Belief Formation in a Signaling Game without Common Prior

Alex Possajennikov*

University of Nottingham

February 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 sufficiently 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, belief elicitation

JEL Codes: D83, C72, C91

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' 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?

*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

This paper analyzes 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 (called Sender's 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 (2009) 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 Drouvelis et al. (2009) were derived based on a belief updating process (first applied to signaling games, albeit only for beliefs about strategies, in Brandts and Holt, 1996), but whether beliefs are really updated in the way the model suggest could not be answered without direct observations of them.

In the experiment reported in this paper, subject 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 (2007) 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 both players. The type of the Sender 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" (see 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 context of the analyzed game, the principle is more applicable to the beliefs about Sender's types. Beliefs about strategies can also be subjected to this principle; however, some reasoning can be used to determine which strategy is more likely.

Thus, the main research questions of this paper concerning beliefs is whether initial beliefs are close to uniform, how they 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. Subjective probabilities are updated slowly indicating that the initial beliefs have a sizeable weight. Beliefs about types may be indeed updated slower than 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 non known ones, due to path dependence in one of the treatments. This happens because starting from uniform initial beliefs would take the play to a different equilibrium than starting from known correct probabilities of Sender's types, if the uniform initial beliefs are not updated fast.

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 the 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 the 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 is 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 the Sender's type affect 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 Drouvelis et al. (2009) 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 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

[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.

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 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.

Drouvelis et al. (2009) show that there are no statistically detected differences between observed play in treatment when the value of p is known or not for $p = 1/2$ or $p = 3/4$. When $p = 1/4$, there are differences in play depending on whether p is known or not, although not as clean as predicted by the naive adjustment theory. One possible explanation is that 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 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 Drouvelis et al. (2009), 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 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

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

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 Drouvelis et al. (2009) 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

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

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. Palfrey and Wang, 2009).

The experiment contains treatments with and without the known probabilities of the 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.

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 Drouvelis et al. (2009). Points were converted to pounds at the rate of £0.05 for 10 points.

The new 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 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 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 Drouvelis et al. (2009). 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 pattern 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 Drouvelis et al. (2009). 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. The non-parametric tests based on matching groups as independent observations confirm this impression.

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.

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, it is a consequence of rather few observations as Type 2 and

