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

Alex Possajennikov*

University of Nottingham

October 2011

Abstract

Using belief elicitation, the paper investigates the formation and the evolution of beliefs in a signaling game, in which the common prior on Sender's type is not induced. Beliefs are elicited both about the type of the Sender and strategies of the players. Results show that players often start with diffuse beliefs and update them in view of observations but not radically enough. An interesting result is that beliefs about types are updated slower than beliefs about strategies. In the medium run, for some specification of game parameters, this leads to outcomes being sufficiently different from the outcomes of the game in which a common prior is induced. It is also shown that elicitation of beliefs does not change the pattern of play considerably.

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

JEL Codes: C91, D83, C72

1 Formation of Beliefs about Uncertain Events

When making a decision in a situation involving uncertainty, individuals may form beliefs about the probabilities of various outcome of uncertain events. Within game theory, the Harsanyi (1967) approach to games with incomplete information postulates that players'

*School of Economics, University of Nottingham, University Park, Nottingham NG7 2RD, United Kingdom. Tel: +44 115 9515461, fax: +44 115 9514159, email: alex.possajennikov@nottingham.ac.uk

beliefs about the events describing their incomplete information are derived from a commonly known probability distribution. If this probability distribution is not known to the players, how do they form and update beliefs with experience?

This paper reports on an experiment in which the process of forming and updating beliefs is explored. Individuals play a signaling game in which one player, the Sender, has a piece of private information (type) and can send a message to another 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 Sender's type, the message, and the action. To take an appropriate action, the Receiver needs to form beliefs about Sender's type based on the message the Sender sends.

The Receiver can get an idea about the appropriate action inferring something about the Sender's type from the message sent. This inference may not be straightforward and Receiver's prior beliefs are important to form beliefs about type based on message.

Prior beliefs can be explicitly induced by specifying the probabilities of Sender's type. Without explicitly induced prior beliefs, players can learn from observations if the game is repeated often enough. Drouvelis, Müller, and Possajennikov (2011) (henceforth DMP) investigated how behavior can be different depending on whether the probabilities of Sender's types are known or not known before a series of interactions starts. The reason for possible difference is that without explicitly induced prior beliefs, players can use different prior beliefs and employ originally different strategies. Path dependence can then lead to possibly different medium to long run outcomes, even if learning from observations allows to approximate the probabilities of Sender's types.

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 (first applied to signaling games, albeit only for beliefs about strategies, in Brandts and Holt, 1996). However, whether beliefs are really updated in the way the model suggests could not be answered without direct observations of them.

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. Belief elicitation procedures have been used in experiments before (e.g. Nyarko and Schotter, 2002 and Costa-Gomes and Weizsäcker, 2008). Rutström and Wilcox (2009) discuss the methodological issues of the influence of belief elicitation procedure on the actual play. Whether belief elicitation affected play is tested in this paper (it does not appear so). While there are several procedures for eliciting beliefs, reviewed in Palfrey and Wang (2009), the most common quadratic scoring rule is used in the experiment reported.

Beliefs are elicited both about the type of the Sender and about strategies of the players. Sender's type is determined exogenously by a random device, thus it represents an "objective" uncertainty. Strategies of the players, on the other hand, are likely to be determined endogenously within the game. The strategic uncertainty is, thus, "subjective" and may depend on models the players use to determine the behavior of the opponent. Nickerson (2004, Ch. 8) argues that beliefs about "objective" uncertainty take more time to be revised. Since in the experiment both types of beliefs are observed, it should be possible to check whether some beliefs are formed and updated faster than others.

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 context, the principle is more applicable to beliefs about Sender's types. Beliefs about strategies can also be subject to this principle; however, some reasoning can be used to determine which strategy is more likely to be played by the opponent.

Thus, the main research questions of this paper are whether initial beliefs are close to uniform, how beliefs are updated, and whether some beliefs are updated faster than others. The data suggest that beliefs about Sender's types are indeed start close to uniform; even beliefs about strategies are not far from the uniform distribution. Beliefs are updated as observations accumulate in the natural direction of the frequency of events. However, updating is slow indicating that initial beliefs have a sizeable weight. Beliefs about types indeed appear to be updated slower than beliefs about strategies.

Given these properties of belief updating, the play in the game exhibits differences between the situations with known probabilities of Sender's types and unknown ones, due to path dependence in one of the treatments. This happens because starting from uniform initial beliefs takes the play to a different equilibrium than starting from known correct probabilities of Sender's types, if initial beliefs are not updated fast. In the other treatments, in the long-run there is no difference in behavior. Therefore the uncovered process of belief formation and updating has sometimes important consequences for long-run outcomes, and the paper identifies situations where this process matters and where it does not matter.

2 The Signaling Game and Belief Elicitation

Individuals were asked to play the signaling game given by the payoff tables below.

	Receiver		
	Type t_1	a_1	a_2
Sender	m_1	15, 10	80, 80
	m_2	25, 10	50, 50

	Receiver		
	Type t_2	a_1	a_2
Sender	m_1	80, 80	15, 30
	m_2	50, 50	25, 30

In the game, the type of the Sender (Player 1) is determined randomly, with the probability of Type t_1 being p and that of Type t_2 being $1-p$. Three value of p are considered, $p = 1/4$, $p = 1/2$, and $p = 3/4$. The Sender, knowing his type, chooses one of two messages, m_1 or m_2. The Receiver (Player 2) 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 the Receiver's action and given in the tables. The first number is the payoff of the Sender and the second number is the payoff of the Receiver.

