

Belief Formation in a Signaling Game without Common Prior: An Experiment*

Alex Possajennikov[†]
University of Nottingham

Final version accepted for *Theory and Decision*, June 2017
Published version available at <http://dx.doi.org/10.1007/s11238-017-9614-z>

Abstract

Using belief elicitation, the paper investigates the process of belief formation and evolution in a signaling game in which a common prior is not induced. Both prior and posterior beliefs of Receivers about Senders' types are elicited, as well as beliefs of Senders about Receivers' strategies. In the experiment, subjects often start with diffuse uniform beliefs and update them in view of observations. However, the speed of updating is influenced by the strength of initial beliefs. An interesting result is that beliefs about the prior distribution of types are updated slower than posterior beliefs, which incorporate Senders' strategies. In the medium run, for some specifications of game parameters, this leads to outcomes being significantly different from the outcomes of the game in which a common prior is induced. It is also shown that elicitation of beliefs does not considerably change the pattern of play in this game.

Keywords: beliefs, signaling, experiment, learning, belief elicitation

JEL Codes: C91, D83, C72

1 Introduction

When making a decision in a situation involving uncertainty, individuals may form beliefs about the probabilities of various outcomes of uncertain events. Within game theory, situations involving uncertainty (about elements other than the actions of other players) are

*Acknowledgments: I would like to thank the School of Economics, University of Nottingham, for financial support and CeDEx for providing access to the infrastructure to run the experiment. At different stages of the project, the paper benefitted from presentations at various conferences, including Foundations of Utility and Risk (FUR) 2016 conference. I thank the editor of this special issue, Ganna Pogrebna, for providing an opportunity for papers presented at the 2016 FUR conference to be considered for publication. I am grateful to an anonymous referee for comments that led to improvements in the paper and to Maria Montero for suggestions to make the exposition better.

[†]Centre for Decision Research and Experimental Economics (CeDEx), School of Economics, University of Nottingham, University Park, Nottingham NG7 2RD, United Kingdom. Tel: +44 115 9515461, email: alex.possajennikov@nottingham.ac.uk

represented by games with incomplete information. The Harsanyi (1967) approach to such games postulates that players' beliefs about the events describing their information are derived from a commonly known probability distribution. This approach proved fruitful by reducing the complexity of the situation. However, in many realistic situations the players do not necessarily know, or do not have a common belief about, the probability distribution governing the uncertainty. If this distribution is not known to the players, how do they form (and update) beliefs about it?

In a strategic situation, other players' behavior is also uncertain (at least initially) from the point of view of a player. Effects of such uncertainty (as opposed to risk arising from a known probability distribution) on behavior and beliefs of players in one-shot strategic interactions have been investigated in experimental settings in, e.g., Eichberger et al. (2008), Ivanov (2011), Kelsey and le Roux (2015) and Li et al. (2017). The findings from these works appear mixed: attitudes towards ambiguity of others' behavior vary across players and settings, making predictions about its effects difficult to generalize.

This paper contributes to the literature by exploring the entire process of belief formation and updating in a strategic situation with incomplete information. It reports on an experiment in which individuals play a signaling game. One player, the Sender, has a piece of private information (type) and can send a message to the other player, the Receiver. The Receiver sees the message but not the type of the Sender and takes an action. The payoffs of both players depend on the Sender's type, the message and the action. To take an appropriate action, the Receiver needs to form beliefs about the Sender's type based on the message the Sender sends. The Receiver can get an idea about the appropriate action by inferring something about the Sender's type from the message sent. This inference may not be straightforward and the Receiver's prior beliefs about the distribution of types are important to form beliefs about the type based on the message received.

Prior beliefs about types can be explicitly induced by specifying the probabilities of the possible types of the Sender. Without explicitly induced prior beliefs, if the game is played often enough, players can learn from observations.¹ Drouvelis, Müller, and Possajennikov (2012) (henceforth DMP) investigated how behavior in the signaling game can be different depending on whether the probabilities of the Sender's types are known or not known before a series of interactions starts. The reason for the possible difference is that without explicitly induced common prior beliefs, players can use different beliefs and thus initially employ different strategies. Path dependence can then lead to different medium to long run outcomes, even if learning from observations allows to approximate the actual probabilities.²

¹Decisions from experience, and their differences from decisions from description, have been reviewed and investigated in psychology literature by, among others, Hertwig and Erev (2009), Gonzalez and Dutt (2011) and Ludvig and Spetch (2011). Often an additional complication in those studies is that sampling size is endogenous; in contrast, in our case the sampling size is fixed and known.

²Chen et al. (2007) also investigated a setting with incomplete information, an auction, with and without induced knowledge of the prior distribution of players' valuations. They found differences in behavior between these two cases in initial rounds but not in later rounds.

In this paper, it is further investigated how beliefs are initially formed and updated in such situations. This is important because a model of behavior in a game with uncertainty cannot be complete without specifying beliefs and their updating. Indeed, predictions about behavior in DMP were derived based on a belief updating process (with starting point based on level-1 behavior in the level-k theory, originated in Stahl and Wilson (1994, 1995), and first applied to signaling games, albeit only for beliefs about strategies, in Brandts and Holt, 1996). However, the question of whether beliefs are really updated in the way the model suggests could not be answered without observing them more directly.

In the experiment reported in this paper, subjects made choices in a signaling game, as well as reported their beliefs at regular intervals. Belief elicitation was incentivized with the quadratic scoring rule. This rule has been used for belief elicitation in game experiments with a small number of actions by, among others, Nyarko and Schotter (2002), Costa-Gomes and Weizsäcker (2008) and Hyndman et al. (2012).³ Rutström and Wilcox (2009), Blanco et al. (2010) and Armantier and Treich (2013) discuss the methodological issues of the possible interaction between belief elicitation and actual play. Whether belief elicitation affected play is tested in this paper (to a large extent, it does not appear so).

In the experiment, Receivers reported beliefs about the prior probability of the type of the Sender in their current (random) match and, after receiving a message, about the posterior probability of the type. Senders, after sending a message, reported beliefs about the matched Receiver’s action in response to this message. The Sender’s type is determined exogenously by a random device, thus its prior probability represents “objective” uncertainty. Which messages are sent by which types (and thus the posterior probabilities of the types), and which actions are taken in response to messages, on the other hand, is determined endogenously within the game. This strategic uncertainty is “subjective” and may depend on the models the players use to determine the behavior of the opponent. Drawing from psychological research, Nickerson (2004, Ch. 8) argues that beliefs about “objective” uncertainty take more time to be revised than beliefs about an individual’s performance. Since in the experiment both types of beliefs are observed, it is possible to see whether some beliefs are updated faster than others in an interactive setting.

Without more explicit information about the resolution of uncertainty, the “principle of insufficient reason” (e.g., Sinn, 1980, and references therein) states that if there is no reason to believe that one event is more likely than another, then they should be assigned equal probability. In the signaling game context, the principle is more applicable to beliefs about types. Beliefs about strategies can also be subject to this principle; in the level-k theory this is the starting point, representing level-1 belief that the behavior of the opponent is level-0 (uniform distribution). However, further levels of reasoning can also be used to determine which strategy is more likely to be played by the opponent, even

³Different procedures for eliciting beliefs are compared and reviewed in, e.g., Palfrey and Wang (2009), Trautmann and van de Kuilen (2015) and Hollard et al. (2016). The quadratic scoring rule appears to be a reasonable compromise between the reliability of a rule and its complexity.

Type t_1	$\Pr(t_1) = p$	Receiver		Type t_2	$\Pr(t_2) = 1 - p$	Receiver	
Sender		a_1	a_2	Sender		a_1	a_2
	m_1	15, 10	80, 80		m_1	80, 80	15, 30
	m_2	25, 10	50, 50		m_2	50, 50	25, 30

Figure 1: The Signaling Game

without experience. Comparing initial beliefs about prior and posterior distributions of Senders' types, or about Receivers' strategies, one can tell whether there is a difference in the formation of beliefs about different types of uncertainty.

