the urn (see Figure 2). After replacement, participant j in the group receives a private signal s_j from the same urn. Participant i then has to guess whether participant j saw a light or a dark signal s_j .

¹³To control for color effects, we also included situations with $Pr(D) = 3/8$ and $Pr(D) = 1/8$. For simplicity, we pool equivalent (counter-balanced) situations throughout the paper, *e.g.*, receiving a d -signal in $Pr(D) = 5/8$ and receiving an l -signal in $Pr(D) = 3/8$. With a slight abuse of notation, we subsume both situations in the corresponding result tables under “ d -signal in $Pr(D) = 5/8$ ”.

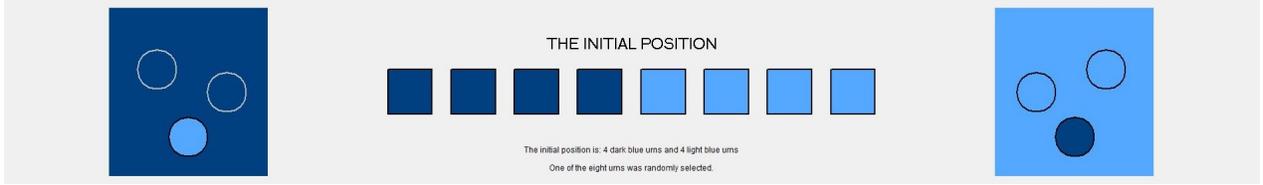


Figure 1: Representation of the prior distribution in the experiment in the 4/8-situation (four dark and four light urns).



Figure 2: Representation of the drawn signal in the experiment (one light ball).

In all situations of our laboratory experiment, a participant's own type is informative about the state of the world in which she and her partner both are in. Therefore, projecting one's own type onto the other participant can be rational. In the example in Figures 1 and 2, the prior probability $Pr(D)$ is $4/8 = 1/2$. After receiving the (informative) signal, it is thus rational for participant i to project her own signal onto the other participant by guessing that $s_j = s_i$. This is also the case for the 5/8 prior but not for the 7/8 prior. We let the participants perform the Guessing Task 18 times (six times per prior).¹⁴

In the Guessing-Task setting, we thus measure projection as follows. Rational projection occurs when the participants guess the same color as their signal in: (i) close to 100% of the cases when the correct Bayesian posterior is larger than 50% and (ii) close to 0% of the cases when the correct Bayesian posterior is smaller than 50% (allowing up to 10% trembles, we would expect $\geq 90\%$ same-color guesses for the first type of situation and $\leq 10\%$ for the second one). Over-projection can be distinguished from rational projection in two ways: by asking for participants' beliefs (see below), and by observing same-color guesses in close to 100% of the cases when the correct Bayesian posterior is smaller than 50%. Under-projection occurs if the participants *do not* guess the same color as their signal in close to 100% of the cases when the correct Bayesian posterior is larger than 50%. To identify information neglect as the source of under-projection, we need to ask participants about the probability they assign to the other player having the same type.

¹⁴The participants performed the Guessing Task in six sequences of the following order: 4/8-situation and 5/8-situation (and 7/8-situation in Treatment 1) followed by counter-balanced 4/8-situation and counter-balanced 5/8-situation (and counter-balanced 7/8-situation in Treatment 1). The counter-balanced situations included light and dark instead of dark and light balls and urns. The guesses were incentivized as follows: One out of all Guessing Tasks was randomly selected, and participants received €6 if the corresponding guess was correct.

3.1.2 Belief-Elicitation Task

After the Guessing Task, we elicit the participants' beliefs. We ask the participants for the probability that the other participant's ball is dark or light for each situation they may face, and for the probability that the urn is dark or light if they saw a dark or a light signal.¹⁵ Using these four questions for each situation, we can see whether participants update at all, whether they update their beliefs in the right or the wrong direction, and at which step they might fail to update correctly.

3.2 Interventions

After finding evidence for information neglect rather than projection in Treatment 1, we conduct two subsequent treatments as robustness checks. In Treatments 2 and 3, we include four iterations of the Guessing Task each. In each of the four Guessing-Task iterations of Treatments 2 and 3, the participants perform the Guessing Task 12 times (six times per prior, including counter-balancing). We try to facilitate projection (rational updating) by three interventions each. Each intervention is designed to get participants to realize that they could learn something about the other participant's type from their own type. To make space for these interventions we restrict our attention to fewer priors in these treatments. While we implement all three priors (4/8, 5/8 and 7/8) in Treatment 1, we implement only 4/8 and 5/8 in Treatments 2 and 3. Please refer to Table 1 for an overview of the three treatments.

Treatment 1	Treatment 2: Robustness checks	Treatment 3: Robustness checks
		<i>Intervention 0: Physical urns</i>
Guessing Task	Guessing Task	Guessing Task
<i>Belief elicitation I</i>	<i>Intervention 1: w/o replacement draws</i>	<i>Intervention 1: Belief elicitation III</i>
	Guessing Task	Guessing Task
	<i>Intervention 2: Belief elicitation II</i>	<i>Intervention 2: Strategy elicitation</i>
	Guessing Task	Guessing Task
	<i>Intervention 3: Strategy elicitation</i>	<i>Intervention 3: Several draws</i>
	Guessing Task	Guessing Task

Table 1: Overview of the three treatments.

¹⁵The participants performed the Belief-Elicitation Tasks in the following order: ball-probability questions for the 4/8-situation and for the 5/8-situation (and for the 7/8-situation in Treatment 1), followed by urn-probability questions for the 4/8-situation and for the 5/8-situation (and 7/8-situation in Treatment 1). The beliefs in Treatment 1 were incentivized as follows: Participants received €0,50 for each belief that did not differ by more than 2.5 percentage points from the correct probability. To give participants an additional incentive to report their beliefs precisely, they additionally received another €0,50 for each belief that did not differ by more than 0.5 percentage points from the correct probability in Treatment 2. In Treatment 3, we used the loss-aversion-adjusted scoring rule introduced by Offerman and Palley (2016). The belief-elicitation rules in Treatments 2 and 3 serve as robustness checks for the results from Treatment 1. They both are designed to avoid belief reports at 50% whenever the participant's true belief does not coincide with 50%.



Figure 3: Props used for the demonstration of the instructions.

Treatment 2. The three interventions are a repetition of the Guessing Task where the balls are drawn without replacement, the Belief-Elicitation Task, and a strategy-elicitation stage. Intervention 1, the Guessing Task without replacement, is meant to draw the participants’ attention to the fact that the balls are drawn from the same urn, and that there is replacement in the main task. This should bring the participants’ attention to the fact that drawing two balls of the same color is more likely than drawing two balls of different colors. In contrast to the Guessing Task where the participants mostly need to “go after their signals” or project, they have to “go against their signals” or under-project in Intervention 1.¹⁶

Intervention 2 of Treatment 2, the Belief-Elicitation Task, is designed to get participants to recognize that the urns have different posterior probabilities after observing their specific signal. Note that in Treatment 1, we include the Belief-Elicitation Task only as a diagnostic tool after the Guessing Task. In Treatment 2, we are interested in what happens to the Guessing-Task choices after the Belief-Elicitation part. In Intervention 3, we ask participants to write down the reasoning behind their reported probabilities from Intervention 2. Intervention 3 is intended to prompt participants to think again—and more deeply—about the task, giving them another opportunity to adjust their subsequent Guessing-Task choices.

Treatment 3. We start Treatment 3 by adding a demonstration using physical urns to the instructions at the beginning of the treatment. The participants first read on-screen instructions about the Guessing Task and the experiment in general, similarly as in Treatments 1 and 2. Then, right before the beginning of the first Guessing-Task iteration, the participants again learn about the process in the Guessing Task, only this time it is demonstrated with physical urns and commented on by the experimenter. For this purpose, we use a number of dark and light table-tennis-sized balls for demonstration. Furthermore, we use non-transparent bags as urns. They are black on the outside and dark or light on the inside. The appearance and numbers of the balls and bags are composed in a way that mimics the appearance and numbers described in the on-screen instructions right before (see Figure 3).

In Treatment 3, the three in-between interventions are—again—a Belief-Elicitation Task, the strategy-

¹⁶As for the Guessing Task, participants perform Intervention 1 involving only without-replacement draws 12 times: six times per prior, including counter-balancing.

elicitation task, and a guessing task with several draws. The rationale for the first two interventions is two-fold. First, they are meant to capture any changes induced by our pre-play demonstration of the situation. Second, this time we use the L-adjusted scoring rule introduced by Offerman and Palley (2016) for the Belief-Elicitation Task, which approximates a proper elicitation procedure without drawing the belief reports towards 50%.¹⁷

Intervention 3 of Treatment 3, the Guessing Task with several draws, should draw the participants' attention to the fact that signals are informative. After seeing the first-drawn signal from the randomly drawn urn (as in the Guessing Task), we show participants three additional with-replacement signals from this urn before they have to make their guess about the color of the other participant's ball. Seeing many balls of the same color should make participants realize that one of the urn types is more likely to be the randomly drawn one. Thus, we may expect participants to transfer this insight when they return to the single-signal Guessing Task afterwards.