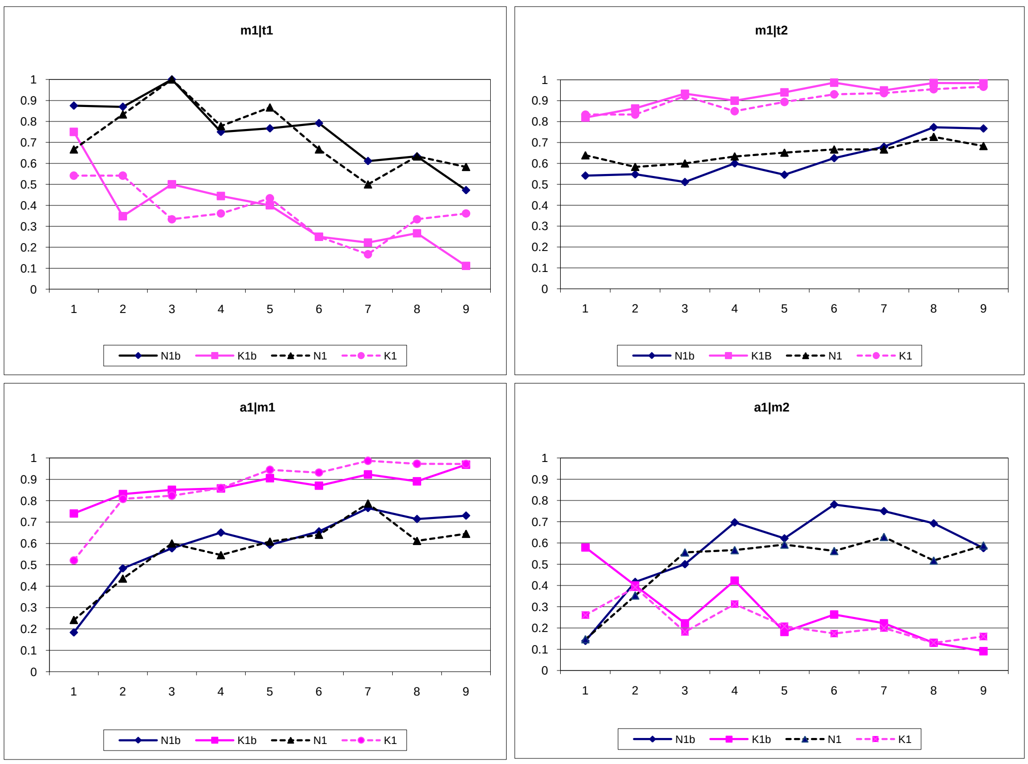

Figure 1: Strategies of players in treatments with $p = 1/4$

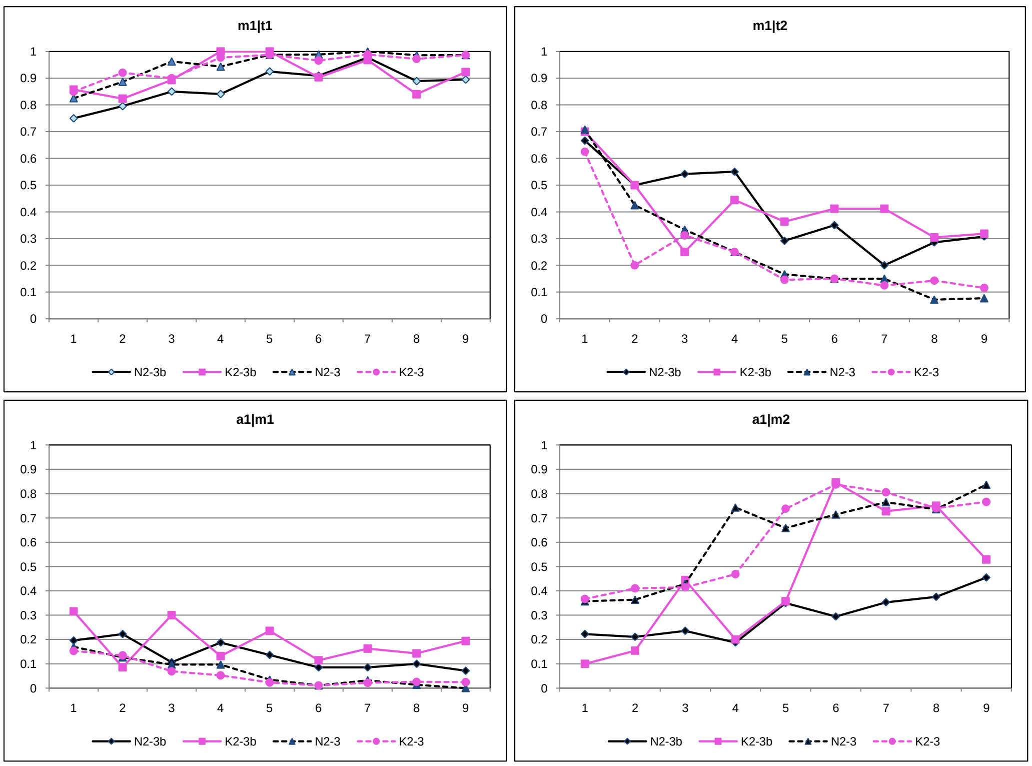

Figure 2: Strategies of players in treatments with $p = 1/2$ and $p = 3/4$

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 the right panels, the panels that are based on more observations do not present perceptible differences.