Thus, the main research questions of this paper are how beliefs are formed (whether initial beliefs, both about types and about strategies, are close to being uniform), how beliefs are updated, and whether some beliefs are updated faster than others. The data suggest that beliefs about Senders' types, both prior and posterior, indeed start close to being uniform; even beliefs about Receivers' strategies are not far from the uniform distribution. As observations accumulate, beliefs are updated in the natural direction of the frequency of events. However, updating is not as fast as simple frequency count would suggest, indicating that initial beliefs may have a sizeable weight in the updating process. Beliefs about the posterior distribution of types appear to be updated faster than about their prior distribution; Receivers' learning of Senders' strategies thus has an effect on how fast beliefs are updated. Senders' beliefs about Receivers' strategies also appear to be updated faster than Receivers' prior belief about Senders' types.

Given these properties of belief updating, the observed play in the game exhibits differences between the situations with known probabilities of Senders' types and unknown ones, due to path dependence in one of the treatments. This happens because starting from the uniform initial beliefs the play is taken to a different equilibrium than starting from known correct probabilities of Senders' types, if initial beliefs about types are not updated too fast. In the other treatments, in the long run there is no noticeable difference in behavior between the cases of known probabilities of types and of unknown ones. Therefore, the uncovered process of belief formation and updating has sometimes important consequences for long-run outcomes.

2 The Signaling Game and Belief Elicitation

Individuals were asked to play the signaling game given by the payoff tables in Figure 1 (see also Drouvelis, Müller, and Possajennikov, 2012, DMP). The first number in a cell in the tables is the payoff of the Sender (Player 1) and the second number is the payoff of the Receiver (Player 2). In the game, the type of the Sender is determined by a random draw, with the probability of Type t_1 being p and that of Type t_2 being $1 - p$. The Sender, knowing

the type, chooses one of two messages, m_1 or m_2 . The Receiver observes the message sent by the Sender but not the Sender's type and takes one of two actions, a_1 or a_2 . Payoffs depend on the Sender's type and message and on the Receiver's action.

Three values of p are considered in the experiment, $p = 1/4$, $p = 1/2$, and $p = 3/4$. For each of the values of p , the game has two separating equilibria $[(m_1, m_2), (a_2, a_1)]$ and $[(m_2, m_1), (a_1, a_2)]$, where the first element is the message of the Sender if the Sender is type t_1 , the second one is the message if the Sender is type t_2 , the third element is the action of the Receiver after receiving message m_1 , and the last element is the action after receiving message m_2 .⁴

Apart from different values of p , the other treatment difference in the experiment is that in some treatments this value is commonly known to the players, while in the other treatments the value is not revealed to them. In this way it can be investigated how the information about the probability distribution of the Sender's types affects adjustment towards equilibrium.

The payoffs in the game were chosen so that a naive adjustment process, discussed in Brandts and Holt (1996) and extended in DMP to situations without a commonly known prior distribution, converges to equilibrium $[(m_2, m_1), (a_1, a_2)]$ in the treatment with known value $p = 1/4$, while in the other treatments the process converges to equilibrium $[(m_1, m_2), (a_2, a_1)]$. This happens if beliefs about types are not updated or updated only slowly.

The naive process starts with a belief that the strategy of the opponent is uniform (belief of a level-1 player in the level-k theory). With such a belief, both types of the Sender prefer to play m_1 .⁵ If $p = 1/4$ (and known), the best response of the Receiver to the uniform strategy of the Sender is a_1 against both messages. If such a play is observed, Type 1 Sender then switches to m_2 and in response the Receiver switches to a_2 against m_2 . Equilibrium $[(m_2, m_1), (a_1, a_2)]$ is reached. If $p = 1/2$ or $p = 3/4$, the best response of the Receiver against the uniform belief about the strategy of the Sender is a_2 against both messages. Now it is Type 2 Sender that would want to switch to m_2 , and then the Receiver switches to a_1 in response to m_2 , reaching equilibrium $[(m_1, m_2), (a_2, a_1)]$.⁶

If p is unknown, naive beliefs are that each type is equally likely. In this case the process will start like the process described above with $p = 1/2$. If this belief about the value of p is not updated, or updated only slowly, the adjustment path to the equilibrium $[(m_1, m_2), (a_2, a_1)]$ can be followed, as if $p = 1/2$ is known.

DMP show that there are no statistically detected differences in the observed play between treatments in which the value of p is known or not for $p = 1/2$ or $p = 3/4$. For

⁴There is also an equilibrium in partially mixed strategies, for each value of p . However, these equilibria are unstable under many specifications of adjustment dynamics and indeed are not observed in the data.

⁵If risk- and ambiguity-neutral; these assumptions are maintained throughout the discussion, also for the Receiver.

⁶The presence of players with a higher level of reasoning than level-1 would accelerate the process but not change the final outcome.

$p = 1/4$, there are differences in play depending on whether this value of p is known or not, although these differences are not as clean as predicted by the naive adjustment theory above. One possible explanation is that the overall direction of adjustment depends on how fast different beliefs are updated. If the adjustment of beliefs about types is much slower than that of beliefs about strategies, the path in the previous paragraph is followed. On the other hand, if type beliefs are revised faster, Receivers may realize sooner that Type 1 is less likely than Type 2 and follow the adjustment path for $p = 1/4$.

In DMP, beliefs were not elicited although it was shown that behavior in the initial periods of the treatments without a commonly known value of p was not statistically different from behavior in the treatment with known value $p = 1/2$. While this provides indirect evidence for the naive theory of belief formation, to understand belief initialization and adjustment better, it is important to observe beliefs directly, as noted, for example, in Nyarko and Schotter (2002).

To perform this direct check on the formation and adjustment of beliefs, in this paper beliefs are elicited during the course of play. The novel angle is that since the signaling game under consideration involves a genuinely random move (with an unknown distribution), players have to form and update beliefs about uncertain events that are conceptually different. The random move by Nature is an example of objective uncertainty, with a stationary distribution.⁷ By contrast, the uncertainty about the strategies of the opponent is random only from the point of view of the player, and its distribution may be changing as the opponent learns how to play the game. Nickerson (2004, Ch. 8) reports psychological evidence about different speed of belief formation depending on whether uncertainty is objective or about a person's performance. Nevertheless, the evidence is not about behavior in a strategic situation and the analysis presented in this paper is a further step towards understanding how players in a game deal with different kinds of uncertainty.

In the experiment belief elicitation is incentivized via the quadratic scoring rule, as, e.g., in Nyarko and Schotter (2002), Costa-Gomes and Weizsäcker (2008) and Hyndman et al. (2012). While this works only for risk-neutral players, payoffs are sufficiently low and there are many periods so that risk neutrality is not an implausible assumption.⁸ In contrast to other papers that used belief elicitation, beliefs are elicited not every period but every few periods. This is done to concentrate subjects' efforts on this task rather than making it routine. It also allows the subjects to gain more observations to base their guess on. Although it reduces the number of observations, the likely extra effort for the task and the better base for the guess may be sufficient to hope that the reported beliefs are a good representation of the ones subjects actually hold.

⁷The stationarity of the distribution was emphasized in the experiment instructions.

⁸Offerman et al. (2009) and Andersen et al. (2012), among others, propose corrections of belief reports to account for different attitudes to risk and ambiguity. These corrections, however, require participants to perform additional tasks. To avoid the additional complexity, and since the focus is on the comparison of beliefs, all likely affected by risk and ambiguity attitudes in a similar manner, no such corrections are applied in this paper.

3 Experiment and Belief Elicitation Design

The design of the experiment in DMP is followed, with the addition of belief elicitation. The signaling game is described in the previous section. Subjects were assigned the role of either the Sender or the Receiver, and made the corresponding decisions. In addition to these decisions, in some periods the subjects were asked to report their beliefs about the matched Sender’s type or about the matched Receiver’s action.

Belief elicitation was based on the following procedure. Suppose that a player has beliefs about a binary random variable X . The beliefs are that $X = 1$ with probability q and $X = 0$ with probability $1 - q$. A player is asked to report q . After a value of the random variable X is realized, the quadratic scoring procedure gives payoff

$$\pi = A \cdot \left(1 - \frac{1}{2} ([q - I(X = 1)]^2 + [1 - q - I(X = 0)]^2) \right), \quad (1)$$

where $I(\cdot)$ is the indicator function that takes value 1 if its argument is true and 0 otherwise. Given this payoff, and assuming risk neutrality, it is optimal to report the true belief q (see, e.g., Trautmann and van de Kuilen, 2015).