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 type t_1, the second 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.[1]

Apart from the differences in the value of p, the other treatment difference in the experiment is that in some treatments this value is commonly known to the players, while in other treatments the value is not revealed to them. In this way it can be investigated how the information about the probability of Sender's type affects adjustment to 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 commonly known prior distribution and slow updating of prior beliefs about types, converges to the equilibrium $[(m_2, m_1), (a_1, a_2)]$ in the treatment with $p = 1/4$ and known, while in the other treatments the process converges to the equilibrium $[(m_1, m_2), (a_2, a_1)]$.

The naive process starts with a belief that the strategy of the opponent is uniform. With such a belief, both types of the Sender prefer to play m_1. When $p = 1/4$, the best response

[1]There 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 not observed in the data.

of the Receiver to the uniform strategy of the Sender is a_1 against both messages. Type 1 Sender then switches to m_2 and in response the Receiver switches to a_2 against m_2. The equilibrium $[(m_2, m_1), (a_1, a_2)]$ is reached. When $p = 1/2$ 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. The equilibrium $[(m_1, m_2), (a_2, a_1)]$ is reached.

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 very slowly, the play can follow the adjustment path to the equilibrium $[(m_1, m_2), (a_1, a_2)]$, as if $p = 1/2$ is known.

DMP show that there are no statistically detected differences between observed play in treatments in which the value of p is known or not for $p = 1/2$ or $p = 3/4$. For $p = 1/4$, there are differences in play depending on whether this value of p is known or not, although not as clean as predicted by the naive adjustment theory. One possible explanation is that the overall direction of adjustment depends on the speed of belief revision about the type, relative to the speed of belief revision about strategy. If the adjustment of type beliefs is much slower than that of the beliefs about strategy, the path in the previous paragraph is followed. On the other hand, if type beliefs are revised faster, the Receiver 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 the behavior in initial periods of treatments without commonly known value of p was not statistically different from behavior in the treatment with known value $p = 1/2$. While this provides an indirect evidence for the naive theory of belief formation, to understand better their initialization and adjustment, it is important to observe beliefs directly, as noted 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, as in Nyarko and Schotter (2002). 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 objective uncertainty, with a stationary distribution.[2] By contrast, the strategic uncertainty about the strategies of the opponent is random only from the view of the player, and its distribution may be

[2]The stationarity of the distribution was emphasized in the instructions.

changing as the opponent learns how to play the game. Nickerson (2004, Ch. 8) reports some evidence about different speed of belief formation depending on whether uncertainty is objective or about a person's performance. Nevertheless, the evidence is not overwhelming and the analysis presented in this paper is a further step towards understanding how players deal with such different kinds of uncertainty.

In the experiment belief elicitation is incentivized via a quadratic scoring rule, as e.g. in Nyarko and Schotter (2002) and Costa-Gomes and Weizsäcker (2008). While this works only for risk-neutral players, payoffs are such that risk-neutrality is not an implausible assumption.

In contrast to other papers that used belief elicitation, in the experiment beliefs are elicited not every period but every few periods. This is done in an effort to concentrate subject's efforts on this task rather than making it routine. It also allows 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 good representation of real ones.

3 Experiment and Belief Elicitation Design

The design of the experiment in DMP is followed, with addition of belief elicitation. The signaling game is described in the previous section. Subjects were assigned the role of either Sender or Receiver, and made corresponding decisions.

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. The quadratic scoring procedure gives payoff

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

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. Palfrey and Wang, 2009).

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

In the treatments in which the value of p is unknown, prior beliefs 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 Receiver's actions 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. 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. See the instructions (in Appendix A) for more details.

The value $A = 50$ and belief elicitation not every period were chosen for several reasons. To get enough incentives to think about beliefs, 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. Due to budgetary constraints, such high payoffs for beliefs were not possible if beliefs were elicited every period. 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. Also, subjects had more observations between the periods of belief elicitation and thus could have a better basis to form their view of probabilistic events.

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 $p = 3/4$.

The length of the sessions was 36 periods, to allow enough opportunities for learning, while at the same time not too long to make the task tedious. The sessions lasted approximately 90-100 minutes. In each session, the roles of Sender and Receiver were assigned randomly at the beginning. Then 8 or 16 participants were randomly matched within groups consisting of 4 Senders and 4 Receivers. The matching protocol and the type assignment was the same as in DMP. Points were 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 Centre for Decision

Research and Experimental Economics (CeDEx) laboratory at the School of Economics at the University of Nottingham in February-March 2009. There were 3 sessions in treatments $N1$ and $K1$, since these treatments are likely to produce the most interesting treatment difference. For each of the other treatments, one session was run. In each session 16 subjects participated, divided into two matching groups of 4 Senders and 4 Receivers, thus making two independent observations per session (one session, in treatment $K3$, had only 8 participants and one independent observation).

In the best equilibrium of the game, and with best predictions, a subject could earn £16.28. The uniformly random strategy, together with the uniform prediction, would have earned on average £10.16 per player. The average earnings were in fact £11.72 per subject, higher than the uniform way of playing and predicting, but way off the payoff in the best equilibrium and for the best predictions.

The main aim of the experiment was to explore the way the 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 the strategy of Receivers. Thus one part is that the belief of Senders are uniform. Receivers are asked about the posterior beliefs, as well as, in treatments with the unknown value of p, about the prior distribution of the Sender's type. While the prior is a distribution for a simple binary event, posterior beliefs reflect the beliefs about the strategy of the Sender. Thus there are further two parts of the hypothesis: the prior distribution is uniform, and the Sender's strategy is uniform.