Result 1 *Belief elicitation does not change the behavior in treatments with $p = 1/4$. There are insufficient data to make confident conclusions for all strategies in treatments with $p = 1/2$ and $p = 3/4$, but for the most common Sender's type and the most common message faced by the Receiver, belief elicitation also does not change the behavior.*

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. The following table presents the average beliefs of Receivers in Period 1 in treatments with unknown p:

Receivers	t_1	$t_1\|m_1$	$t_1\|m_2$
Beliefs	0.53	0.59	0.55
No. of observations	40	30	10

The two last columns represent posterior beliefs. If taken as representative of a player, they are actually inconsistent with the reported prior as the two posterior beliefs cannot be both higher than the prior beliefs. However, these are average beliefs of many players and therefore they are not necessarily representative of any of them. Nevertheless, the fact that they all are not far from 0.5 indicate that initial beliefs are not far from uniform, certainly for the prior and possibly for the strategies.

The beliefs of the Senders in Period 1 are

Senders	$a_1\|m_1$	$a_1\|m_2$
Beliefs	0.48	0.47
No. of observations	30	10

The average beliefs of Senders are also quite close to 0.5.

It is understandable that the belief about t_1 is close to uniform as the Sender type realization comes from objective uncertainty. However, even for posterior beliefs of the

Receiver and Sender's belief about the strategies of the Receiver, beliefs are close to uniform. Perhaps the situation appears to the subjects sufficiently complicated at the beginning so that they fall back to the rule of thumb that events ought to be equally likely.

Result 2 *Initial beliefs in treatments with unknown value of p are close to uniform. Thus holds for the prior probability of Sender's type, for posterior such probability, and for the belief by Senders about Receivers strategies.*

4.3 Belief adjustment

To make a picture of how beliefs evolve and reflect the realization of types and strategies, the following procedure is used. Recall that beliefs were elicited every five periods, starting with period 1, making a total of eight belief responses. Between the periods of belief elicitation, a player could observe five realization of variables.

First it is checked whether belief report make sense in that players then choose actions that are best given these beliefs. For Receivers it is straightforward to find out whether the action chosen is best given their beliefs about types of Sender after observing the message sent. For Senders it is not possible to determine whether the message chosen is a best response since beliefs for the unsent message are not elicited. Consistency is checked in a weak form by checking whether the message could be best given some beliefs for the unsent message. For the game played, Senders would not be playing the best message 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.

In treatment $N1$ Receivers choose best action 75% of the time. While not overwhelming, it is better that the random action choice. Sender messages are consistent with possible best response 91% of the time. In treatment $K1$ Receivers choose best response 80% of the time, and Senders' behavior is consistent with best response in 98% of cases. In the other treatments Receivers choose best response $77-84\%$ of the time, while Senders actions are consistent with best response $92-98\%$ of the time.

After sending a message, Senders were asked what their beliefs were regarding Receiver's action. Therefore they saw five observations though possibly for the message different from the one they sent in the elicitation period. The table below reports the average beliefs of Senders about Receivers' strategies in each period in which they were elicited, in treatments

$N1$ and $K1$.

Senders	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

Receivers were asked about their beliefs about the type of Sender before receiving a message and after receiving a message. Thus both prior and posterior beliefs were elicited. There is again eight observation of responses, and the players could observe five realizations of types and associated messages in between. The table below shows the average beliefs of Receivers about Sender's type, both prior to seeing the message and posterior to it, in treatments $N1$ and $K1$ for each period of belief elicitation.