The experiment contains treatments with and without known prior probabilities of Senders’ types. In treatments in which the probabilities are not known, Receivers are asked about their beliefs about the matched Sender’s type before a message is received (prior beliefs) and after they have received a message (posterior beliefs). In treatments in which the value of p is known, Receivers are asked only about their posterior beliefs. In all treatments, Senders are asked about the probability of the matched Receiver’s action after they have sent the message.

In treatments in which the value of p is unknown, prior beliefs about the Sender’s type represent beliefs about an event that is independent of the opponent’s actions. On the other hand, posterior beliefs of Receivers and beliefs of Senders about the Receiver’s action concern events that are affected by the actions of the opponent. Formation and adjustment of beliefs may be different depending on the distinction between “objective” events and events influenced by the opponent.

In the experiment, beliefs were elicited according to rule (1) with $A = 50$. An experimental session lasted 36 periods, with subjects randomly matched every period. Beliefs were elicited in Period 1 (initial beliefs), and then every 5 periods (i.e., in periods 1, 6, 11, 16, 21, 26, 31, 36), about the events described in the previous paragraphs.⁹

The decisions not to elicit beliefs every period and to set $A = 50$ were made for several reasons. Subjects observed feedback about types and decisions in their own match only, thus having one observation about the realizations of relevant random variables each period. Having more observations between the periods of belief elicitation gives the subjects a better basis to form their view of relevant probabilistic events. Additionally, being asked to report

⁹See a sample of the experiment instructions in Section A of Supplementary Materials for more details.

their beliefs only at certain periods makes this task less routine and thus hopefully induces sufficient effort to think about it. To get enough incentives to think about beliefs, the payoffs for getting them right are comparable with those from playing the game. The subjects could get a maximum of 50 points from correctly predicting the type or the action of the other player, while in the game 50 was the second-highest payoff.¹⁰

The treatment differences are the value of p ($p = 1/4, 2/4, 3/4$), and whether this value is known or not (K or N). In the sequel a treatment is denoted Xy , with $X = K$ if p is known and $X = N$ if not, and $y = 1$ if $p = 1/4$, $y = 2$ if $p = 2/4$, and $y = 3$ if $p = 3/4$. Thus, the treatment without commonly known value of p and $p = 1/4$ is denoted $N1$ and similarly for the other treatments.

The length of the sessions was 36 periods, to allow enough opportunities for learning, while at the same time not too long for the task to become tedious for the subjects. Each session lasted approximately 90-100 minutes. In each session, the Sender and Receiver roles were assigned randomly at the beginning. Then the subjects were randomly matched to play the game in every period. In some periods they were also asked to report their beliefs as described above. At the end of a session, the subjects were paid according to the total amount of accumulated points, with points converted to pounds at the rate of £0.05 for 10 points.

The new (with respect to DMP) set of experiments was done in the CeDEx laboratory at the School of Economics at the University of Nottingham in February-March 2009 and December 2013. The experiment was programmed with z-tree (Fischbacher, 2007) and subjects were recruited using ORSEE (Greiner, 2015). There were 3 sessions in treatments $N1$ and $K1$, since these treatments are likely to produce the most interesting treatment difference. The number of sessions was chosen to have enough observations for the non-parametric tests below. For each of the other treatments, two sessions were run. In each session 16 subjects participated, divided into two matching groups of 4 Senders and 4 Receivers, thus making two largely independent observations per session.¹¹ The matching protocol and the type assignment was the same as in DMP. Table 1 summarizes the design, along with the predictions about the likely long-run equilibrium, discussed in section 2.

In the best equilibrium of the game ($[(m_2, m_1), (a_1, a_2)]$ if $p = 1/4$ and $[(m_1, m_2), (a_2, a_1)]$ if $p = 3/4$), and with best predictions, a subject could earn £16.28. If all subjects used the uniformly random strategy and made the uniform prediction, they would have earned on average £10.16 per subject. The average earnings were in fact £11.96 (between \$17-20, depending on the exchange rate at the time of the experiments) per subject, higher than with uniform play and predictions, but way off the payoff in the best equilibrium and with

¹⁰Of course, eliciting beliefs more often would have allowed to collect more data; however, this can make the task routine and incentives would have to be smaller. Facing the trade-off between paying less every period or having a higher payment every few periods, the latter option was chosen since it gives the subjects more incentives to take the belief reporting task seriously.

¹¹One session, in treatment $K3$, had only 8 participants.

Treatment	Matching groups	Subjects	Prediction
N1	6	48	$[(m_1, m_2), (a_2, a_1)]$
K1	6	48	$[(m_2, m_1), (a_1, a_2)]$
N2	4	32	$[(m_1, m_2), (a_2, a_1)]$
K2	4	32	$[(m_1, m_2), (a_2, a_1)]$
N3	4	32	$[(m_1, m_2), (a_2, a_1)]$
K3	3	24	$[(m_1, m_2), (a_2, a_1)]$
Total		216	

Table 1: Number of subjects and matching groups for each treatment

the best possible predictions.

The main aim of the experiment was to explore the way beliefs are formed and updated. Since beliefs are elicited directly, one can formulate two hypotheses concerning beliefs, one for their initialization and the other for updating.

Hypothesis 1 *Initial beliefs are uniform.*

The hypothesis consists of several parts, depending on the event about which beliefs are elicited. In all treatments, Senders are asked about Receivers' strategies. Thus one part is that the beliefs of Senders are uniform. Receivers are asked about the posterior beliefs about Senders' types, as well as, in treatments with unknown value of p , about the prior probability of the types. While the prior belief is a distribution for a simple binary event, posterior beliefs reflect beliefs about Senders' strategies. Thus there are further two parts of the hypothesis: beliefs about the prior distribution are uniform, and beliefs about Senders' strategies are uniform.

The hypothesis is based on the principle of insufficient reason.¹² If it is rejected, then apparently subjects initialized their beliefs differently discerning some reasons for doing so. The hypothesis is more likely to hold for beliefs about the prior distribution of types, since strategic considerations can lead to different beliefs about the actions of Receivers and the strategies of Senders.

Hypothesis 2 *Beliefs are updated with experience. The subjective probability of experienced outcomes increases.*

There are several ways to operationalize the hypothesis, since there are many ways to update beliefs in the direction of experienced outcomes. The details of hypothesis operationalization are left for the next section.

The third hypothesis is a composite hypothesis controlling for the possible differences in behavior depending on whether beliefs are elicited or not.

¹²Sinn (1980) provides a relatively recent theoretical analysis of it and Binmore et al. (2012) give recent experimental evidence in its favor.

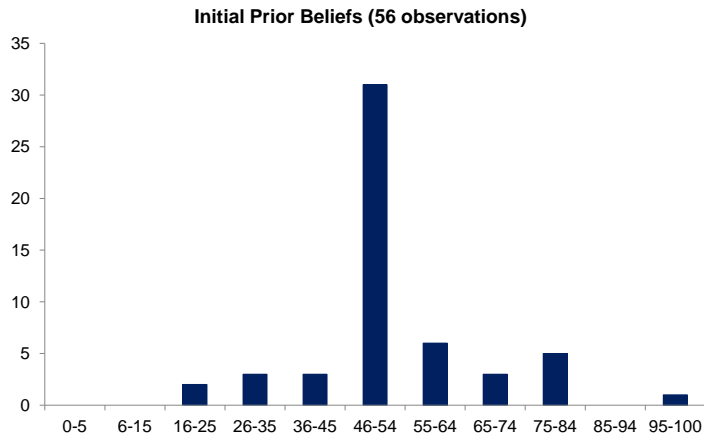


Figure 2: Histogram of initial beliefs about prior in the N treatments

Hypothesis 3 *Behavior with belief elicitation is not different from behavior without belief elicitation.*

The hypothesis compares the data from the new experiment with the data on the same game but without belief elicitation in DMP. There, it was found that there are differences in behavior between treatments $N1$ and $K1$, and there are no differences between treatments with known and unknown prior for other values of p . The hypothesis checks whether the presence of belief elicitation makes the subjects behave more or less strategically than in the absence of belief elicitation.