Finally, note that we add an additional robustness check between Treatment 1 and Treatments 2 and 3. Namely, we provide more detailed decision screens and more elaborate instructions in Treatments 2 and 3, as depicted in Figure 4.¹⁸ In particular, the more elaborate screen in Figure 4(b) stresses the intermediate step of one particular urn having been drawn, as well as the fact that both balls (being positioned on the same line as the urn) come from the same urn.

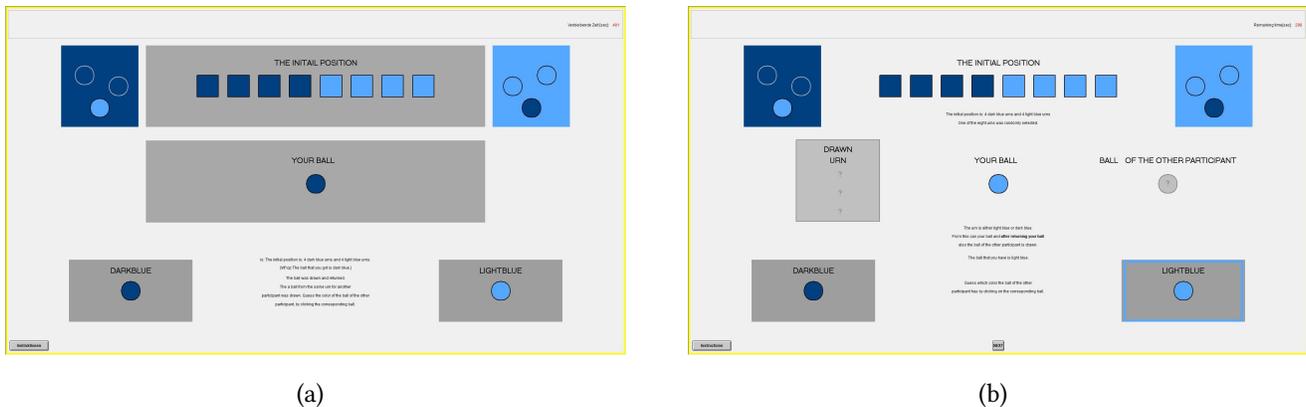


Figure 4: Changes in the decision screen from Treatment 1 (a) to Treatments 2 and 3 (b).

3.3 Procedure

We programmed all treatments in z-Tree (Fischbacher, 2007) and recruited all participants using ORSEE (Greiner, 2015). We conducted eight experimental sessions in total—four for Treatment 1 and two each for Treatments 2 and 3. In Treatment 1, 116 students (average age: 22 years old, standard deviation $sd=2$; 38% male) participated and earned an average payment of € 18 ($sd=5$), including a show-up fee

¹⁷To save time, we only elicited beliefs for the 4/8-prior in Treatment 3.

¹⁸Please see the instructions in Appendix B.

of € 6. The length of the sessions was about 90 minutes.¹⁹ In Treatment 2, 58 students (average age: 23 years, sd=2; 48% male) participated. They earned an average payment of € 13 (sd=4; show-up fee: € 6). Sessions lasted for approximately 70 minutes. In Treatment 3, 54 students participated (average age: 21 years, sd=2; 48% male) and earned an average payment of € 16 (sd=5; show-up fee: € 6). The length of the sessions was around 70 minutes.

4 Results

4.1 (No) Projection

4.1.1 Guessing Task

Starting with Treatment 1, the results show that many participants update their beliefs less than expected by theories of type projection or Bayesian updating, or they do not update at all. In the rightmost column of Table 2, we show how much projection we would expect if participants learned from their signals. For comparison, we show the actual projection frequency in the middle column *Frequency of projection* (we calculate the average frequency of projection for each participant, and then average over all participants). The column shows how often the participants actually guess the same color signal that they received.

Prior	Signal	Frequency of projection	Theoretical predictions
4/8-situation	d or l	76.0%	100% projecting (56%)
5/8-situation	d	94.5%	100% projecting (59%)
	l	36.1%	100% projecting (52%)
7/8-situation	d	94.7%	100% projecting (64%)
	l	6.9%	0% projecting (41%)

Table 2: Results of the Guessing Task in Treatment 1. The percentages in parentheses in the column to the right are the correct posterior probabilities for the other player having the same color.

Let us focus first on the simplest situation, $\Pr(D) = 4/8$. The result of 76% differs strongly from the 100% benchmark. Keeping in mind that random choice would mean 50% same-color guesses—and that projection would mean that participants *always* go with the signal—76% same-color guesses is a low rate. Another result from the non-50%-prior situations also indicates that the participants are biased

¹⁹Participants in the sessions for Treatment 1 also participated in a completely unrelated pilot treatment for another study after Treatment 1. Instructions for the additional treatment were distributed only after Treatment 1 was finished.

towards the prior information: The 36% in the 5/8-situation differs even more strongly from the 100% benchmark. The other percentages in the 5/8 and 7/8-situations, on the contrary, are quite close to the respective benchmarks. However, following the prior coincides with rational projection in these situations.

4.1.2 Belief-Elicitation Task

The results from the Belief-Elicitation Task complement the results from the Guessing Task in Treatment 1 and show even stronger evidence for information neglect and against projection, especially if we again focus on the simplest 4/8-situation as in Figure 5. The panel on the left shows the participants' beliefs that the other participant shares their type (the correct Bayesian posterior being 56% as indicated by the dashed line). The panel on the right shows the participants' beliefs that the urn's color corresponds to the color of their own ball. More than 60% of the participants report beliefs of 50% for both balls and urns, which coincides with the 4/8 prior. A small minority of participants report ball-beliefs of 67%. These participants seem to be 100% sure about the urn they are facing—or neglect that a 67% probability can only be correct if they were completely certain about the randomly drawn urn.

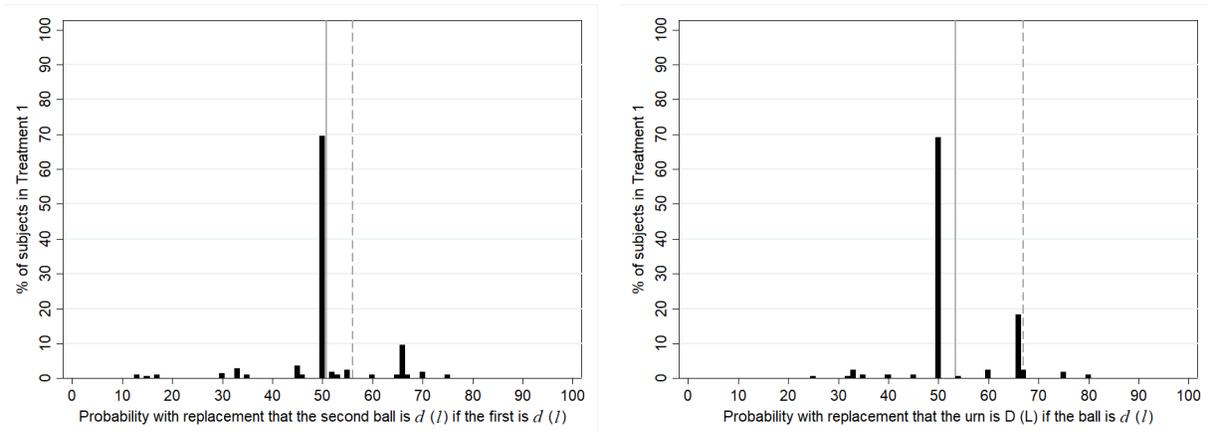


Figure 5: Results of the Belief-Elicitation Task in Treatment 1 (with ball-beliefs on the left and urn-beliefs on the right). Theoretical predictions are represented by the dashed line, average beliefs by the solid line.

In both cases, the beliefs differ from the theoretical prediction (dashed line; Wilcoxon signed-rank tests, p -values < 0.001). Roughly 70% of the participants report beliefs of 50% (for both the urn and the ball color; 55% of the participants report beliefs of 50% in all four probability questions), that is, more than half of the participants appear to believe that they cannot learn anything about the unknown selected urn from their received signal.

For the situations with asymmetric priors, fewer participants re-state the exact prior (see Figures

A1–A3 in Appendix A), which is unsurprising as the prior has to be computed first and does not coincide with 50%. To get an idea of the prevalence of information neglect and correct belief updates, we count (1) the number of stated beliefs (about both urn and ball colors) for the three situations of Treatment 1 which are close to the correct posterior (± 2.5 percentage points) and (2) the number of beliefs that are either close to the prior (± 2.5 percentage points) or exactly 50% (see Table 3).

Situation (prior)	Close to correct posterior (1)	Close to prior or exactly 50% (2)	Ratio of correct updates to info neglects (1)/(2)
4/8-situation	52	325	0.16
5/8-situation	57	194	0.29
7/8-situation	58	149	0.39

Table 3: Number of stated beliefs per situation that are (1) close to the correct posterior (± 2.5 percentage points) or (2) either close to the prior (± 2.5 percentage points) or exactly 50%, for the different situations. Within situations we pool the beliefs after d and l signals and the belief about urn color.