The hypothesis is based on the principle of insufficient reason (see, e.g. Sinn, 1980, for a relatively recent analysis of it). If it is rejected, there are some reasons to initialize beliefs differently. The hypothesis is more likely to hold for prior beliefs about types, since strategic considerations can lead to different beliefs about actions and strategies.

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.

Hypothesis 3 *The behavior in the experiment with belief elicitation is not different from the behavior without it.*

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 patterns of play are different in the present experiment.

The hypothesis serves as a check on procedures. Players may behave differently depending on whether they are asked about their beliefs or not. If the hypothesis is not rejected, then beliefs elicitation is not changing the way the game is played.

4 Experiment Results

4.1 Behavior with and without eliciting beliefs

To begin, behavior in the experiment with belief elicitation is analyzed and compared with the behavior without the elicitation of beliefs. Thus, Hypothesis 3 is analyzed first.

Figure 1 shows the average strategies in treatments with $p = 1/4$, both in the new experiment with belief elicitation (solid lines) as well as such strategies without belief elicitaion (dotted lines) from DMP. The solid and dotted lines of the same color are rather close one to another in each panel. Thus the differences in play between cases in which beliefs are elicited and in which they are not appear minimal.

Table 1 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.[3] In the table, "*b*" refers to the treatment with elicited beliefs while "*nb*" to treatments without belief elicitation. The first two rows of the table indeed confirm that there are no statistically significant differences between corresponding treatments in proportions of the times given strategies are played.

[3] Same results holds for tests based on all periods or on the last eight periods.

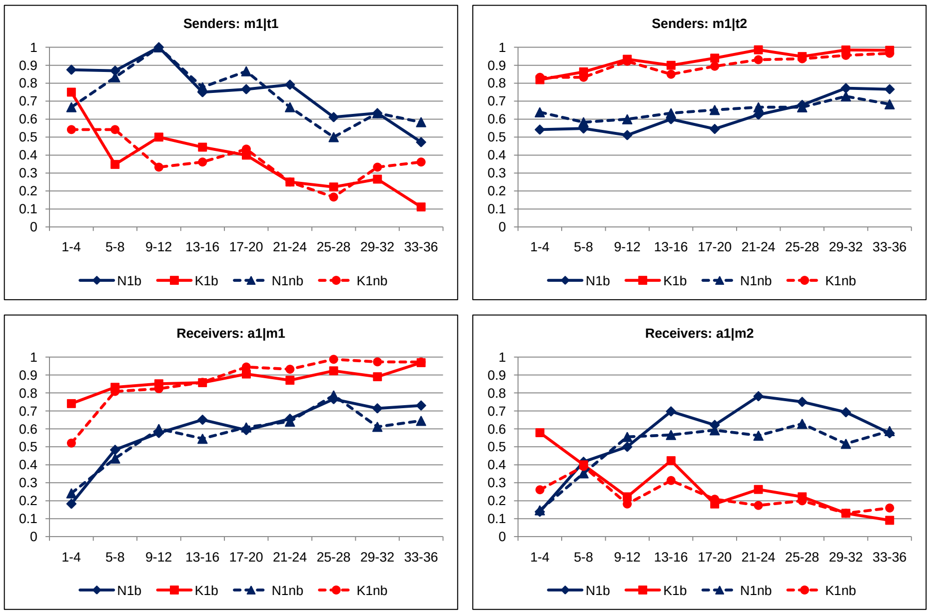

Figure 1: Strategies of players in treatments with $p = 1/4$

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	0.71 vs 0.68	0.72 vs 0.67	0.70 vs 0.58
$K1b$ vs $K1nb$	0.20 vs 0.30	0.97 vs 0.95	0.91 vs 0.97	0.16 vs 0.16
$N1b$ vs $K1b$	0.63 vs** 0.20	0.71 vs** 0.97	0.72 vs** 0.91	0.70 vs** 0.16

** - significant difference at $\alpha = 0.05$; 6 observations per treatment

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

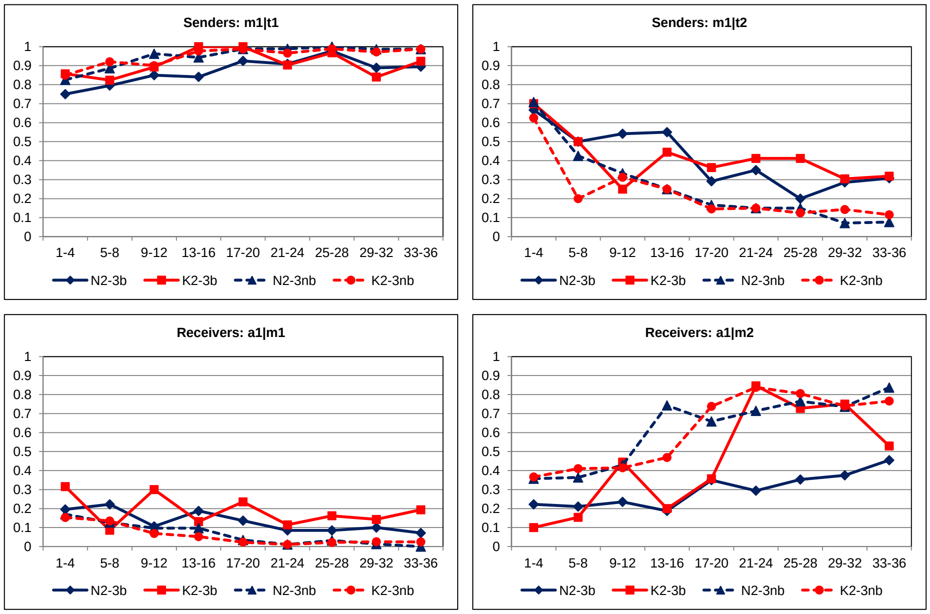

Figure 2: Strategies of players in treatments with $p = 1/2$ and $p = 3/4$

Figure 1 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 1.[4]