4 Experiment Results

4.1 Initial beliefs

To begin, initial beliefs, elicited in Period 1, are analyzed. Recall that in the treatments in which the actual probabilities of Senders' types is not revealed to the subjects (N treatments), Receivers are asked to report their belief about the matched Sender's type before seeing any message or taking any action (prior beliefs). Additionally, in all treatments, after Senders have chosen a message, they are asked to predict what action the matched Receiver will take in response to this message (action beliefs). Receivers, meanwhile, see the matched Sender message and are asked to choose an action and, at the same time, make a statement about the Sender's type (posterior beliefs).

For prior beliefs, absent any other information, the most natural guess, based on the principle of insufficient reason, is that each of the two types is equally likely. Figure 2 presents the histogram of 56 observations of reported initial prior beliefs about the Sender's type in the N treatments.

Most ($31/56 = 57\%$) of the reported beliefs lie within the interval 46-55%, i.e. close to

	<i>N</i> treatments		Treatment <i>K1</i>	
	$t_1 m_1$ (41 obs)	$t_1 m_2$ (15 obs)	$t_1 m_1$ (20 obs)	$t_1 m_2$ (4 obs)
Mean	0.56	0.48	0.31	0.38
St.Dev.	(0.25)	(0.24)	(0.24)	(0.25)

Table 2: Initial posterior beliefs about the types of Senders

0.5 probability of Type 1 (in fact, all 31 observations in this interval are exactly at 50%). The average reported prior belief about the type is 0.54 (standard deviation 0.15). The one-sample t-test and the Wilcoxon signed-rank test do not reject the hypothesis that the median reported belief is equal to 0.5 at 5% significance level.¹³ Thus the prior belief about Senders' types appears centered on 0.5 and, according to the histogram, is concentrated on this value.

Senders may be expected to behave differently in treatments *K* with different values of p but not so much in treatments *N*, where Receivers do not know (yet) which type of Senders is more likely to send which message. Thus, posterior beliefs of Receivers may differ across the *K* treatments but not across the *N* treatments, which is indeed the case for the reported beliefs. The reported posterior beliefs of Receivers about the types of Senders in Period 1 can then be pooled in the *N* treatments. Table 2 shows the mean and the standard deviation of these pooled beliefs and compare them with posterior beliefs in treatment *K1*.¹⁴

The initial posterior beliefs about types in the *N* treatments are also not far from 0.5, although the standard deviation is higher than for the initial prior beliefs. The Wilcoxon-Mann-Whitney rank-sum test does not find a significant difference between the posterior beliefs for the two different messages, and the signed-rank test for paired observations does not detect a significant difference between the reported prior and posterior beliefs about types.

In treatment *K1* there is also no significant difference between the reported posterior type beliefs across the two messages. Recall that in the *K1* treatment the common prior $p = 0.25$ is induced. Although the average reported posterior beliefs are higher, they are not significantly different from 0.25 by the signed-rank test.

Thus, there is little evidence that the average initial posterior beliefs of Receivers take into account the possible separation of Sender's types by messages. The reported beliefs are consistent with Senders pooling, including with the possibility of both types of Senders choosing one of the two messages uniformly randomly.

In the *N* treatments, Senders know their own realized type but they do not know the prior probability of each type and they cannot expect Receivers to react differently

¹³Detailed test results for these and other tests in this section are given in Section B.1 of Supplementary Materials.

¹⁴For posterior beliefs in *K2* and *K3*, observations similar to those about posterior beliefs in *K1* can be made.

	N treatments		Treatment K1	
	$a_1 m_1$ (41 obs)	$a_1 m_2$ (15 obs)	$a_1 m_1$ (20 obs)	$a_1 m_2$ (4 obs)
Mean	0.48	0.43	0.65	0.55
St.Dev.	(0.31)	(0.21)	(0.22)	(0.10)
Actual play	0.17	0.20	0.75	0.50

Table 3: Initial beliefs of Senders about the actions of Receivers

depending on the value of p , thus the N treatments are pooled together. Among the K treatments, where Senders can expect a different reaction of Receivers, treatment $K1$ is again chosen for illustration. The reported beliefs of Senders about the actions of Receivers in Period 1 are shown in Table 3.

In treatments without an induced common prior, the average beliefs of Senders are quite close to 0.5, although they are heterogeneous as the standard deviation is high. Non-parametric tests do not find a significant difference in the median of these beliefs by message, or from the uniform belief $\Pr(a_1|m) = \Pr(a_2|m) = 0.5$. Note that these beliefs are not very accurate: the last row shows the proportions of actions actually played by Receivers and they are much lower than the beliefs reported by Senders.

In treatment $K1$, Senders report beliefs that action a_1 is going to be taken more often than action a_2 by Receivers. These beliefs are sensible because, knowing that Type 2 is more likely, Receivers indeed get a higher payoff by choosing a_1 . These beliefs also reflect to some extent the actual proportion of choices of action a_1 . It appears that Senders did make some adjustment for strategic sophistication of Receivers already in Period 1 if the common prior probability of Type 1 $p = 1/4$ was induced. With an unknown prior though, Senders' beliefs are close to a 50-50 chance of Receivers taking either action.

Result 1 *For initial beliefs:*

- (i) *Initial beliefs of Receivers about the prior probability of Senders' types are close to uniform in treatments with an unknown value of p ;*
- (ii) *Initial posterior beliefs of Receivers about Senders' types are not different from the initial prior beliefs about the types;*
- (iii) *Initial beliefs of Senders about the actions of Receivers are close to uniform in the treatments with an unknown value of p but put more weight on a_1 in treatment $K1$.*

The results are thus consistent with the naive approach story that many subjects hold level-1 beliefs. Receivers mostly follow the “principle of insufficient reason” in forming their prior belief about Senders' types. They do not see Senders separating their message by type all in the same way, thus their average posterior beliefs are consistent with all Senders choosing the same message, or all Senders choosing messages uniformly randomly.

Senders' average beliefs are also consistent with Receivers choosing either of the two actions with the same probability in the N treatments; in the K treatments Senders appear more sophisticated though, partially predicting a higher likelihood of certain actions by Receivers.

Beliefs are consistent with optimizing behavior if a subject's chosen message or action is a best response to the reported belief. In the experiment Receivers play best response to the reported posterior beliefs in Period 1 73% of the time, higher than a random choice.¹⁵ For Senders it is not possible to determine whether their choice of message is indeed a best response because they are not asked for their beliefs about the action of the Receiver in response to the non-chosen message. One possibility is to consider whether no beliefs about actions after the non-chosen message would make the message played consistent with best response.¹⁶ Since one can often find beliefs making the choice of message consistent with best response, only 4% of messages and reported beliefs of Senders in Period 1 are clearly inconsistent with best response. Alternatively, one can assume that in Period 1 Senders have the same beliefs about Receivers' actions after both messages. If this assumption is adopted, 70% of Senders' chosen messages and reported beliefs in Period 1 are consistent with best response.

4.2 Belief adjustment

The previous section analyzed beliefs in Period 1. Subjects in the experiment were also asked to report beliefs in a number of subsequent periods, and this section looks at these reports.

4.2.1 Beliefs about the prior probability of Senders' types

Figure 3 illustrates the evolution of the average belief of Receivers about the prior (i.e., before seeing the message sent to them) probability of Senders' types for the three N treatments (recall that beliefs were elicited in rounds 1, 6, 11, 16, 21, 26, 31, 36). The average reported beliefs about the prior are represented by solid lines; the dotted lines are explained below. Starting from beliefs about the probability of type t_1 close to 0.5 for all three treatments, reported beliefs generally move in the right direction (downwards for $p = 1/4$ and upwards for $p = 3/4$, although movements for $p = 1/2$ and $p = 3/4$ are more erratic because they are based on fewer observations (16 subjects in each of $N2$ and $N3$) than for $p = 1/4$ (24 subjects)). Non-parametric tests for $N1$ and $N3$ treatments

¹⁵This is also evidence that hedging is not a large problem, since it would imply choosing the other action, not the best response.

¹⁶For the game played, the worst-case scenario for the Sender after m_1 is to get 15 and after m_2 is to get 25. Thus, m_2 is always consistent with best response since there always exists a belief that m_1 leads to a lower payoff. After m_1 the expected payoff is $15 \cdot \Pr(a_1) + 80 \cdot (1 - \Pr(a_1))$ for type t_1 and $80 \cdot \Pr(a_1) + 15 \cdot (1 - \Pr(a_1))$ for type t_2 . The chosen message m_1 would be inconsistent with best response if the reported beliefs $\Pr(a_1)$ in response to it were more than $11/13$ for type t_1 and less than $2/13$ for type t_2 .