4.2 Nudging Interventions

Before we discuss the effects of our interventions, please recall that we would see “50%” of the participants making the rational choice if they were guessing light or dark completely randomly and, in accordance with the theoretical predictions, we would expect to see “100%” in all the iteration columns in 4 and 5. What we see instead is that the share stays below 82

The only intervention that helps to some extent is the Belief-Elicitation Task. In both Treatments 2 and 3, it raises the percentages in the Guessing Task by around 5-8 percentage points (Wilcoxon signed-rank tests between individual percentages of Bayesian choices in the Guessing Tasks before and after the Belief-Elicitation intervention yield $p = 0.051$ for the difference between the second and third iterations in Table 4, and $p = 0.014$ for that between the first and second iterations in Table 5). About 10-16% of the “information neglect” people seem to change their behavior from neglect to projection, and their performance does not deteriorate again afterwards. Otherwise, neither asking for strategy descriptions, nor having a task without replacement, nor having a task with multiple draws, nor physically enacting the situation helps (all p -values > 0.400 and > 0.477 , Wilcoxon signed-rank tests for Treatments 2 and 3, respectively). In light of these results, we conclude that none of the interventions improves the participants’ performance in the Guessing Task by much.

Note that it is not the case that people do not respond to any interventions—they do. They rarely guess their own signal when the ball is not replaced (see column 4 of Table 4), and they guess the correct (more likely) color most of the time when they see multiple signals (of the same color, see the bold-faced figures in Tables 8 and 9 below). However, they do not seem to transfer anything to our

main Guessing Task (compare column 5 of Table 4, and the final column of Table 5). In the following, we discuss behavior in the belief- and strategy-elicitation interventions in more detail.

Prior	Signal	1 st iter.	<i>Interv. 1</i>	2 nd iter.	<i>Interv. 2</i>	3 rd iter.	<i>Interv. 3</i>	4 th iter.
4/8-situation	<i>d</i> or <i>l</i>	71.0%	21.7%	69.0%	s.4.2.1	76.1%	s.4.2.2	73.0%
5/8-situation	<i>d</i>	89.8%	57.0%	88.3%	s.4.2.1	94.8%	s.4.2.2	94.2%
	<i>l</i>	30.3%	2.7%	22.3%	s.4.2.1	28.2%	s.4.2.2	24.6%

Table 4: Results of the four Guessing-Task iterations and three Interventions in Treatment 2.

Prior	Signal	1 st iter.	<i>Interv. 1</i>	2 nd iter.	<i>Interv. 2</i>	3 rd iter.	<i>Interv. 3</i>	4 th iter.
4/8-situation	<i>d</i> or <i>l</i>	71.3%	s.4.2.1	79.0%	s.4.2.2	80.2%	s.4.2.3	81.2%
5/8-situation	<i>d</i>	91.2%	s.4.2.1	90.6%	s.4.2.2	96.8%	s.4.2.3	96.8%
	<i>l</i>	37.1%	s.4.2.1	31.4%	s.4.2.2	32.0%	s.4.2.3	28.5%

Table 5: Results of the four Guessing-Task iterations and three Interventions in Treatment 3, after the Intervention 0 with physical urns.

4.2.1 Belief-Elicitation interventions

The situation concerning the reported beliefs in Treatment 2 barely improved compared to Treatment 1. Let us again first focus on the simplest 4/8-situation, as in Figure 6. Roughly 60% of the participants report beliefs of 50%. Again, a small minority of participants report beliefs of 66% or 67%.

In contrast, there is a notable change in reported beliefs in Treatment 3 compared to Treatments 1 and 2. Only 40.7% of the participants report beliefs of 50% in Treatment 3 for the ball-probabilities, as depicted in the left panel of Figure 7 below. A minority of participants (14.8% and 13.0% for *d* and *l* accordingly) report beliefs of 66% or 67%.

The largest difference between the treatments is in terms of the beliefs about the urns. Only 16.7% to 18.5% (in *l* and *d*, respectively, as depicted in the right-hand panel of Figure 7) report beliefs of 50% for the urn-probabilities in Treatment 3. Instead, 25.9% of participants report beliefs of 66% or 67%. This substantial change in stated beliefs is hard to explain under the assumption that the distributions of the true beliefs are the same under the different belief-elicitation rules.²⁰ Instead, the results suggest

²⁰Under the elicitation rules in Treatment 1 (2), risk- or loss-aversion can maximally explain a shift toward 50% of 2.5 (0.5) percentage points. Hence, adjusting for these potential biases could only explain a shift of the same magnitude away

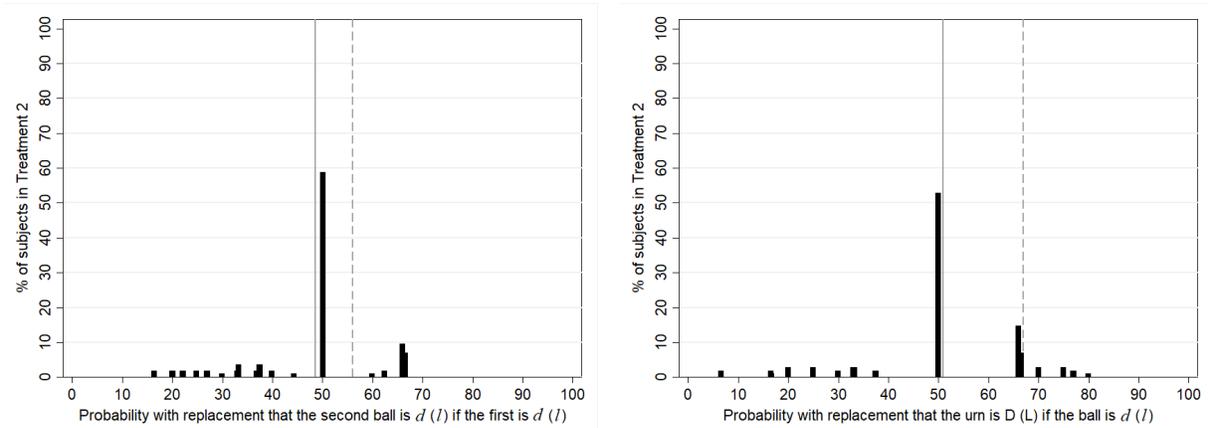


Figure 6: Results of the Belief-Elicitation Task in Treatment 2 (with ball-beliefs on the left and urn-beliefs on the right). Theoretical predictions are represented by the dashed line, average beliefs by the solid line.

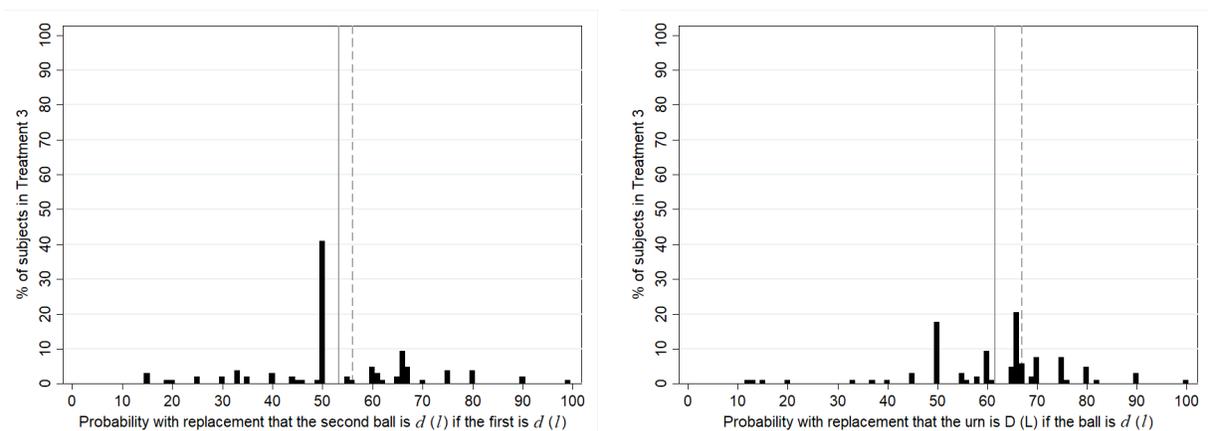


Figure 7: Results of the Belief-Elicitation Task in Treatment 3 (with ball-beliefs on the left and urn-beliefs on the right). Theoretical predictions are represented by the dashed line, average beliefs by the solid line.

that the participants' beliefs are influenced differently by the elicitation method itself. Potentially, the L-adjusted scoring rule nudges some participants to think differently about the problem.²¹

The results from the belief-elicitation under the L-adjusted scoring rule notwithstanding, it is unclear why so many participants still report 50% ball-probability beliefs in all three experiments. Related to this, it is also unclear why not more participants alter their behavior in the Guessing Task in any of the three treatments, even after the helping interventions. The participants' elicited strategies (Intervention 3 in Treatment 2 and Intervention 2 in Treatment 3) allow us to shed some light on both of these questions.

from 50%. However, if the stated beliefs under the L-adjusted rule in Treatment 3 were true, they would imply far bigger shifts.

²¹An alternative explanation would be that the physical demonstration of the situation had an influence on the beliefs. However, this seems unlikely as it had no influence on the behavior in the Guessing Task.