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

Although the use of messages as Type 2 Sender and the use of actions as Receiver after message m_2 appear erratic in the figure, it is a consequence of rather few observations as Type 2 and after message m_2. In these treatments, Senders are more often Type 1, and as such they overwhelmingly play m_1, which is almost exclusively answered by a_1. The two left panels of Figure 2 capture this from many observations of such behavior. Thus even if there are apparent differences in some panels, the overall trend is the same in all panels,

[4]Again, tests on all periods and the last eight periods produce the same results.

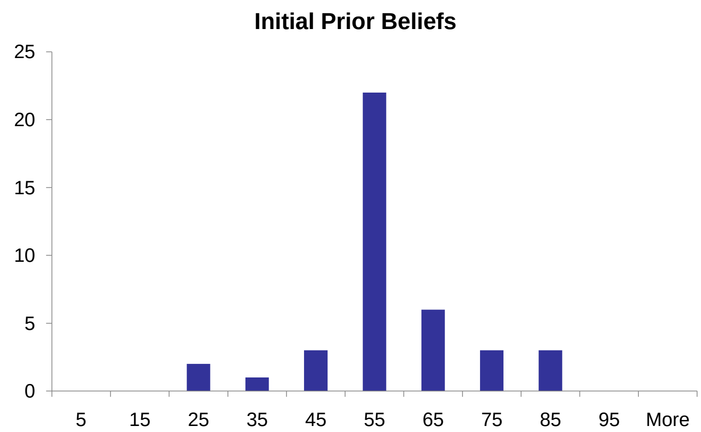

Figure 3: Histogram of initial beliefs about prior in treatments with unknown prior

and the differences are small in the panels that are based on more observations.

Result 1 *Belief elicitation does not change the behavior in treatments with $p = 1/4$. There is a difference in behavior between treatments with known and unknown prior if $p = 1/4$ and no such differences if $p = 1/2$ or $p = 3/4$. There are insufficient data to make confident conclusions for differences all strategies between treatments with or without belief elicitation with $p = 1/2$ and $p = 3/4$, but such differences appear small.*

4.2 Initial beliefs

For treatments in which the actual probability of Sender's type was not revealed, the most natural guess, based on the principle of insufficient reason, is that each of the two types is equally likely. Figure 3 presents the histogram of 40 observations of reported initial prior beliefs in the N treatments.

Most of the reported beliefs lie within the interval 45-55%, i.e. close to 0.5 probability of Type 1. The average reported prior belief is

Prior Beliefs	t_1 (40 obs)
Mean	0.53
St.Dev.	(0.13)

One-sample t-test and Wilcoxon signed-rank test do not reject the hypothesis that the median reported belief is equal to 0.5. Thus the prior belief is centered on 0.5 and, according to the histogram, is concentrated on this value.

Reported posterior beliefs of Receivers in Period 1 in the N treatments and in treatment $K1$ are given in the following table:[5]

Posterior	N treatments		Treatment $K1$	
beliefs	$t_1 \| m_1$ (30 obs)	$t_1 \| m_2$ (10 obs)	$t_1 \| m_1$ (20 obs)	$t_1 \| m_2$ (4 obs)
Mean	0.59	0.55	0.31	0.38
St.Dev.	(0.24)	(0.21)	(0.24)	(0.25)

Posterior beliefs in the N treatments are not far from 0.5. The Wilcoxon-Mann-Whitney test does not find significant difference between the posteriors for the two different messages, and the signed-rank test for paired observations does not detect a significant difference between reported prior and posterior beliefs.

The last two columns of the table report posterior beliefs in treatment $K1$. There is also no significant difference between the type beliefs after 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 no significantly different from 0.25 by the signed-rank test.

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

The beliefs of senders about actions of receivers in Period 1 are

Strategy	N treatments		Treatment $K1$	
beliefs	$a_1 \| m_1$ (30 obs)	$a_1 \| m_2$ (10 obs)	$a_1 \| m_1$ (20 obs)	$a_1 \| m_2$ (4 obs)
Mean	0.48	0.47	0.65	0.55
St.Dev.	(0.31)	(0.23)	(0.22)	(0.10)
Actual play	0.07	0.30	0.75	0.50

In the N treatments, the average beliefs of senders are quite close to 0.5, although they are heterogeneous (standard deviation is high). Non-parametric tests do not find a significant

[5]There are too few observation for the other treatments thus the analysis is done only for these treatments.

difference in these beliefs by message, or from the uniform belief 0.5 on type t_1. Note though that 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, Receiver indeed gets a higher payoff by choosing a_1. These beliefs also reflect to some extent the actual proportion of action a_1, at least after message m_1. It appears that senders did make some adjustment for strategic consideration of receivers already in Period 1 if the common prior $p = 1/4$ was induced. With unknown prior though, senders' beliefs are close to 50-50 chance of receiver taking either action.