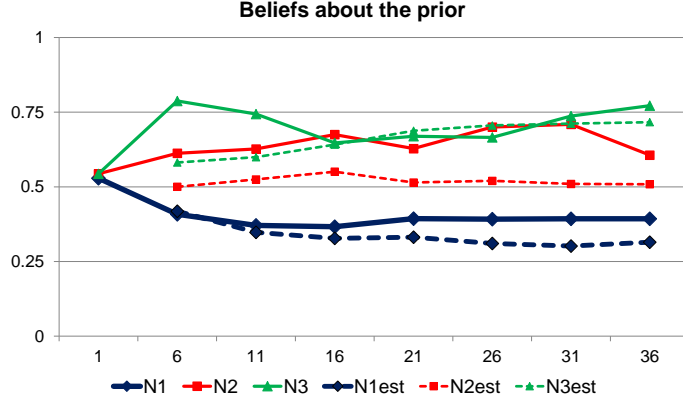


Figure 3: Evolution of reported and predicted beliefs about the prior probability of type t_1

confirm that beliefs in the last period are different from those in the first period.¹⁷ Thus it appears that beliefs about the prior probability are adjusted in the direction of experienced outcomes, at least on average.

To analyze the process of belief adjustment further, several models of belief evolution based on observations are compared. These models of empirical beliefs are

- *Baseline*. Beliefs are equal to the proportion of the times the Sender was type t_1 in a given Receiver’s set of observations. Let A_1^τ denote the count of observations of type t_1 and A_2^τ the count for type t_2 up to period τ . If type t_i is observed in period τ , then $A_i^{\tau+1} = A_i^\tau + 1$, $A_j^{\tau+1} = A_j^\tau$ for $j \neq i$. Beliefs are $q_1^\tau = A_1^\tau / (A_1^\tau + A_2^\tau)$. The initial counts are $A_1^0 = A_2^0 = 0$.
- *Forgetting* (Cheung and Friedman, 1997). This process behaves like the baseline process except that the counts are discounted: $A_i^{\tau+1} = \gamma A_i^\tau + 1$, $A_j^{\tau+1} = \gamma A_j^\tau$ for $j \neq i$. If $\gamma < 1$, then observations further back in the past have less weight in the total count, i.e., they are getting “forgotten”.
- *Initial strength* (Brandts and Holt, 1996). This process is like the baseline process except that the initial counts are not 0 but $A_1^0 = A_2^0 = A$, where A is estimated from the data. Larger values of A would mean that new observations have less weight compared with initial beliefs, i.e., beliefs are updated slower.
- *Forgetting and initial strength*. The process combines both the forgetting parameter γ and the initial beliefs strength A .

The forgetting parameter γ and the initial beliefs strength A are estimated from the comparison of the beliefs predicted by the model with the reported beliefs by minimizing the

¹⁷For $N2$, there is no statistically significant difference between beliefs in Periods 1 and 36. The test results for this and subsequent subsections are reported in Section B.2 of Supplementary Materials.

(392 obs)	One previous period	Empirical				50-50
		Base	Forgetting $\gamma = 0.97$	Init. strength $A_{Pr} = 4.93$	$\gamma = 0.99$ $A_{Pr} = 5.30$	
SSE	111.47	29.65	29.55	26.05	26.04	34.44

Table 4: SSE scores for models of adjustment of prior beliefs about types

sum of squared errors (SSE) between the prediction of the model and the reported beliefs, using data pooled from all players and restricting initial beliefs to be uniform.¹⁸ The results of the estimations and the obtained minimized SSE scores are reported in Table 4.

The table contains also the SSE scores for two other benchmark models. One is the one-previous-period model where beliefs are equal to the observation from the previous period (i.e., equal 1 if the Sender was type t_1 in the previous period and 0 otherwise). Another model, reported in the last column, is the one that predicts probability 0.5 all the time.

It can be seen from the table that the baseline model and the forgetting model do not improve much on the 50-50 prediction. However, models with an initial strength of beliefs do better, and the one with forgetting is not very different from the one without forgetting. It appears that the best model is the one that allows for forgetting and with the strength on initial beliefs $A_{Pr} = 5.3$. Since each new observation has weight 1, the value 5.3 indicates how slowly beliefs about the “objective” probability of Senders’ types change.

The dotted lines in Figure 3 correspond to average beliefs of this best-fit estimated model. The model tracks beliefs in treatments N1 and N3 reasonably well, but not so much in treatment N2.

4.2.2 Posterior beliefs about types

The model with a certain strength on initial beliefs seems to fit the data best among the considered models for prior beliefs about types. If this model also explains the evolution of posterior beliefs about types or beliefs about strategies, one can compare the different speeds of belief revision since the parameter A can be seen as a measure of this speed.

For treatments $N1$ and $K1$ (for which there are more observations than for $N2$ and $K2$, or for $N3$ and $K3$), the evolution of the average posterior beliefs of Receivers is illustrated in Figure 4 (solid lines). The figure reflects type separation in treatment $K1$, while the picture is much more mixed in treatment $N1$, as non-parametric tests confirm.¹⁹ The dotted lines again illustrate the average beliefs of the best-fit model found below.

Posterior beliefs appear to start close to 0.5 in treatment $N1$ and somewhere between 0.5 and 0.25 in treatment $K1$ and then move in the direction of experienced outcomes (which

¹⁸Since there are only a few observations per subject, estimating parameters for each subject is infeasible.

¹⁹Indeed, play did not converge clearly to either of the separating equilibria in treatment $N1$; see Figure 6 in Section 4.3 for the evolution of the average proportions of each strategy played.

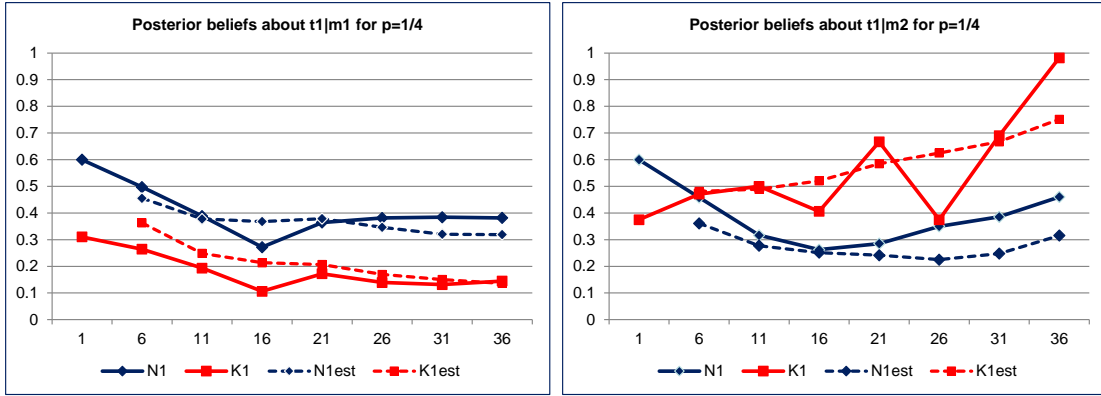


Figure 4: Evolution of reported and predicted posterior beliefs about types in treatments with $p = 1/4$

(756 obs)	Reported beliefs	One previous period	Empirical		
			Base	$\gamma = 0.97$ $A_{P_s} = 2.29$	50-50
SSE		161.47	66.32	55.43	100.42
Best resp.	0.84	0.73	0.82	0.82	0.58

Table 5: SSE scores for models of adjustment of posterior beliefs about types

are reflected in the dotted lines representing an empirically based adjustment model). Non-parametric tests show that there are differences in the reported posterior beliefs in Period 1 and in Period 36 for most of the comparisons (except for beliefs about $t_1|m_2$ in treatment $N1$, for which there are fewer observations). Subjects in the experiment seem to learn something about the posterior type probabilities over time.