4.2.2 Strategy-elicitation interventions

How do the participants explain their answers in the Belief-Elicitation Task? We asked the participants to describe in free text their reasoning process behind the answers they provided in the Belief-Elicitation part regarding the 4/8-situation. We classify them into several categories and compare the answers of the participants who report 50% beliefs with those who report non-50% beliefs, as depicted in Tables 6 and 7. The reasoning strategies between these two groups differ substantially, as the participants who report non-50% beliefs most often mention learning from the signal in both Treatment 2 and Treatment 3 (e.g., “I assumed that a dark urn was chosen, as it contains more dark balls; so there are 3 balls, 2 of which are dark”).

Strategy categories	Subjects with 50% beliefs (in percent)		with non-50% beliefs (in percent)	
	Ball beliefs	Urn beliefs	Ball beliefs	Urn beliefs
1. Learn from the signal	-	-	46	48
2a. Equal probabilities	32	58	4	4
2b. Equal number of balls	32	29	-	-
2c. No difference after replacement	27	3	4	-
3a. Scarce information	3	7	17	18
3b. False or confused	6	3	29	30
<i>N</i>	34	31	24	27

Table 6: Results of the Intervention 3 (Strategy elicitation) in Treatment 2.

Strategy categories	Subjects with 50% beliefs (in percent)		with non-50% beliefs (in percent)	
	Ball beliefs	Urn beliefs	Ball beliefs	Urn beliefs
1. Learn from the signal	-	-	63	77
2a. Equal probabilities	36	20	3	-
2b. Equal number of balls	55	30	3	-
2c. No difference after replacement	-	-	-	-
3a. Scarce information	-	10	6	14
3b. False or confused	-	30	16	2
3c. Other	9	10	9	7
<i>N</i>	22	10	32	44

Table 7: Results of the Intervention 2 (Strategy elicitation) in Treatment 3.

The participants who report 50% beliefs, however, most often mention the following three reasons behind their answers. Firstly, they argue that the probability of drawing a dark urn is the same as drawing a light urn (e.g., “Each urn has the same probability of being chosen, hence the uniform distribution

50-50”). Secondly, they count the colors of the balls in both possible types of urns together (e.g., “A light and a dark urn constitute a total of 3 light balls and 3 dark balls; the probability is thus 50%”). Thirdly, they argue that there is no difference which signal has been drawn, given that the draws are with replacement (e.g., “The ball was put back in the urn, such that 6 balls were in the game all the time; of these, three were light and three were dark”). Some of the free-text answers were difficult to categorize meaningfully, either due to scarce content in the written texts, confusion, or other reasons.

The strategies mentioned in Treatment 3 can be categorized similarly to the strategies in Treatment 2. However, the shares of the respective categories differ. In particular, the participants reporting non-50% beliefs mention learning from the signal more often in Treatment 3 than in Treatment 2.

4.2.3 The several-draws intervention

In Intervention 3 of Treatment 3, we provide the participants with three additional signals from this urn before they make their guess about the color of the other participant’s signal. The participants do indeed appear to project their signals more after receiving three or four balls of the same color than after receiving one ball in the single-draw Guessing Task (compare Tables 8 and 9 with Table 5 above). That is, they seem to understand that the same urn is used for all the four draws—and that this urn is then also used for the other participant in their group. Note, however, that we do not see 100% projection rates even when participants get to see three or four balls of the same color.²²

Signal	<i>d</i>	<i>l</i>	Prediction	Probability
4 <i>d</i> , 0 <i>l</i>	25	5	100%	65%
%	83.3%	16.7%		
3 <i>d</i> , 1 <i>l</i>	68	7	100%	60%
%	90.7%	9.3%		
2 <i>d</i> , 2 <i>l</i>	56	60	-	50%
%	48.3%	51.7%		
1 <i>d</i> , 3 <i>l</i>	8	70	0%	40%
%	10.3%	89.7%		
0 <i>d</i> , 4 <i>l</i>	2	23	0%	35%
%	8.0%	92.0%		

Table 8: Results of the Intervention 3 with four ball draws, 4/8-situation.

²²These data suggest that around 10% of the sample are, most likely, not paying any attention at all (corresponding to the number of people who make the wrong choice even in the “several-draws” task).

Signal	<i>d</i>	<i>l</i>	Prediction	Probability
4 <i>d</i> , 0 <i>l</i>	30	3	100%	65%
%	90.9%	9.1%		
3 <i>d</i> , 1 <i>l</i>	55	5	100%	62%
%	91.7%	8.3%		
2 <i>d</i> , 2 <i>l</i>	48	13	100%	54%
%	78.7%	21.3%		
1 <i>d</i> , 3 <i>l</i>	11	39	0%	43%
%	22.0%	78.0%		
0 <i>d</i> , 4 <i>l</i>	28	92	0%	36%
%	23.3%	76.7%		

Table 9: Results of the Intervention 3 with four ball draws, 5/8-situation.

4.3 A single-agent control treatment

Having documented a robust information-neglect phenomenon within the standard two-player setup, we wanted to test whether the results would carry over to a setting in which participants had to guess the color of their *own* second ball after seeing the first one. This control treatment tests for the influence of having another player as well as of making the sequential nature of the ball draws even more salient. Because we ran the sessions in early March 2022 (when the covid-19 omicron-wave was racing for its peak), we ran the sessions online using z-Tree unleashed Duch et al., 2020. Participants were drawn from the LakeLab participant pool at the University of Konstanz In order to have a clean treatment comparison, we also ran a baseline treatment, randomizing participants into the two treatments (56 single-agent participants vs. 54 baseline participants) within-sessions. The setup consisted of our standard Guessing task, followed by “Belief elicitation III” (as in Treatment 3) and another repetition of the Guessing task.

While there was a lower projection rate than in our previous data in the first Guessing Task in the baseline treatment (58.5% projection), the single-agent treatment led to a projection rate that was very similar to our earlier laboratory sessions (74.4% vs. 76.0%, 71.0%, and 71.3% in Treatments 1-3). This apparent difference was still present after the belief-elicitation task (67.0% vs. 77.4%, respectively). Again, we document a moderate improvement after the Offerman-Palley (2016) belief-elicitation task in the baseline treatment (58.5% vs. 67.0%), but not so much in the single-agent treatment (74.4% vs. 77.4%). When we look at the beliefs, however, there are no treatment differences at all. The average belief that the second ball has the same color as the first ball is 49.5% in the baseline treatment and 48.7% in the single-agent treatment. Moreover, the percentages of beliefs that are exactly “50%” are 39.8% and 40.2%, respectively. Hence, the information-neglect result also carries over to a scenario without another player.

5 Discussion and conclusions

The motivation of our investigation was twofold. On the one hand, we started from the observation of two growing strands of the literature that make starkly contrasting predictions regarding the use of information. The first strand is the projection literature which suggests that people typically infer too much from their own type about the types of other players. The second strand of the literature that we dub “information neglect” literature starts from the opposite end and assumes some form of information neglect. On the other hand, one of the central elements of virtually every imperfect-information model (and many experimental studies thereof) presupposes a rational use of new information. Hence, it is important to examine this element and explore which of the behavioral effects would dominate in such a setup, over-projection or information neglect—or whether they might cancel each other out in the average.

In our very simple setting, we find no evidence for (over-)projection. Instead, we find strong evidence for information neglect. Importantly, we demonstrate that people do not simply underweight new information uniformly but neglect the information contained in their own type altogether. Focusing on the situation with a 4/8-prior, for example, we see that up to 70% of the participants report beliefs that imply that they cannot learn anything from their signal.

These results raise the question of what the reason for so much information neglect might be. Our single-agent control treatment demonstrates that it is not the social nature of the situation of our main treatments. While the experimental social-learning literature (e.g., Weizsäcker, 2010; Eyster and Rabin, 2010) shows how one’s own information is weighted differently than (inferred) information of others, in our study subjects neglect information irrespective of whether there is another player or not.

It is more fruitful turn back to the literature on the Monty-Hall problem in search for an answer to our question. Burns (2017) argues that people follow a “no-change principle”, under which probabilities remain the same after the host opens one of the doors. He calls this phenomenon “no-change principle” because the probabilities that he elicits from the participants suggest a reasoning that implies that “there has been no change to these doors, so they maintain their equal status” (Burns, 2017, p.1701). The participants in our experiment seem to make the same mistake. Their answers to our questions regarding their reasoning clearly reflect a “no-change principle”, too. *The ball has been put back, hence nothing has changed and the probabilities must be the same as the prior probabilities* (categories 2a–2c in Tables 6 and 7), appears to be the faulty logic that many of our participants follow. In this context, our results suggest that information neglect can occur also in set-ups that are much simpler than the Monty-Hall problem, including typical set-ups that involve inferences about others from one’s own type.

Acknowledgments

We would like to thank Jordi Brandts, Yves Breitmoser, Carlos Alós-Ferrer, Dorothea Kübler, Urs Fischbacher and the research group at the Thurgau Institute of Economics (TWI) for their helpful comments, as well as the participants in a number of conferences and workshops. This work was supported by the Foundation for Science and Research of the Canton Thurgau in Switzerland.