Result 2 *Initial beliefs about the prior are close to uniform in treatments with unknown value of p. Initial posterior beliefs about Sender's type are not different from initial prior beliefs. Initial beliefs about actions of Receiver are close to uniform in the treatments with unknown value of p but put more weight on a_1 in treatment $K1$.*

To see that subjects took reporting of beliefs in Period 1 seriously, one can check whether they are consistent with the chosen message or action. Receivers play best response to the reported posterior beliefs in Period 1 76% of the time.

For senders it is not possible to determine whether their choice of message is indeed best response because they are not asked for their beliefs about the action of receiver in response to the non-chosen message. One possibility is to consider whether no beliefs about action after the non-chosen message would make the message played inconsistent with best response.[6] Since one can often find beliefs making the choice of message consistent with best response, only 5% of messages and beliefs in Period 1 are inconsistent with best response. Alternatively, one can assume that in Period 1 senders have the same beliefs about receivers' action after both messages. If this assumption is adopted, 70% of Senders' chosen messages and reported beliefs in Period 1 are consistent with best response.

[6]For the game played, Senders' message would be inconsistent with best response if they were choosing m_1 and if their beliefs of action a_1 in response to it were more than 11/13 for type t_1 and less than 2/13 for type t_2.

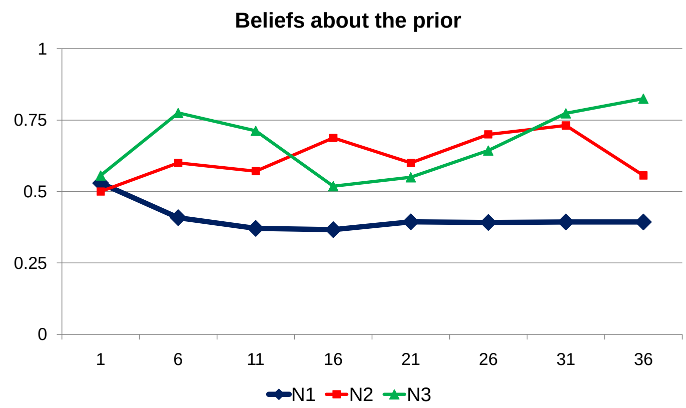

Figure 4: Evolution of beliefs about the prior

4.3 Belief adjustment

4.3.1 Beliefs about the prior

The table below shows the evolution of average belief of receivers about the prior (before seeing the message sent to them) probability of Sender's type for the three N treatments.

Prior beliefs about the probability of type t_1

Period	$N1$	$N2$	$N3$
1	0.53	0.50	0.56
6	0.41	0.60	0.78
11	0.37	0.57	0.71
16	0.37	0.69	0.52
21	0.39	0.60	0.55
26	0.39	0.70	0.64
31	0.39	0.73	0.77
36	0.39	0.56	0.83

Figure 4 illustrates these beliefs graphically.

Starting from the beliefs 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 (8 subjects in each of $N2$ and $N3$) than for $p = 1/4$ (24 subjects). Non-parametric tests for $N1$ treatment confirm that beliefs in the last period are different from those in the first period.[7] Thus it appears that beliefs are adjusted in the direction of experienced outcomes.

To analyze further the process of belief adjustment, 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 sender was type t_1 in a given receiver's observations. Let A_1^τ be the count for type t_1 and A_2^τ the count for type t_2 in period τ. If type t_i is observed in period τ, then $A_i^{\tau+1} = A^\tau + 1$, $A_j^{\tau+1} = A^\tau$ for $j \neq i$. The beliefs are $q_1^\tau = A_1^\tau/(A_1^\tau + A_2^\tau)$. The initial counts are $A_i^0 = 0$.

- *Forgetting.* This process behaves like the baseline process except that the counts are discounted: $A_i^{\tau+1} = \gamma A^\tau + 1$, $A_j^{\tau+1} = \gamma A^\tau$ 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 beliefs.* This process is like the baseline process except that the initial counts are not 0 but $A_i^0 = A$, where A is estimated from the data. Larger values of A would mean that new observations have less weight compared with the 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 sum of squared errors (SSE) between the prediction and the reported beliefs. The results of the estimations and the obtained minimized SSE scores are reported below

(280 obs)	Cournot	Empirical				50-50
		Base	Forgetting $\gamma = 0.97$	Init. strength $A_{Pr} = 14.7$	$\gamma = 0.99$ $A_{Pr} = 16.4$	
SSE	80.10	22.68	22.57	18.64	18.63	23.59

[7]There are too few observations for $N2$ and $N3$ treatment to make definite conclusions about such a difference.

The table contains also SSE for two other benchmark models. One is the Cournot model where beliefs are equal to the observation from one previous period (i.e. equal 1 if Sender was type t_1 in the previous period and 0 otherwise). Another model, reported in the last column, is the one that simply predict 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 the initial strength of beliefs do better, although the one with forgetting is not much different from the one without forgetting. It appears that the best model is one with the strength on initial beliefs $A_{Pr} = 14.7$. Since each new observation has weight 1, the value 14.7 indicates how slow the beliefs change.

4.3.2 Posterior beliefs about types and beliefs about strategies

The model with the strength on initial beliefs seems to fit the data best of the considered models for the prior beliefs. 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 measures of this speed.

For treatments $N1$ and $K1$, for which there are more observations, the evolution of average posterior beliefs of receivers is given in the following table.