To see which adjustment model fits best, the same models as for the prior beliefs about types were considered, with the results presented in Table 5.²⁰ The model with an initial strength of beliefs has the lowest SSE score, while allowing for forgetting does not improve it much.²¹ An interesting observation is that the estimated initial belief strength parameter of this model, $A_{P_s} = 2.29$ is considerably lower than the corresponding parameter for the prior beliefs about types, $A_{P_r} = 5.3$. It appears that the posterior beliefs about types are updated faster than the prior ones, possibly because the posterior beliefs incorporate

²⁰Scores are based on all treatments, not only on those with $p = 1/4$. The SSE score for the forgetting model is very similar to that of the baseline model; the score for the initial strength model without forgetting is similar to that of this model with forgetting (see Section B.2 of Supplementary Materials). Thus only the baseline and the full (initial strength and forgetting) scores are reported.

²¹It is the predictions of the model with $\gamma = 0.97$ and $A_{P_s} = 2.29$ that are also shown in Figure 4 by dotted lines.

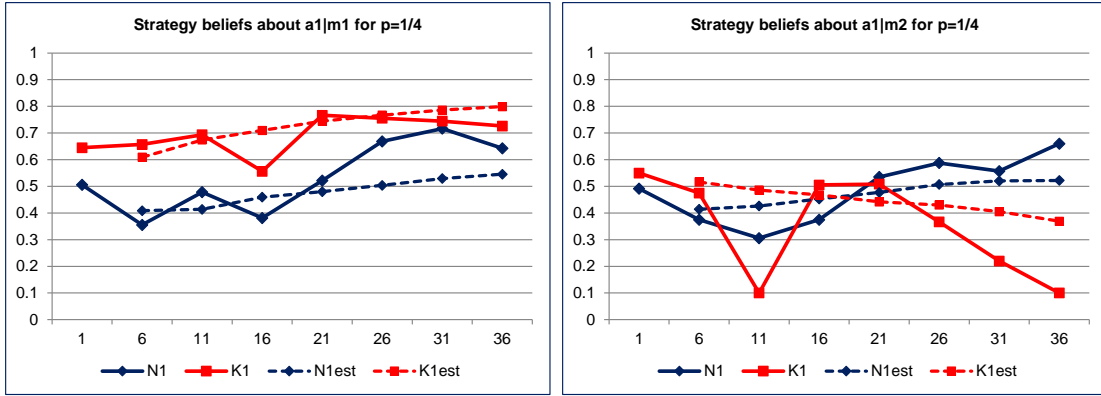


Figure 5: Evolution of reported and predicted beliefs about actions in treatments with $p = 1/4$

beliefs about strategies as well, which are updated faster than beliefs about the objectively uncertain process of type determination.

The table also reports the proportion of choices that were best responses to the reported beliefs (column “Reported beliefs”) or that would be best responses to the beliefs predicted by the models. Receivers chose best response to their reported beliefs 84% of the time, while if their beliefs were following the best adjustment model, their actions would have been best responses 82% of the time. This is close to 84%, thus the adjustment model reflects the reported posterior beliefs to some extent.

4.2.3 Beliefs about strategies

Senders in the experiment reported beliefs about the matched Receiver’s action in response to the message sent. For treatments $N1$ and $K1$, Figure 5 illustrates the evolution of average reported beliefs, together with the predictions of the best adjustment model (dotted lines, labelled $N1est$ and $K1est$). There is again a clearer separation of beliefs about Receivers’ responses for treatment $K1$ than for treatment $N1$, confirmed by statistical tests. Indeed, the observed play in the $K1$ treatment comes close to one of the separating equilibria while in the $N1$ treatment in most of the matching groups there is no convergence (see again Figure 6 in Section 4.3 for the evolution of the average proportions of strategies).

Strategy beliefs start from close to 0.5 in treatment $N1$ but from a higher value in treatment $K1$. Then they move to some extent in the direction of experienced outcomes although this movement is less clear than for the prior or posterior beliefs about types. Indeed, non-parametric tests detect a statistical difference between the reported strategy beliefs in periods 1 and 36 only for beliefs about $a_1|m_2$ in treatment $K1$. Nevertheless, the adjustment models above can be applied to strategy beliefs as well.

The same models as for the prior and posterior beliefs about types were considered,

(756 obs)	One previous period	Empirical		50-50
		Base	$\gamma = 1.00$ $A_{St} = 2.59$	
SSE	164.05	74.75	61.52	93.65

Table 6: SSE scores for models of adjustment of beliefs about strategies

with the results reported in Table 6.²² The lowest SSE score is again achieved by the model with an initial strength of beliefs (allowing for forgetting does not change the SSE score much). The estimated strength parameter of this model, $A_{St} = 2.59$ is again lower than the corresponding parameter for the prior beliefs, $A_{Pr} = 5.3$. Thus, it appears that on average beliefs about strategies are updated faster than beliefs about the prior probability of Senders’ types. However, the SSE score is higher than that for posterior beliefs, indicating that even the best-fit model for beliefs of Senders about Receivers’ actions does not fit the reported beliefs as well as for the posterior beliefs of Receivers.

The observations can be summarized in the following result:

Result 2 *From the beliefs reported in different time periods*

- (i) *On average, beliefs adjust towards observed realizations of the relevant events;*
- (ii) *On average, the model with a weight on initial beliefs explains the reported beliefs better than the other models;*
- (iii) *The weight on initial beliefs is larger for beliefs about the prior probabilities of Senders’ types than for beliefs about the posterior probabilities of the types or about Receivers’ actions.*

The last part of the result resembles the psychological evidence in Nickerson (2004, Ch. 8) that beliefs about a person’s performance are updated faster than beliefs about an “objectively” uncertain process. The prior probability of types is objectively uncertain, while the posterior probability of types and the probability of a given action of the Receiver depend on the behavior of the players. In the strategic situation under consideration, beliefs about probabilities of events that depend on players’ decisions are updated faster, which is represented by a lower weight on initial beliefs about such events.

Overall in the experiment, Receivers played best response to their beliefs 80% of the time. For Senders, it is not possible to determine whether their messages are fully best-reply consistent with their reported beliefs because beliefs about the Receiver’s action after the

²²Scores are again based on all treatments, not only on those with $p = 1/4$. The SSE score for the forgetting model is very similar to that of the baseline model; the score for the initial strength model without forgetting is similar to that of this model with forgetting (see Section B.2 of Supplementary Materials). Thus only the baseline and the full (initial strength and forgetting) scores are reported.

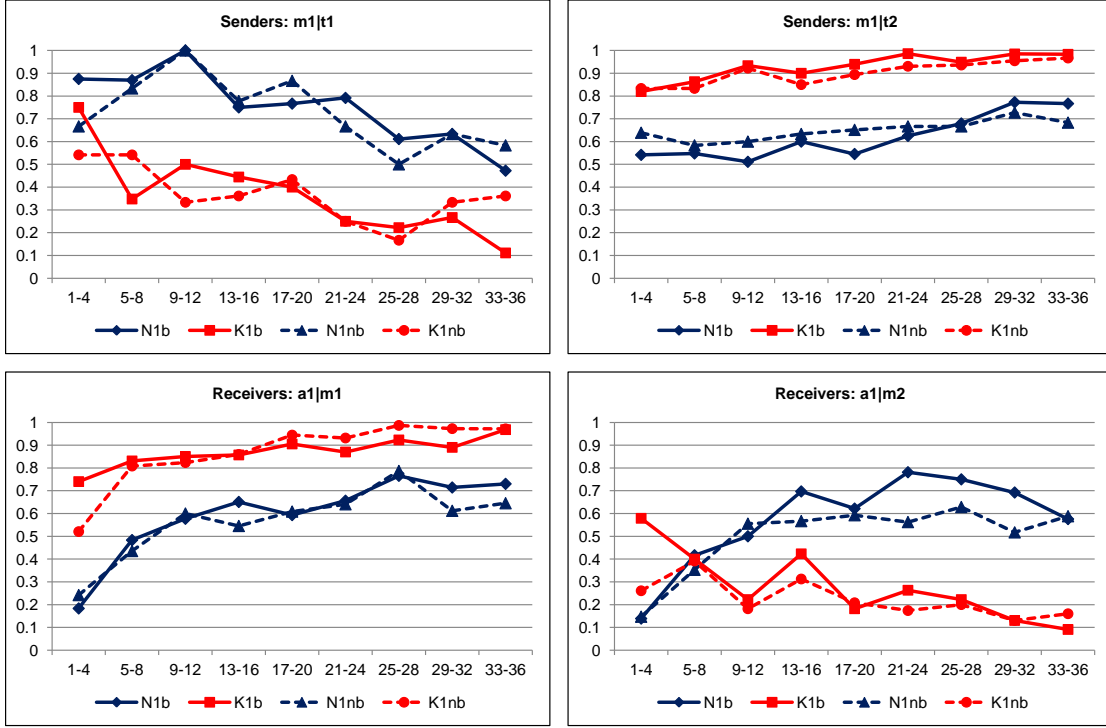


Figure 6: Strategies of players in treatments with $p = 1/4$

non-chosen message were not elicited. Only 5% of Senders' messages and reported beliefs in all periods and all treatments are inconsistent with having some beliefs after the non-chosen message that would make the chosen message a best response to the reported beliefs. It is also worth noting that subjects' payoffs from belief statements were 37.75 points on average (34.47-39.48 depending on treatment). Reporting belief 0.5 would have earned a subject 38 points for sure, while reporting beliefs corresponding to the baseline model of empirical beliefs (i.e., reporting the empirical frequencies of types or actions observed so far) would have earned 40.97 points on average. It appears that subjects tried to make guesses but their attempts were not very successful.