References

- Agranov, M., Goeree, J. K., Romero, J., and Yariv, L. (2017). What makes voters turn out: The effects of polls and beliefs. *Journal of the European Economic Association*, 16(3):825–856.
- Akerlof, G. A. (1970). The market for “lemons”: Quality uncertainty and the market mechanism. *Quarterly Journal of Economics*, 84(3):488–500.
- Aksoy, O. and Weesie, J. (2012). Beliefs about the social orientations of others: A parametric test of the triangle, false consensus, and cone hypotheses. *Journal of Experimental Social Psychology*, 48(1):45–54.
- Al-Nowaihi, A. and Dhami, S. (2015). Evidential equilibria: Heuristics and biases in static games of complete information. *Games*, 6(4):637–676.
- Ambuehl, S. and Li, S. (2018). Belief updating and the demand for information. *Games and Economic Behavior*, 109:21–39.
- Angrisani, M., Guarino, A., Jehiel, P., and Kitagawa, T. (2021). Information redundancy neglect versus overconfidence: A social learning experiment. *American Economic Journal: Microeconomics*, 13(3):163–97.
- Bar-Hillel, M. (1980). The base-rate fallacy in probability judgments. *Acta Psychologica*, 44(3):211–233.
- Benjamin, D. J. (2019). *Errors in probabilistic reasoning and judgment biases*, volume 2. Elsevier B.V.
- Blanco, M., Engelmann, D., Koch, A., and Normann, H. (2014). Preferences and beliefs in a sequential social dilemma: A within-subjects analysis. *Games and Economic Behavior*, 87:122–135.
- Bouton, L., Castanheira, M., and Llorente-Saguer, A. (2016). Divided majority and information aggregation: Theory and experiment. *Journal of Public Economics*, 134:114–128.
- Breitmoser, Y. (2019). Knowing me, imagining you: Projection and overbidding in auctions. *Games and Economic Behavior*, 113:423–447.

- Brusco, S. (1998). Unique implementation of the full surplus extraction outcome in auctions with correlated types. *Journal of Economic Theory*, 80(2):185–200.
- Burns, B. D. (2017). Equiprobability principle principle or “no change” principle ? Examining reasoning in the Monty Hall Dilemma using unequal probabilities. *39th Annual Conference of the Cognitive Science Society*, pages 1697–1702.
- Conlin, M., O’Donoghue, T., and Vogelsang, T. (2007). Projection bias in catalog orders. *American Economic Review*, 97(4):1217–1249.
- Coutts, A. (2018). Good news and bad news are still news: experimental evidence on belief updating. *Experimental Economics*, pages 1–27.
- Cremer, J. and McLean, R. P. (1985). Optimal selling strategies under uncertainty for a discriminating monopolist when demands are interdependent. *Econometrica*, 53(2):345–361.
- Cremer, J. and McLean, R. P. (1988). Full extraction of the surplus in Bayesian and dominant strategy auctions. *Econometrica*, 56(6):1247.
- Danz, D., Kübler, D., Mechtenberg, L., and Schmid, J. (2015). On the failure of hindsight-biased principals to delegate optimally. *Management Science*, 61(8):1938–1958.
- Danz, D., Madarasz, K., and Wang, S. (2018). The biases of others: Anticipating information projection in an agency setting.
- Dawes, R. (1989). Statistical criteria for establishing a truly false consensus effect. *Journal of Experimental Social Psychology*, 25(1):1–17.
- De Filippis, R., Guarino, A., Jehiel, P., and Kitagawa, T. (2022). Non-bayesian updating in a social learning experiment. *Journal of Economic Theory*, 199:105188. Symposium Issue on Ambiguity, Robustness, and Model Uncertainty.
- Duch, M. L., Grossmann, M. R., and Lauer, T. (2020). z-tree unleashed: A novel client-integrating architecture for conducting z-tree experiments over the internet. *Journal of Behavioral and Experimental Finance*, 28:100400.
- Duffy, J., Hopkins, E., and Kornienko, T. (2021). Lone wolf or herd animal? information choice and learning from others. *European Economic Review*, 134:103690.
- Edwards, W. (1968). Conservatism in human information processing. *Formal Representation of Human Judgment*.
- El-Gamal, M. A. and Grether, D. M. (1995). Are people bayesian? uncovering behavioral strategies. *Journal of the American statistical Association*, 90(432):1137–1145.

- Enke, B. and Zimmermann, F. (2019). Correlation Neglect in Belief Formation. *Review of Economic Studies*, 86(1):313–332.
- Esponda, I. and Vespa, E. (2014). Hypothetical thinking and information extraction in the laboratory. *American Economic Journal: Microeconomics*, 6(4):180–202.
- Eyster, E. and Rabin, M. (2005). Cursed equilibrium. *Econometrica*, 73(5):1623–1672.
- Eyster, E. and Rabin, M. (2010). Naïve herding in rich-information settings. *American Economic Journal: Microeconomics*, 2(4):221–243.
- Fehrler, S. and Hughes, N. (2018). How transparency kills information aggregation: Theory and experiment. *American Economic Journal: Microeconomics*, 10(1):181–209.
- Fischbacher, U. (2007). Z-Tree: Zurich toolbox for ready-made economic experiments. *Experimental Economics*, 10(2):171–178.
- Fischhoff, B. and Beyth-Marom, R. (1983). Hypothesis evaluation from a bayesian perspective. *Psychological Review*, 90(3):239.
- Friedman, D. (1998). Monty Hall’s three doors: Construction and deconstruction of a choice anomaly. *American Economic Review*, 88(4):933–946.
- Gilbert, D., Pinel, E., Wilson, T., Blumberg, S., and Wheatley, T. (1998). Immune neglect: A source of durability bias in affective forecasting. *Journal of Personality and Social Psychology*, 75(3):617.
- Gillman, L. (1992). The car and the goats. *American Mathematical Monthly*, 99(1):3–7.
- Goeree, J. K. and Großer, J. (2007). Welfare reducing polls. *Economic Theory*, 31(1):51–68.
- Greiner, B. (2015). Subject pool recruitment procedures: Organizing experiments with ORSEE. *Journal of the Economic Science Association*, 1(1):114–125.
- Grether, D. M. (1980). Bayes Rule as a Descriptive Model: The Representativeness Heuristic. *The Quarterly Journal of Economics*, 95(3):537.
- Guarnaschelli, S., McKelvey, R. D., and Palfrey, T. R. (2000). An experimental study of jury decision rules. *American Political Science Review*, 94(2):407.
- Hanna, R., Mullainathan, S., and Schwartzstein, J. (2014). Learning through noticing: Theory and evidence from a field experiment. *The Quarterly Journal of Economics*, 129(3):1311–1353.
- Holt, C. A. and Smith, A. M. (2009). An update on Bayesian updating. *Journal of Economic Behavior and Organization*, 69(2):125–134.

- Holt, C. A. and Smith, A. M. (2016). Belief elicitation with a synchronized lottery choice menu that is invariant to risk attitudes. *American Economic Journal: Microeconomics*, 8(1):110–139.
- Kahneman, D. and Tversky, A. (1972). Subjective probability: A judgment of representativeness. *Cognitive Psychology*, 3(3):430–454.
- Kluger, B. D. and Wyatt, S. B. (2004). Are judgment errors reflected in market prices and allocations? Experimental evidence based on the Monty Hall problem. *Journal of Finance*, 59(3):969–997.
- Kübler, D. and Weizsäcker, G. (2004). Limited depth of reasoning and failure of cascade formation in the laboratory. *The Review of Economic Studies*, 71(2):425–441.
- Levy, G. and Razin, R. (2015). Correlation neglect, voting behavior, and information aggregation. *American Economic Review*, 105(4):1634–1645.
- Madarász, K. (2012). Information projection: Model and applications. *The Review of Economic Studies*, 79(3):961–985.
- Madarász, K. (2015). Projection equilibrium: Definition and applications to social investment and persuasion. *Mimeo*, (January):1–48.
- Martínez-Marquina, A., Niederle, M., and Vespa, E. (2019). Failures in contingent reasoning: The role of uncertainty. *American Economic Review*, 109(10):3437–3474.
- Meehl, P. E. and Rosen, A. (1955). Antecedent probability and the efficiency of psychometric signs, patterns, or cutting scores. *Psychological Bulletin*, 52(3):194.
- Messé, L. and Sivacek, J. (1979). Predictions of others' responses in a mixed-motive game: Self-justification or false consensus? *Journal of Personality and Social Psychology*, 37(4):602.
- Miller, J. B. and Sanjurjo, A. (2019). A bridge from monty hall to the hot hand: The principle of restricted choice. *Journal of Economic Perspectives*, 33(3):144–162.
- Nalebuff, B. (1987). Puzzles. *Journal of Economic Perspectives*, 1(1):157–63.
- Offerman, T. and Palley, A. (2016). Lost in translation: An off-the-shelf method to recover probabilistic beliefs from loss-averse agents. *Experimental Economics*, 19(1):1–30.
- Offerman, T., Sonnemans, J., and Schram, A. (1996). Value orientations, expectations and voluntary contributions in public goods. *Economic Journal*, pages 817–845.
- Page, S. E. (1998). Let's make a deal. *Economics Letters*, 61:175–180.
- Palacios-Huerta, I. (2003). Learning to open Monty Hall's doors. *Experimental Economics*, 6(3):235–251.