Posterior beliefs	N1		K1	
Period	$t_1\|m_1$	$t_1\|m_2$	$t_1\|m_1$	$t_1\|m_2$
1	0.60	0.60	0.31	0.38
6	0.50	0.46	0.26	0.47
11	0.39	0.42	0.19	0.50
16	0.27	0.26	0.11	0.41
21	0.36	0.29	0.17	0.67
26	0.38	0.35	0.14	0.38
31	0.38	0.39	0.13	0.69
36	0.38	0.46	0.14	0.98

One can see from the numbers that there is type separation in treatment $K1$, as one of the separating equilibria is played, while the picture is much more mixed in treatment $N1$. Indeed, few matching groups converged clearly to either of the separating equilibria in this treatment. Figure 5 shows the evolution of posterior beliefs graphically, together with the

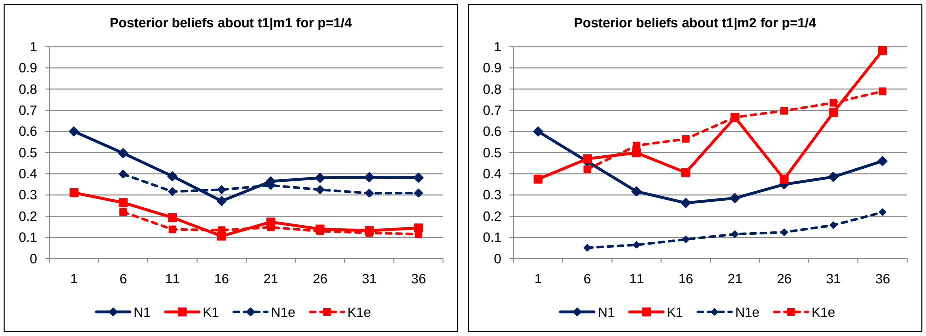

Figure 5: Evolution of posterior beliefs in treatments with $p = 1/4$

predictions of the best adjustment model (dotted lines), which is explained in more details below.

Posterior beliefs appear to start close to 0.5 (lower for the $K1$ treatment) and then move generally in the direction of experienced outcomes (which 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 (expect for beliefs about $t_1|m_2$ in treatment $N1$). Subjects seem to learn something about the posterior beliefs over time.

To see which adjustment model fits best, the same models as for the prior beliefs were considered, with the following results:[8]

(532 obs)	Reported	Cournot	Empirical Base	Empirical $\gamma = 0.99$ $A_{Ps} = 5.8$	50-50
SSE		121.90	54.55	45.45	69.65
Best resp.	0.80	0.71	0.77	0.76	0.48

The model with an initial strength of beliefs has the lowest SSE score. An interesting observation that the estimated strength parameter of this model, $A_{Ps} = 5.8$ is considerably

[8]Scores 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. Thus only the baseline and the full (initial strength and forgetting) scores are reported.

lower than such a parameter for the prior beliefs, $A_{Pr} = 14.7$. It appears that posterior beliefs are updated faster than prior ones, possibly because posterior beliefs incorporate beliefs about strategies as well, which are updated faster than beliefs about "objective" uncertain process.

The table also reports the proportion of choices that were best responses to reported beliefs (column "Reported") or that would be best responses to beliefs predicted by the model. Receivers chose best response to their reported beliefs 80% of the time, while if their beliefs were following the best adjustment model, their actions would have been best responses 76% of the time. This is close to 80%, thus the adjustment model reflects reported beliefs to some extent.

Senders in the experiment reported beliefs about receivers' action in response to the message sent. For treatments $N1$ and $K1$, the following table presents the evolution of these beliefs about strategies (on average).

Strategy beliefs	N1		K1	
Period	$a_1\|m_1$	$a_1\|m_2$	$a_1\|m_1$	$a_1\|m_2$
1	0.51	0.49	0.65	0.55
6	0.36	0.38	0.66	0.48
11	0.48	0.31	0.69	0.10
16	0.38	0.38	0.56	0.51
21	0.52	0.54	0.77	0.51
26	0.67	0.59	0.76	0.37
31	0.72	0.56	0.74	0.22
36	0.64	0.66	0.73	0.10

There is again a clearer separation of beliefs about receivers' responses for treatment $K1$ than for treatment $N1$, because the play in the $K1$ treatment converges to one of the separating equilibrium while in the $N1$ treatment in most of the matching groups there is no convergence. Figure 6 illustrates the evolution of average reported beliefs, together with the predictions of the best adjustment model (dotted lines, see below).

Strategy beliefs also start close to 0.5 in treatment $N1$ and from a higher value in treatment $K1$. Then they move to some extent in the direction of experienced outcome although this movement is less clear than for the prior or posterior beliefs about types. Indeed, non-parametric tests detect statistical difference between reported strategy beliefs in

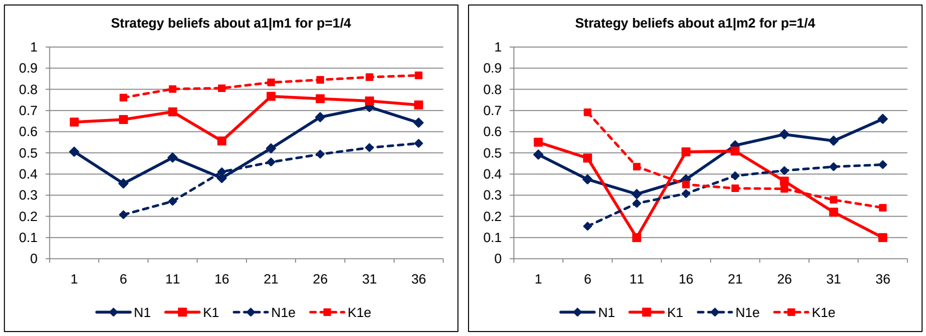

Figure 6: Evolution of beliefs about actions in treatments with $p = 1/4$

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 were considered, with the following results:[9]

(532 obs)	Cournot	Empirical		50-50
		Base	$\gamma = 1.02$ $A_{St} = 6.6$	
SSE	125.14	53.45	44.71	61.27

The lowest SSE score is again achieved by the model with an initial strength of beliefs. The estimated strength parameter of this model, $A_{St} = 6.6$ is close to such a parameter from the estimation of posterior beliefs $A_{Ps} = 5.8$ and is lower than such a parameter for the prior beliefs, $A_{Pr} = 14.7$. It seems that beliefs about strategies are updated faster than beliefs about the prior probability of Sender's type.