4.3 Behavior with and without eliciting beliefs

In this section, behavior in the experiment with belief elicitation is analyzed and compared with behavior without elicitation of beliefs. Since belief elicitation may change the way subjects behave in a game, it is important to see if this happens in the signaling game under consideration.

Figure 6 shows the average strategies in treatments with $p = 1/4$, both in the new experiment with belief elicitation (solid lines, denoted with "b") as well as such strategies without belief elicitation (dotted lines, denoted with "nb") from DMP. The solid and dotted lines of the same color are rather close to each other in each panel. Thus, the differences

Proportions of strategies for $p = 1/4$ (Periods 21-36) and comparison tests				
	Senders		Receivers	
	$m_1 t_1$	$m_1 t_2$	$a_1 m_1$	$a_1 m_2$
$N1b$ vs $N1nb$	0.63 vs 0.60 (0.936)	0.71 vs 0.68 (0.873)	0.72 vs 0.67 (0.749)	0.70 vs 0.58 (0.337)
$K1b$ vs $K1nb$	0.20 vs 0.30 (0.256)	0.97 vs 0.95 (0.858)	0.91 vs 0.97 (0.360)	0.16 vs 0.16 (1.000)
$N1b$ vs $K1b$	0.63 vs 0.20 (0.014)**	0.71 vs 0.97 (0.005)***	0.72 vs 0.91 (0.019)**	0.70 vs 0.16 (0.003)***

Note: p -values in parentheses. 6 observations per treatment. For $N1b$ vs $N1nb$ and $K1b$ vs $K1nb$, $H_0 : Prop_{X1b} = Prop_{X1nb}$. For $N1b$ vs $K1b$, $H_0 : Prop_{N1b} \leq Prop_{K1b}$ for $m_1|t_1$ and $a_1|m_2$, $H_0 : Prop_{N1b} \geq Prop_{K1b}$ for $m_1|t_2$ and $a_1|m_1$. ** - $p < 0.05$; *** - $p < 0.01$.

Table 7: Non-parametric tests of differences between treatments for $p = 1/4$

in play between the cases in which beliefs are elicited and in which they are not appear minimal.

Table 7 shows the results of non-parametric tests based on matching groups as independent observations for the latter part of the sessions (Periods 21-36), when behavior is more stable.²³ The first two rows of the table indeed confirm that there are no statistically significant differences between the corresponding treatments in the proportions of the times with which strategies are played.

Figure 6 also shows that for $p = 1/4$ there is a difference between the treatment in which p is known and the treatment in which p is unknown. This difference is preserved in the new set of experiments with belief elicitation, and is also confirmed by non-parametric statistical tests in Table 7.

Strategies in treatments with $p = 1/2$ and $p = 3/4$ are similar and thus the data for these treatments are pooled in Figure 7, where the average strategies in these treatments with belief elicitation are shown as solid lines while the dotted lines show the average strategies without belief elicitation.

The overall trend appears similar in all panels, even if there are apparent differences in some panels. For some strategies ($m_1|t_2$ and $a_1|m_2$), non-parametric tests detect significant differences between some treatments with and without belief elicitation while for other strategies such differences are not detected.²⁴ For the comparison between treatments with known and unknown value of p , no differences are found, as in DMP for the game without belief elicitation. Thus, the results are mixed but overall the differences in behavior with

²³The same results holds for tests based on all periods or on the last eight periods. The data on which the tests are based are given in Section B.3 of Supplementary Materials, also for other tests in this section.

²⁴The tests are reported in Section B.4 of Supplementary Materials.

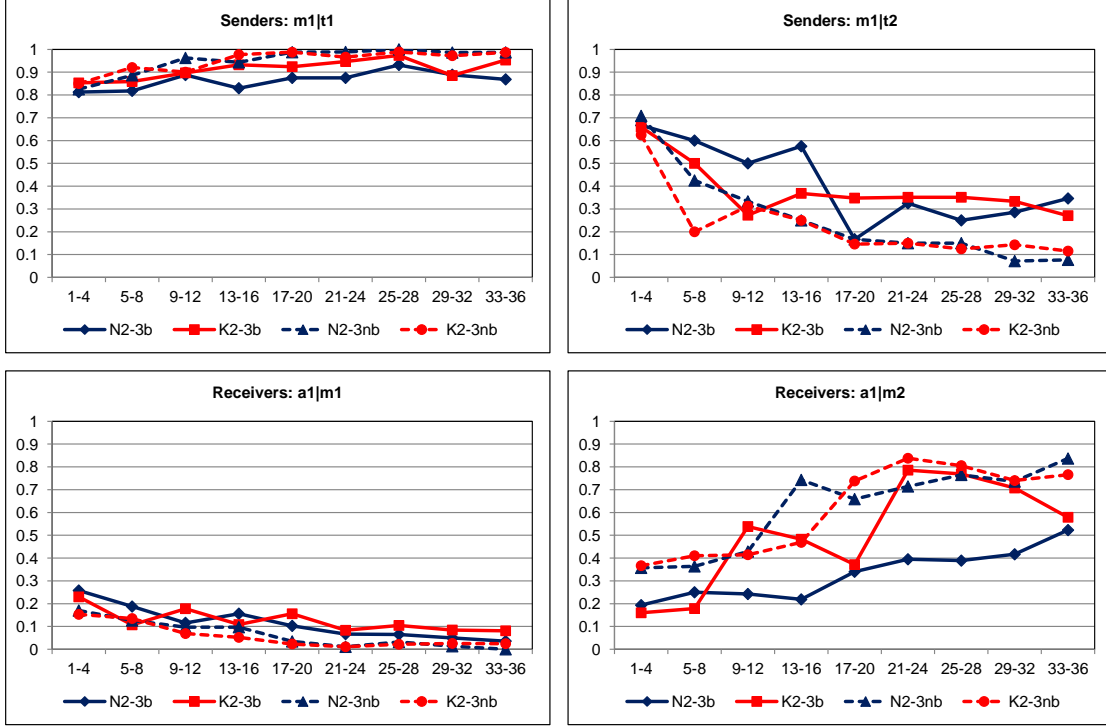


Figure 7: Strategies of players in treatments with $p = 1/2$ and $p = 3/4$

and without belief elicitation in treatments with $p = 1/2$ or $p = 3/4$ appear small.

Result 3 *For behavior in the game*

- (i) *Belief elicitation does not change behavior in treatments with $p = 1/4$; there are differences in behavior between treatments with known and unknown prior probability of types if $p = 1/4$;*
- (ii) *For $p = 1/2$ or $p = 3/4$, there are no differences in behavior between treatments with known and unknown probability of types; for some strategies there are (small) differences in behavior depending on whether beliefs are elicited while for other strategies there is no such differences.*

While one possible explanation for the latter result is hedging in the experiment with belief elicitation (some players may occasionally choose a strategy that is not a best response to their beliefs while compensating the possible loss with points earned from belief reports), a more likely explanation appears to be confusion from the more complicated tasks with belief elicitation. In a post-experiment questionnaire, some subjects said that they made choices randomly, and the proportions in the less frequent situations of type t_2 and message m_2 reflect this. In any case, the effect of elicitation appears small and confined to situations where play approaches a pure strategy. In situations where play is less predictable, as in treatments with $p = 1/4$, no effect of belief elicitation on play is observed.