- Read, D. and Van Leeuwen, B. (1998). Predicting hunger: The effects of appetite and delay on choice. *Organizational Behavior and Human Decision Processes*, 76(2):189–205.
- Ross, L., Greene, D., and House, P. (1977). The false consensus effect: An egocentric bias in social perception and attribution processes. *Journal of Experimental Social Psychology*, 13(3):279–301.
- Selvin, S. (1975a). A problem in probability [Letter to the editor]. *The American Statistician*, 29(1):67.
- Selvin, S. (1975b). On the Monty Hall problem [Letter to the editor]. *The American Statistician*, 29(3):134.
- Siddiqi, H. (2009). Is the lure of choice reflected in market prices? Experimental evidence based on the 4-door Monty Hall problem. *Journal of Economic Psychology*, 30(2):203–215.
- Slembeck, T. and Tyran, J.-R. (2004). Do institutions promote rationality? An experimental study of the three-door anomaly. *Journal of Economic Behavior & Organization*, 54(3):337–350.
- Taylor, C. R. and Yildirim, H. (2010). Public information and electoral bias. *Games and Economic Behavior*, 68(1):353–375.
- Tolvanen, J. (2017). Political ambiguity in a world of heterogeneous beliefs. *Mimeo*.
- Tor, A. and Bazerman, M. H. (2003). Focusing failures in competitive environments: Explaining decision errors in the Monty Hall game, the acquiring a company problem, and multiparty ultimatums. *Journal of Behavioral Decision Making*, 16(5):353–374.
- Weizsäcker, G. (2010). Do We Follow Others when We Should? A Simple Test of Rational Expectations. *American Economic Review*, 100(5):2340–2360.

Appendices (for online publication)

A Further results

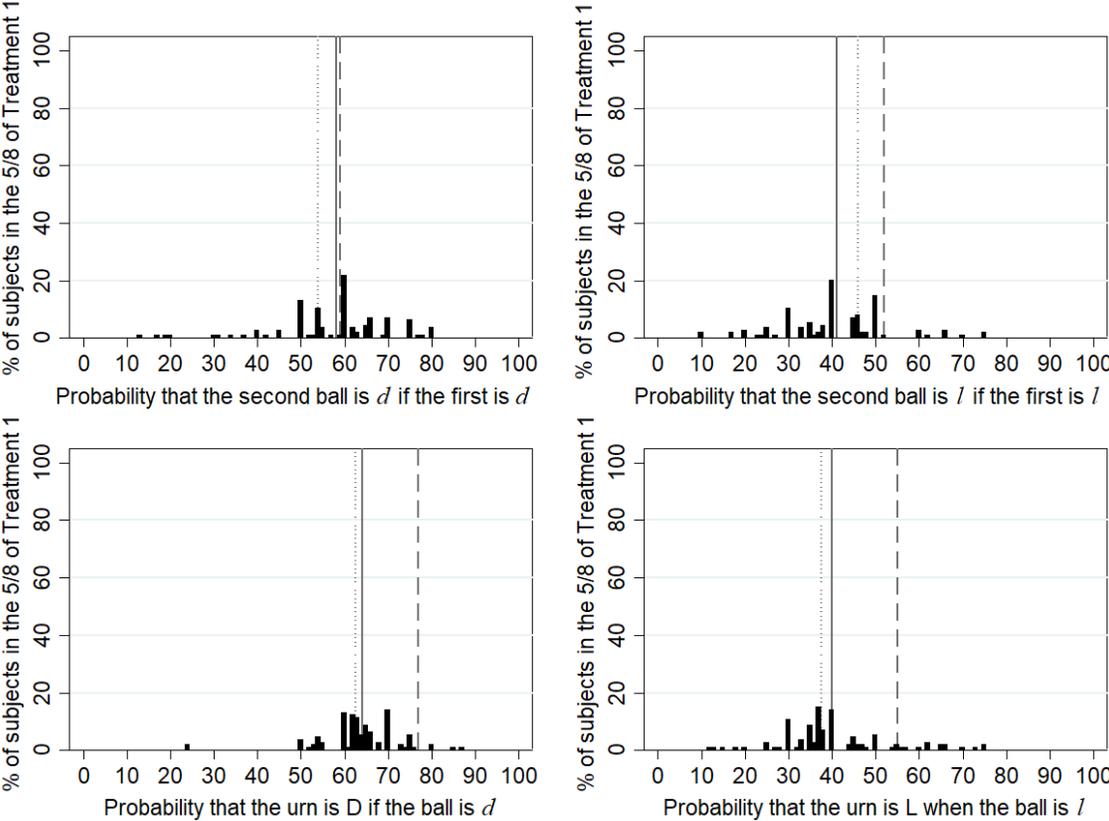


Figure A1: Results of the Belief-Elicitation Task in Treatment 1 (with ball-probability beliefs on the left and urn-probability beliefs on the right). Theoretical predictions are represented by the dashed, priors by the dotted, and average beliefs by the solid line.

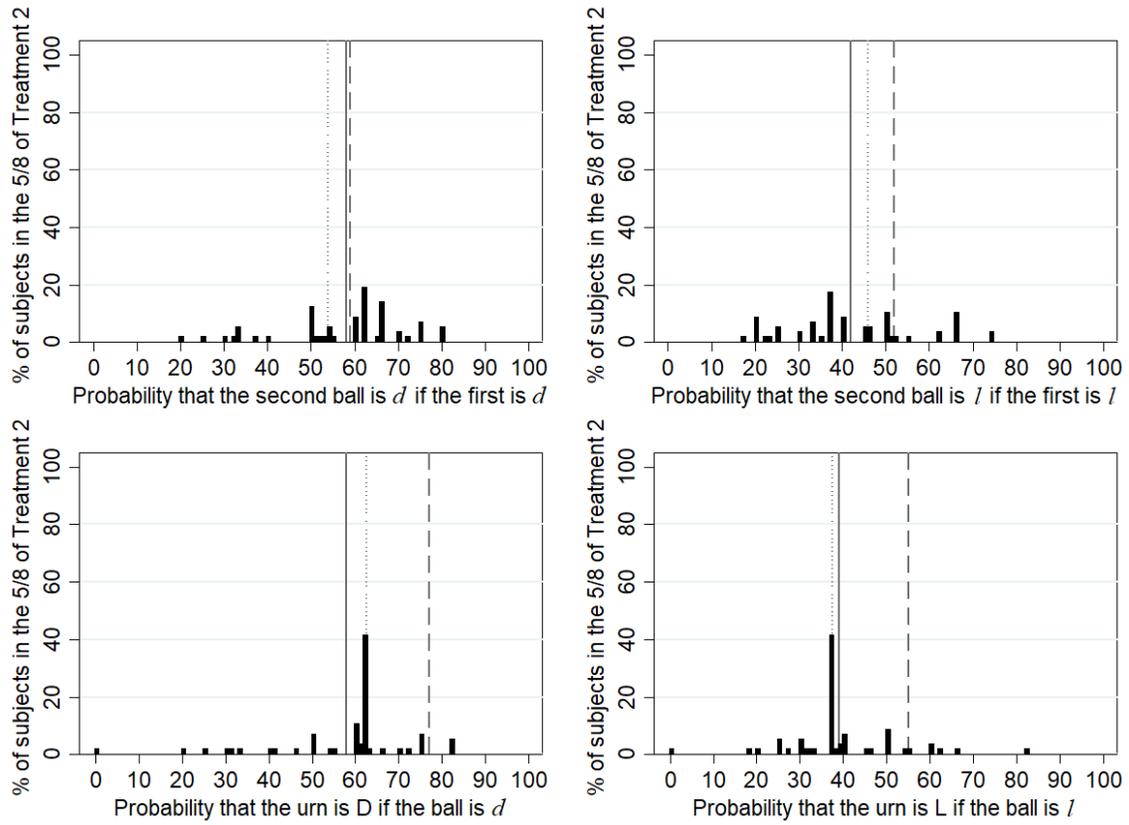


Figure A2: Results of the Belief-Elicitation Task in Treatment 2 (with ball-probability beliefs on the left and urn-probability beliefs on the right). Theoretical predictions are represented by the dashed, priors by the dotted, and average beliefs by the solid line.