The observations can be summarized in the following result:

Result 3 *Beliefs adjust towards observed realizations of the relevant events. In aggregate, the model with a weight on initial beliefs explains reported beliefs better than the other models. The weight on initial beliefs is larger for beliefs about the prior than for beliefs about the posteriors or about the strategies.*

[9] Scores 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. Thus only the baseline and the full (initial strength and forgetting) scores are reported.

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 about an "objectively" uncertain process. Prior probability of types is "objectively" uncertain, while such posterior probability and probability of a given action of receiver depend on the opponent, thus beliefs about such probabilities are updated faster, which is represented by a lower weight on initial beliefs.

As observed above, receivers played best response to their beliefs 80% of the time. For senders, it is not possible to determine whether their messages were fully consistent with their reported beliefs because beliefs about Receiver's action after 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 non-chosen message that would make the chosen message best response to reported beliefs. It is also worth noting that subjects' payoffs from belief statements were 36-37 points on average (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 would have earned 39-40 points. It appears that subjects tried to make guesses but their attempts were not very successful.

5 Conclusion

In a situation in which probabilistic information is not provided, subjects learn about it from experience. The results reported in this paper show that, roughly, belief adjustment starts from a uniform prior and adjusts towards experienced outcomes. The model that fits observed data best is the one with some weight on initial beliefs, with beliefs incorporating new observations slowly.

The paper uses a novel approach in that beliefs are elicited only at some periods. This allowed subjects to make the experience between elicitation rounds smoother and thus get smoother reported beliefs. It makes belief elicitation less prominent for the subjects thus helping to keep their behavior similar to a similar experiment without belief elicitation.

It is confirmed that when no information is provided to the subjects, their beliefs concentrate around the uniform distribution. Beliefs are then updated generally towards the observed frequencies. Commonly subjects did play a best response to their beliefs showing that belief reporting and the choice of strategies tasks were taken seriously.

There are some differences in adaptation of beliefs about impersonal events (determination of types) and about strategies. Subjects have a prior about the impersonal process and change it in the direction of the observed frequencies slowly. For strategies the influence of prior 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.

There are some issues that are not addressed yet in the analysis, particularly the heterogeneity of subjects. Obviously, subjects can have different priors and update them use different parameters or even processes. The extension to heterogeneous subjects is left for future research.

The results of the paper advance the understanding of belief formation processes. It is done here on the example of a signaling game, for which the importance of the common prior assumption is also demonstrated. With a theory of belief adjustment, it may be easier to understand behavior in other economic situations involving uncertainty as well.

A Instructions for the treatment with unknown value of p

Please read these instructions carefully. Please do not talk to other people taking part in the experiment and remain quiet throughout. If you have a question, please raise your hand. We will come to you to answer it.

In this experiment you can earn an amount of money, depending on which decisions you and other participants make. The experiment consists of 36 rounds, in each of which you can earn Points. Your payout at the end of the experiment is equal to the sum of Points you earn in all rounds, converted to pounds. For every 10 Points you will be paid 5p.

Description of the experiment

Participants are assigned the role of either "A-participant" or of "B-participant". In each round of the experiment, all participants are matched randomly in pairs, one from each role. A random draw determines the type of the A-participant, which can be either "Type 1" or "Type 2". The random draw is such that with an $X\%$ chance the A-participant is of Type 1, and with a $(100 - X)\%$ chance of Type 2. There is a new random draw each round, and the value of X is constant over all rounds of the experiment. After the random draw, the A-participant is informed about his/her type and decides between options "C" and "D". After that, the B-participant is informed about which option was chosen by the A-participant, but not about the type of the A-participant, and chooses between options "E" and "F". The payoffs of the two participants are determined according to the tables overleaf on page 2.

In some rounds of the experiment, the B-participant is asked to predict the type of the matched A-participant, both before and after the A-participant has chosen an option, and the A-participant is asked to predict the option that will be chosen by the matched B-participant. You are asked "What is the chance that the participant is of Type 1 / chooses option E" and "What is the chance that the participant is of Type 2 / chooses option F". You answer with two numbers Y and Z between 0% and 100%, and the sum of the two numbers should be 100. The points you earn depend on your prediction and on the actual type or option chosen by the participant according to the formulas overleaf on page 3.

[In the treatments with known value of p, X was explicitly given, e.g. 75. In the last paragraph, the word "before" was deleted, i.e. the B-participant was asked only after the A-participant has chosen an option.]

Payoffs

Payoffs from the choice of options

The payoffs of both participants depend on the A-participant's type, the option chosen by the A-participant and the option chosen by the B-participant.

The A-participant's payoffs

The payoffs of the A-participant (in blue) in each round are given in the following two tables (along with the B-participant's payoffs in red). For the A-participant of Type 1, payoffs are given by the table on the left, and for the A-participant of Type 2, by the table on the right.