5 Conclusion

In situations in which information about probabilities governing stochastic events is not provided at the beginning, subjects can learn about them from experience. To explain the results of the experiment in this paper, models based on the history of observations were considered. The results show that, roughly, beliefs often start from a uniform distribution and then adjust towards experienced outcomes. For the prior probability distribution of types, which is stationary, this process seems natural. For the posterior probabilities of Senders' types and for the probabilities of Receivers' actions, beliefs can in addition incorporate models of other players' behavior, such as experience-weighted attraction learning (Camerer and Ho, 1999, and its application to signaling games, Anderson and Camerer, 2000). Nevertheless, for parsimony the paper focuses on purely observational models.

Among the models that were considered, the model that fits the observed data best is the one with some weight on initial beliefs, with beliefs incorporating new observations slowly. There are some differences in the adaptation of beliefs about impersonal events (the determination of types) and about strategies. Subjects may have an initial belief about the impersonal process and change it in the direction of the observed frequencies slowly. For strategies the influence of the initial belief is weaker. Strategies are conscious choices of the opponent and it may make sense to realize that the opponent is also learning thus pre-conceived ideas about his or her behavior should get less weight.

The paper uses an approach in which beliefs are elicited only at some periods. This allowed subjects to have more data between elicitation rounds and thus get smoother reported beliefs. It may also make belief elicitation less prominent for the subjects thus helping to keep their behavior similar to the one in the same game but without belief elicitation (even though there appear to be small differences between behavior in games with and without elicitation when play approaches a pure equilibrium). Subjects also often played a best response to their beliefs showing that belief reporting and the choice of strategies tasks were taken seriously.

The analysis in the paper focuses on the models that fit data better on average. Subjects may be heterogeneous in their initial beliefs and update them using different parameters or even processes. While the extension to heterogeneous subjects is clearly potentially interesting, it would require more data collected for each subject. The present analysis gives a step for understanding the process of belief formation and updating in aggregate.

The results of the paper advance the understanding of belief formation processes and try to discriminate among alternative models of belief formation and adjustment. It is done here on the example of a signaling game, for which the importance of the common prior assumption is also demonstrated. With the theory of belief adjustment that is found to provide the best fit to the data in this paper, it may be easier to understand behavior in other economic situations involving uncertainty as well.

References

- Andersen, S., Fountain, J., Harrison, G.W., Hole, A.R., Rutström, E.E. (2012) “Inferring Beliefs as Subjectively Uncertain Probabilities”, *Theory and Decision* 73, 161-183.
- Anderson, C.M., Camerer, C.F. (2000) “Experience-Weighted Attraction Learning in Sender-Receiver Signaling Games”, *Economic Theory* 16, 689-718.
- Armantier, O., Treich, N. (2013) “Eliciting Beliefs: Proper Scoring Rules, Incentives, Stakes, and Hedging”, *European Economic Review* 62, 17-40.
- Binmore, K., Stewart, L., Voorhoeve, A. (2012) “How Much Ambiguity Aversion? Finding Indifferences between Ellsberg’s Risky and Ambiguous Bets”, *Journal of Risk and Uncertainty* 45, 215-238.
- Blanco, M., Engelmann, D., Koch, A.K., Normann, H.-T. (2010) “Belief Elicitation in Experiments: Is There a Hedging Problem?”, *Experimental Economics* 13, 412-438.
- Brandts, J., Holt, C.A. (1996) “Naive Bayesian Learning and Adjustment to Equilibrium in Signaling Games”, working paper, Instituto de Análisis Económico (CSIC), Barcelona and University of Virginia.
- Camerer, C.F., Ho, T. (1999) “Experience-Weighted Attraction Learning in Normal Form Games”, *Econometrica* 67, 827-874.
- Chen, Y., Katuščák, P., Ozdenoren, E. (2007) “Sealed Bid Auctions with Ambiguity: Theory and Experiments”, *Journal of Economic Theory* 136, 513-535.
- Cheung, Y.-W., Friedman, D. (1997) “Individual Learning in Normal Form Games: Some Laboratory Results”, *Games and Economic Behavior* 19, 46-76.
- Costa-Gomes, M.A., Weizsäcker, G. (2008) “Stated Beliefs and Play in Normal-form Games”, *Review of Economic Studies* 75, 729-765.
- Drouvelis, M., Müller, W., Possajennikov, A. (2012) “Signaling without a Common Prior: Results on Experimental Equilibrium Selection”, *Games and Economic Behavior* 74, 102-119.
- Eichberger, J., Kelsey, D., Schipper, B. (2008) “Granny versus Game Theorist: Ambiguity in Experimental Games”, *Theory and Decision* 64, 333-362.
- Fischbacher, U. (2007) “z-Tree: Zurich Toolbox for Ready-Made Economic Experiments”, *Experimental Economics* 10, 171-178.
- Gonzalez, C., Dutt, V. (2011) “Instance-Based Learning: Integrating Sampling and Repeated Decisions From Experience”, *Psychological Review* 118, 523-551.

- Greiner, B. (2015) "Subject Pool Recruitment Procedures: Organizing Experiments with ORSEE", *Journal of the Economic Science Association* 1, 114-125.
- Harsanyi, J.C. (1967) "Games with Incomplete Information Played by "Bayesian" Players. Part I. The Basic Model", *Management Science* 14, 159-182.
- Hertwig, R., Erev, I. (2009) "The Description-Experience Gap in Risky Choice", *Trends in Cognitive Sciences* 13, 517-523.
- Hollard, G., Massoni, S., Vergnaud, J.-C. (2016) "In Search of Good Probability Assessors: An Experimental Comparison of Elicitation Rules for Confidence Judgments", *Theory and Decision* 80, 363-387.
- Hyndman, K., Özbay, E.Y., Schotter, A., Ehrblatt, W. (2011) "Belief Formation: An Experiment with Outside Observers", *Experimental Economics* 15, 176-203.
- Ivanov, A. (2011) "Attitudes to Ambiguity in One-Shot Normal Form Games: An Experimental Study", *Games and Economic Behavior* 71, 366-394.
- Kelsey, D., le Roux, S. (2015) "An Experimental Study on the Effect of Ambiguity in a Coordination Game", *Theory and Decision* 79, 667-688.
- Li, Z., Loomes, G., Pogrebna, G. (2017) "Attitudes to Uncertainty in a Strategic Setting", *Economic Journal* 127, 809-826.
- Ludvig, E.A., Spetch, M.L. (2011) "Of Black Swans and Tossed Coins: Is the Description-Experience Gap in Risky Choice Limited to Rare Events?", *PLoS One* 6, e20262.
- Nickerson, R.S. (2004) *Cognition and Chance*, Lawrence Erlbaum Associates, Mahwah, NJ.
- Nyarko, Y., Schotter, A. (2002) "An Experimental Study of Belief Learning Using Elicited Beliefs", *Econometrica* 70, 971-1005.
- Offerman, T., Sonnemans, J., van de Kuilen, G., Wakker, P.P. (2009) "A Truth-Serum for Non-Bayesians: Correcting Proper Scoring Rules for Risk Attitudes", *Review of Economic Studies* 76, 1461-1489.
- Palfrey, T.R., Wang, S.W. (2009) "On Eliciting Beliefs in Strategic Games", *Journal of Economic Behavior & Organization* 71, 98-109.
- Rutström, E.E., Wilcox, N.T. (2009) "Stated Beliefs versus Inferred Beliefs: A Methodological Inquiry and Experimental Test", *Games and Economic Behavior* 67, 616-632.
- Sinn, H.-W. (1980) "A Rehabilitation of the Principle of Insufficient Reason", *Quarterly Journal of Economics* 94, 493-506.

Stahl, D.O., Wilson, P.W. (1994) “Experimental Evidence of Players Models of Other Players”, *Journal of Economic Behavior & Organization* 25, 309-327.

Stahl, D.O., Wilson, P.W. (1995) “On Players Models of Other Players: Theory and Experimental Evidence”, *Games and Economic Behavior* 10, 218-254.

Trautmann, S.T., van de Kuilen, G. (2015) “Belief Elicitation: A Horse Race among Truth Serums”, *Economic Journal* 125, 2116-2135.