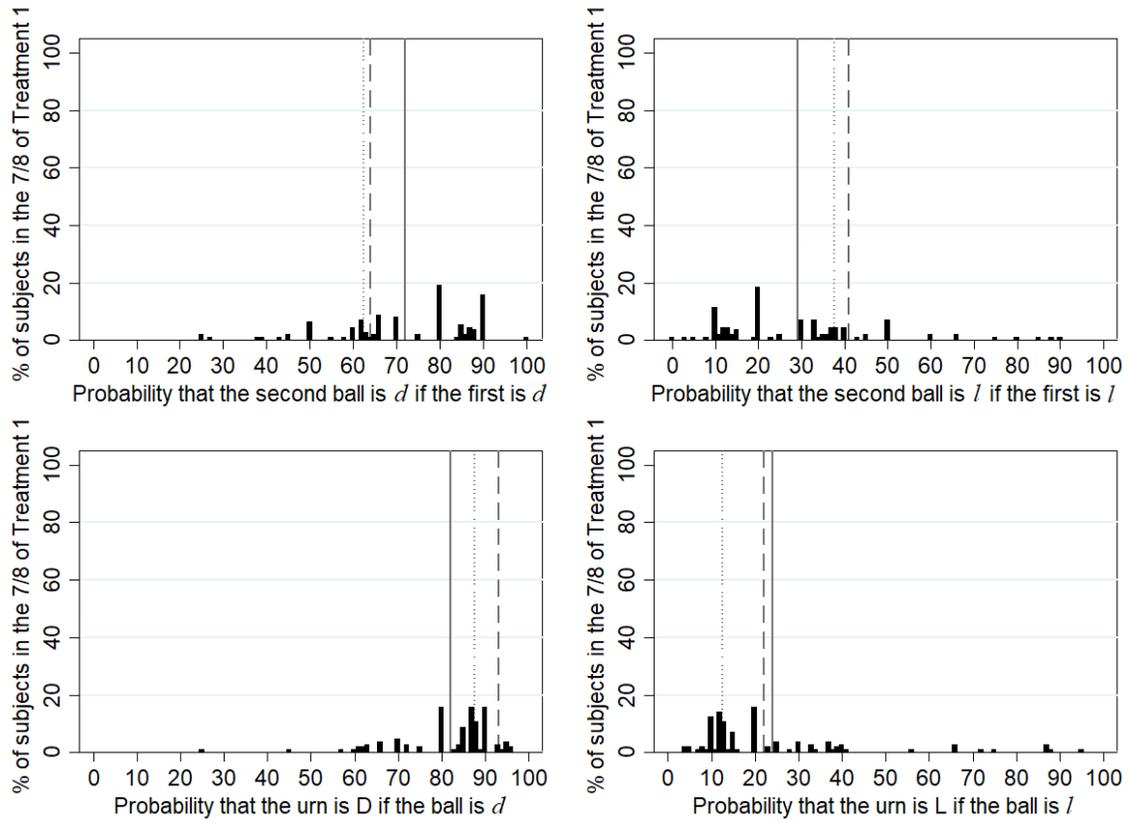


Figure A3: Results of the Belief-Elicitation Task in Treatment 1 (with ball-probability beliefs on the left and urn-probability beliefs on the right). Theoretical predictions are represented by the dashed, priors by the dotted, and average beliefs by the solid line.

B Instructions

The following instructions (translated from German) correspond to the most frequently used text versions over the three treatments. In the meantime, there were a few adjustments between the treatments. Firstly, we changed the layout of the decision screen and the instruction texts before Treatment 2. Secondly, we changed the colors of the light blue and dark blue balls to white and blue balls, respectively, before Treatment 3.

B.1 Instructions for the Guessing Task throughout the three treatments

Overview

Welcome to this experiment. We ask you not to speak to the other participants and to turn off your mobile phone and other mobile technical devices during the experiment. For taking part in today's experiments, you will be paid in cash at the end. The amount of the payout depends partly on chance and partly on your decisions. It is thus important that you carefully read and understand the instructions before starting the experiment.

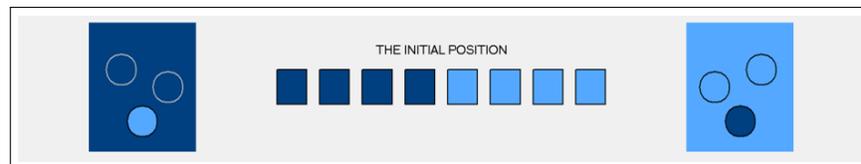
Today's experiment includes **six parts**, each comprising several rounds. In the end, several randomly-drawn rounds will be paid out. From the parts 1, 3, 5 and 6 of the experiment, **one** round is randomly selected and paid out. In addition, **one** randomly drawn round from part 2 and **all** rounds of part 4 are paid out. The not-drawn rounds will not be paid out.

Your payout is based on the points you earned in the rounds, plus 6 euro for completing a subsequent questionnaire. The conversion of the points into euro is done as follows. Each point is worth 50 cents, such that: **10 points = 5.00 euro**. Each participant gets paid privately so that other participants cannot see how much you have earned.

Composition of the experiment

This experiment includes six different parts. The 1st, 2nd, 3rd, 5th and 6th parts each comprise 12 identical rounds. In each round, a decision must be made which can be correct or incorrect. The tasks in these parts are similar in structure. The 4th part consists of eight identical rounds. In each round, there are multiple-choice questions to be answered. **Before starting each part of the experiment, you will read the instructions for that part.** On the next two pages, you will find the instructions for **part 1**.

The initial position



In each round, there are a total of 8 urns of balls. Each urn is either light blue and contains two light blue and one dark blue ball or dark blue and contains two dark blue and one light blue ball.

There can be three different initial positions. Either 3, 4 or 5 of the 8 urns can be dark blue. The remaining urns are light blue. The initial position is displayed at the top of the screen in each round. In the example shown above, you see a dark blue urn on the left, a light blue urn on the right and an initial position in the center with 4 light blue and 4 dark blue urns.

The drawn urn

As you already know, there are two types of urns: the dark blue urns and the light blue urns. The dark blue urn contains 1 light blue and 2 dark blue balls, the light blue urn 2 light blue and 1 dark blue. In each round, one of the 8 urns is randomly drawn from the initial position by the computer. Each urn has the same probability of being drawn. However, you will not know whether the drawn urn is light blue or dark blue.

Draw of the ball(s)

In each round, one or more balls are drawn randomly from the selected urn. You will see the color of the randomly drawn ball(s) on the screen. Whenever several balls are drawn, they are drawn one after the other and immediately returned. Another player will also receive one or more randomly drawn balls from the selected urn. Each ball in the drawn urns always has the same probability of being selected. This applies to the other participant as well as to you. Each participant only sees the color of their ball(s), not the other participant's.

The tasks



In the first part of the experiment, exactly one ball is displayed to each participant. You then have to guess the color of the ball that was drawn for the other participant. Before you make your decision, you will be informed of the (for both identical) initial position and the color of your drawn ball. In the example shown, it is dark blue.

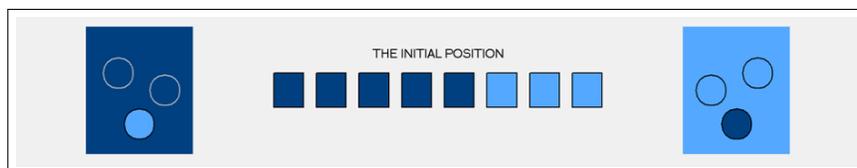
In addition to the normal rounds described so far, there are situations in which neither the urn nor the ball is randomly drawn. The predetermined situations amount for less than 1/6 of all rounds. You cannot distinguish these rounds from the normal rounds. For your payout, however, only those rounds in which randomness decides which urn is drawn and which ball is drawn from this urn are relevant.

The payout

From the 12 rounds of the second part, one round is randomly selected. One decision is randomly selected from the 48 rounds of parts 1, 3, 5 and 6 of the experiments. The predefined situations cannot be selected. If you have given the correct answer, that is, the color of the other participant's ball, you receive 12 points. The payment details for the 4th part can be found at the beginning of part 4.

Questions?

We kindly ask you to take the time to look over the instructions carefully. If you have any questions, please raise your hand. An experimenter will then come to you.



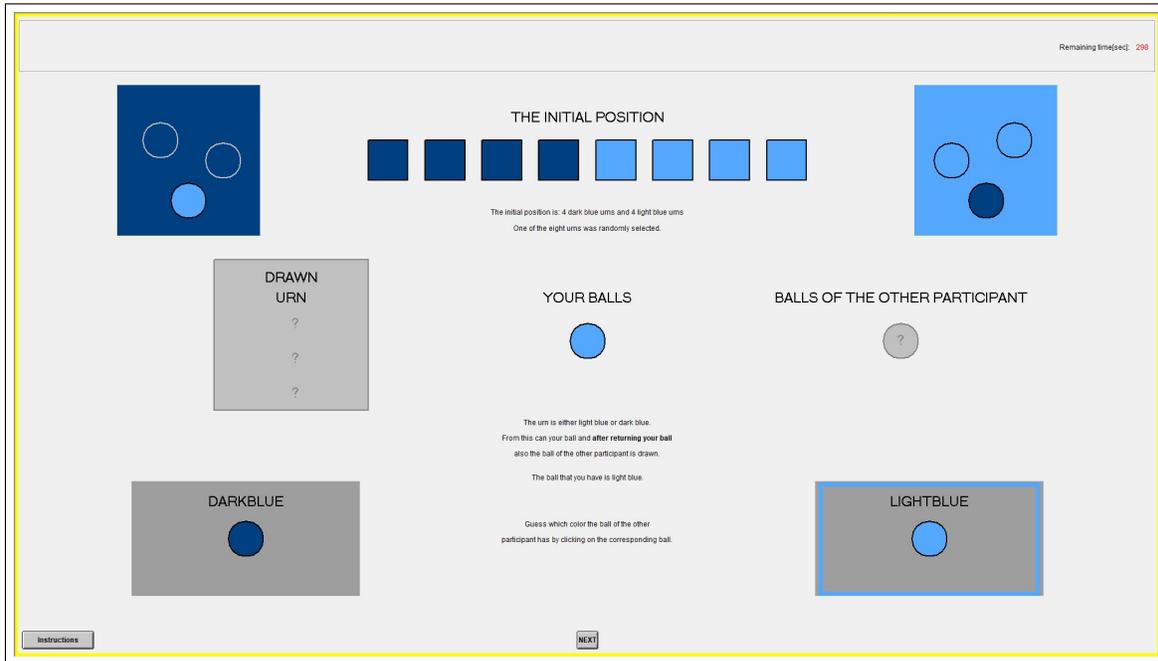
Please answer the following questions about the experiment.