Payoff table for Type 1 of the A-participant:

Decision of the A-participant		Decision of the B-participant: **E**	**F**
	C	**15**, 10	**80**, 80
	D	**25**, 10	**50**, 50

Payoff table for Type 2 of the A-participant:

Decision of the A-participant		Decision of the B-participant: **E**	**F**
	C	**80**, 80	**15**, 30
	D	**50**, 50	**25**, 30

The B-participant's payoffs

The payoffs of the B-participant (in red) in each round are given in the following two tables (along with the A-participant's payoff in blue). If the A-participant chose option "C", the payoffs are given by the table on the left, and if the A-participant chose option "D", by the table on the right.

Payoff table for the B-participant if A-participant chose option "C":

Type of the A-participant		Decision of the B-participant: **E**	**F**
	1	15, **10**	80, **80**
	2	80, **80**	15, **30**

Payoff table for the B-participant if A-participant chose option "D":

Type of the A-participant		Decision of the B-participant: **E**	**F**
	1	25, **10**	50, **50**
	2	50, **50**	25, **30**

Payoffs from predictions

The payoffs of both participants depend on the prediction and on the actual type of, or option actually chosen by, the matched participant.

The A-participant's payoffs

If an A-participant predicts that the chance that the B-participant chooses option "E" is $E\%$ and the chance that the B-participant chooses option "F" is $F\% = (100 - E)\%$, the points earned are

$$\begin{array}{ll} 50 \cdot (1 - (1 - E/100)^2) & \text{if the B-participant actually chooses “E”} \\ 50 \cdot (1 - (1 - F/100)^2) & \text{if the B-participant actually chooses “F”} \end{array}$$

rounded to the nearest integer.

The B-participant's payoffs

If a B-participant predicts that the chance that the A-participant is of Type 1 is $Y\%$ and the chance that the A-participant is of Type 2 is $Z\% = (100 - Y)\%$, the points earned are

$$\begin{array}{ll} 50 \cdot (1 - (1 - Y/100)^2) & \text{if the A-participant actually is of Type 1} \\ 50 \cdot (1 - (1 - Z/100)^2) & \text{if the A-participant actually is of Type 2} \end{array}$$

rounded to the nearest integer.

Note that you get the maximum 50 points when you predict, for example, that the chance of Type 1 is 100% and Type 1 actually happens, or that the chance of Type 1 is 0% and Type 2 actually happens. You get 0 points if you prediction is completely wrong. You get an intermediate number of points if you predict that the chance of each type or of each action is between 0% and 100%. The formulas are designed in such a way that you maximize your expected payoff from your prediction if you state your true belief about the chance of the type of the A-participant, or of the action about to be chosen by the B-participant.

Summary

To give you an overall picture of the rules, the timing of events in each round can be summarized as follows:

1. The computer randomly matches participants in pairs.
2. The computer randomly determines the A-participant's type. With an $X\%$ chance the A-participant is of Type 1 and with a $(100 - X)\%$ chance of Type 2. The value of X is constant over all rounds of the experiment.

3. The A-participant is informed about his/her type. Then the A-participant chooses between options "C" and "D".

4. The B-participant is informed about the choice of the A-participant, but not about his/her type. Then the B-participant chooses between options "E" and "F".

5. Payoffs result as described in the tables above.

6. In some rounds, the participants are asked to predict the type of, or the option that will be chosen by, the matched participant. Payoffs for these predictions are added to the payoffs above.

Number of rounds, role assignment and matching

The experiment consists of 36 rounds.

The role of either the A-participant or the B-participant will be randomly assigned to each participant in the room at the beginning of the experiment. You will then keep the same role during the entire experiment.

In each round the computer will randomly match one A-participant and one B-participant from a group of eight subjects. The matching is completely random, meaning that there is no relation between the participant you have been matched with last round (or any other previous round) and the participant with whom you are matched in the current round.

References

[1] Brandts, J., Holt, C.A. (1996) "Naive Bayesian Learning and Adjustment to Equilibrium in Signaling Games", working paper, Instituto de Analisis Economico, Barcelona, and University of Virginia.

[2] Costa-Gomes, M.A., Weizsäcker, G. (2008) "Stated Beliefs and Play in Normal-form Games", *Review of Economic Studies* 75, 729-765.

[3] Drouvelis, M., Müller, W., Possajennikov, A. (2011) "Signaling without a Common Prior: Results on Experimental Equilibrium Selection", *Games and Economic Behavior*, forthcoming.

[4] Harsanyi, J.C. (1967) "Games with Incomplete Information Played by "Bayesian" Players. Part I. The Basic Model", *Management Science* 14(3), 159-182.

[5] Nickerson, R.S. (2004) *Cognition and Chance*, Lawrence Erlbaum Associates, Mahwah, NJ.

[6] Nyarko, Y., Schotter, A. (2002) "An Experimental Study of Belief Learning Using Elicited Beliefs", *Econometrica* 70, 971-1005.

[7] Rutström, E.E., Wilcox, N.T. (2009) "Stated Beliefs versus Inferred Beliefs: A Methodological Inquiry and Experimental Test", *Games and Economic Behavior* 67(2), 616-632.

[8] Palfrey, T.R., Wang, S.W. (2009) "On Eliciting Beliefs in Strategic Games", *Journal of Economic Behavior & Organization* 71, 98-109.

[9] Sinn, H.-W. (1980) "A Rehabilitation of the Principle of Insufficient Reason", *Quarterly Journal of Economics* 94(3), 493-506.