1. How many light blue balls are there in a light blue urn? [...]
2. In the above example, with what probability is a light blue urn drawn (in %)? [...]
3. How many rounds will be paid out in the second part? [...]

Part 1 [Part 3 / Part 5 / Part 6] of the experiment

In each round of the 1st [3rd, 5th, 6th] part, a ball is randomly drawn from the selected urn. The color of the randomly drawn ball is displayed on the screen. The ball is drawn and returned immediately. Another player also receives a randomly drawn ball from the same selected urn after your ball has been returned.

In the [first] part, exactly one ball is shown to each participant. You then have to guess the color of the ball that was drawn for the other participant. Before making your decision, you will be informed of the (for both identical) initial position and the color of your drawn ball.



Payout [at the end of the experiment]

Participant [...] will now determine the payout-relevant decision by throwing a dice. Please wait in your seat while the payout-relevant initial positions and situations are rolled.

The first dice roll determines the payout-relevant initial position in the [first] part of the experiment: "1" for the initial position with 4, "2" for the initial position with 5 dark blue urns, also "4" for the initial position with 4, "5" for the initial position with 5 light blue urns. If a "3" or "6" is rolled, the roll has to be repeated.

The second dice roll determines the payout-relevant part of the 1st, 3rd, 5th 6th parts: "1" for the first part, "3" for the third part, "5" for the fifth part, "6" for the sixth part. If a "2" or "4" is rolled, the roll has to be repeated.

The third dice roll determines the payout-relevant initial position in the second part of the experiment: "1" for the initial position with 4, "2" for the initial position with 5 dark blue urns, also "4" for the initial position with 4, "5" for the initial position with 5 light blue urns. If a "3" or "6" is rolled, the roll has to be repeated.

The fourth dice roll determines the payout-relevant situation under the previously diced initial position in the [first] part of the experiment: "1" and "4" for situation 1, "2" and "5" for situation 2, "3" and "6" for situation 3. And the fifth dice throw determines the payout-relevant situation under the previously diced initial position in the second part of the experiment: "1" and "4" for situation 1, "2" and "5" for situation 2, "3" and "6" for situation 3.

In the unlikely event that a predefined situation is rolled, the computer will randomly select one of the other situations.

B.2 Interventions

B.2.1 Intervention 1 of Treatment 2

Part 2 of the experiment

In each round of the **2nd part**, a ball is randomly drawn from the selected urn. The color of the randomly drawn ball is displayed on the screen. The ball is drawn and **NOT returned**. Another player also receives a randomly drawn ball after your ball has been returned to the same selected urn.

Each participant thus sees exactly one ball. You then have to guess the color of **the next** ball that was drawn for the other participant. Before making your decision, you will be informed of the (for both identical) initial position and the color of your drawn ball.

B.2.2 Intervention 2 of Treatment 2

Part 4 of the experiment

In the **4th part**, your task is to answer questions. The questions relate to different scenarios. The scenarios will be similar to those in the first, second and thirds parts.

Your task is to specify the probability of a dark blue or a light blue ball being drawn in each scenario, or the probability that the urn is dark blue or light blue in each scenario.

In the **4th part**, you get 1 point for each correct answer. An answer is considered to be correct if the probability that you give is a maximum of 2.5 percentage points away from the correct probability. That is, if the correct probability is x (in %), you get 1 point for each response that is not greater than $x+2.5$ and not smaller than $x-2.5$. In addition, you get 1 point for each answer that is not greater than $x+0.5$ and not smaller than $x-0.5$. Your income is then converted into euro and paid out privately.

In the first and third part of the experiment, your task was to guess the color of the ball drawn for the other participant. Your ball was drawn and immediately [returned / NOT returned] before the other participant received a ball. Now you must indicate the probability that the ball of the other participant has a certain color and the probability that the urn has a certain color.

The initial position is: 4 [5] dark blue urns and 4 [3] light blue urns. If you saw a dark blue ball under the above initial position, what is the probability (in %) that the ball of the other participant was also dark blue? [...] If you saw a light blue ball under the above initial position, what is the probability (in %) that the ball of the other participant was also light blue? [...] If you saw a dark blue ball under the above initial position, what is the probability (in %) that the selected urn was dark blue? [...] If you saw a light blue ball under the above initial position, what is the probability (in %) that the selected urn was light blue? [...]

B.2.3 Intervention 3 of Treatment 2

Thank you, you have almost finished part 5 of the experiment. A few questions about your choices in the past parts of the experiment will now follow.

In the 1st, 3rd and 5th part, your task was to guess the color of the ball drawn for the other participant. The ball was drawn and returned immediately before the other participant had received a ball.

Then in the fourth part, you had to specify the probability that the ball had a certain color and the probability that the urn had a certain color. Now we ask you to answer two questions about your decisions. The initial position was: 4 dark blue urns and 4 light blue urns.

In Part 4, you answered the following question: If you saw a dark blue ball given the above initial position, what was the probability (in %) that the other player's ball drawn after replacing your ball was also dark blue?

Your answer was (in %): [...] Please describe what you considered when answering this question: (please press “Enter” after entering the text) [...]

In Part 4, you also answered the following question: If you saw a dark blue ball under the above initial position, what is the probability (in %) that the drawn urn was dark blue? Your answer was (in %): [...] Please describe what you considered when answering this question: (please press “Enter” after entering the text) [...]

B.2.4 Intervention 1 of Treatment 3

Part 2 of the experiment

In Part 2, your task is to answer questions. The questions relate to different scenarios. The scenarios will be similar to those in the first part. In the first part of the experiment, your task was to guess the color of the ball drawn for the other participant.

Your task is to indicate the probability with which a blue or white ball is drawn in each scenario, or with what probability the drawn urn is blue or white in each scenario. Your payout will then depend on how well you choose your estimate.

The entering of your probability estimate is done in two steps: (1) First enter your estimate in the input field. (2) After that, another menu will be displayed on the right to reflect your input. Here you can also see what your payout will be if the other participant has indeed drawn a ball of the same color or a different color or if the urn has indeed the same color or a different color.

You can change your estimate in two ways: Firstly, you can enter a new estimate in the input field. Alternatively, you can directly click on a line in the menu on the right and update your estimate by clicking on “Customize”. When you are satisfied with your decision, press the red “Confirm” button to proceed to the next task. The button will appear on the right as soon as you have entered your first estimate. The initial position is: 4 [5] blue urns and 4 [3] white urns.

Instructions

In the upper part of the screen, you can see the corresponding initial position. Your task: If you have seen a blue [white] ball, indicate how likely you think it is that the other participant’s ball is also blue [white], given the initial position shown above. Your payout will then depend on how well you choose your estimate.

Entering your probability estimate is done in two steps: (1) First enter your estimate in the input field. (2) After that, another menu will be displayed on the right to reflect your input. Here you can also see what your payout will be if the other participant has indeed drawn a same-colored or different-colored ball.

At this point, you can still change your estimate in two ways: Firstly, you can enter a new estimate in the input box. Alternatively, you can directly click on a line in the menu on the right and update your estimate by clicking on “Customize”. When you are satisfied with your decision, press the red “Confirm” button to proceed to the next task. This will appear on the right as soon as you have entered your first estimate.

If you saw a blue ball given the initial position shown above, with what probability (in %) was the other player’s ball blue as well? [If you saw a blue ball given the initial position shown above, what is the probability (in %) that the urn was also blue?] Give your estimate as an integer between 0 and 100.

B.2.5 Intervention 2 of Treatment 3

In the 1st and 3rd part, your task was to guess the color of the ball drawn for the other participant. The ball was drawn and returned immediately before the other player got a ball. In the 2nd part, you had to specify the probability that this ball had a certain color and with what probability the drawn urn had a certain color. Now we ask you to answer two questions about your decisions. The initial position: 4 blue urns and 4 white urns.

In the 2nd part, you answered the following question: If you saw a blue ball under the initial position shown above, what is probability (in %) that the other player's ball was also blue? Your answer was (in %): [...] Please describe what you considered when answering this question: (after entering your answer, please press "Enter") [...]

In the 2nd part, you also answered the following question: If you saw a blue ball under the initial position shown above, what is the probability (in %) that the drawn urn was blue? Your answer was (in %): [...] Please describe what you considered when answering this question: (after entering your answer, please press "Enter") [...]

B.2.6 Intervention 3 of Treatment 3

Part 5 of the experiment

In each round of the 5th part, a ball is randomly drawn four times from the drawn-out urn and returned back again. You will see the color of the randomly drawn balls on the screen. Each ball is drawn, displayed and returned immediately. Another participant will then also see four randomly drawn balls from the same drawn urn.

You then have to guess the color of the first drawn ball for the other participant. Before making your decision, you will be informed about the (for both identical) initial position and the colors of your drawn balls.