Receivers	N1			K1	
Period	t_1	$t_1\|m_1$	$t_1\|m_2$	$t_1\|m_1$	$t_1\|m_2$
1	0.53	0.60	0.60	0.31	0.38
6	0.41	0.50	0.46	0.26	0.47
11	0.37	0.39	0.42	0.19	0.50
16	0.37	0.27	0.26	0.11	0.41
21	0.39	0.36	0.29	0.17	0.67
26	0.39	0.38	0.35	0.14	0.38
31	0.39	0.38	0.39	0.13	0.69
36	0.39	0.38	0.46	0.14	0.98

Figure 3 displays belief adjustment in treatments $N1$ and $K1$. In the figure, together with the eight average belief responses, the dotted lines show what would the belief response be if each player simply reported the observed frequencies of the relevant event. The solid lines in the figures show the actual belief responses, and the dotted lines show the frequencies of Senders' types or Receivers' actions up to the period of elicitation. The figure confirms

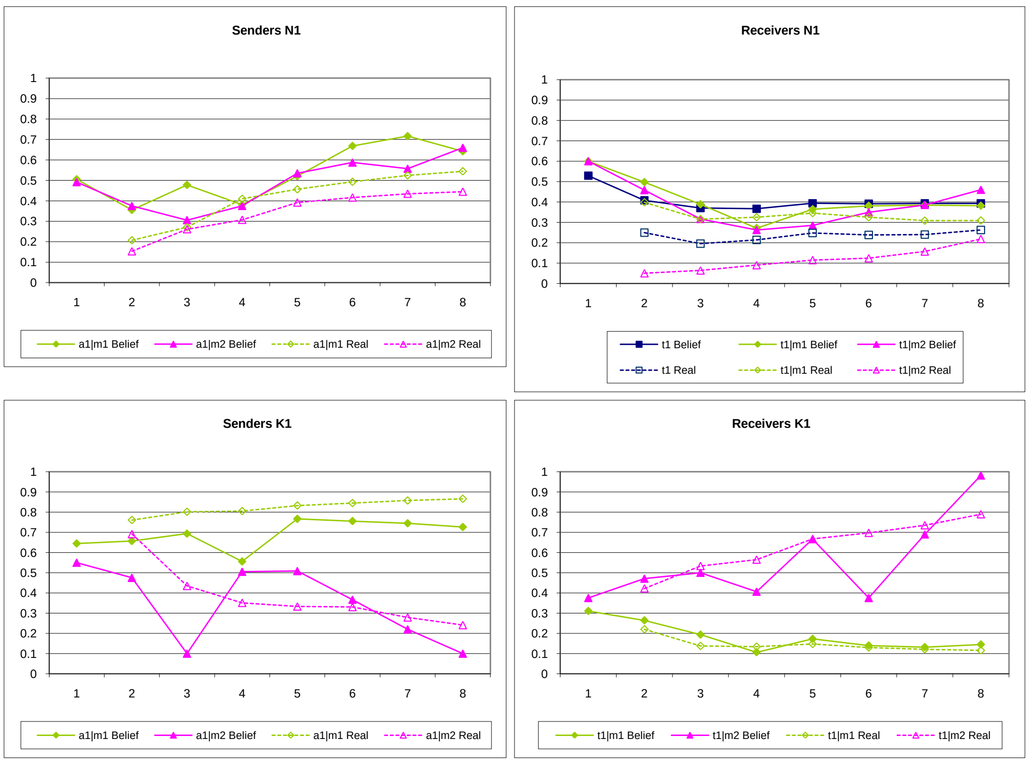

Figure 3: Belief adjustment in treatments with $p = 1/4$

that in treatment $N1$, initial beliefs, both by Receivers about types and by Senders about actions, are close to 0.5. Beliefs then move towards frequencies, and then roughly follow the direction of movement of frequencies. This pattern is consistent with players having a uniform prior which is adjusted in view of experience, albeit slowly.

Note that Senders seem to adjust beliefs close to the frequencies faster than the Receivers. Recall that Senders report beliefs about Receivers' strategies while Receivers has to report beliefs about Senders' types, which are based on objective uncertainty. Even the posterior beliefs of Receivers are based to some degree on prior beliefs about the type which are updated slowly.

In treatment $K1$ beliefs track frequencies quite close, specially for Receivers. Here, prior probabilities are given thus prior beliefs about types do not play a rile. Without having to form an estimation of objective uncertainty, players seem to be able to estimate opponent's strategies reasonably closely. Beliefs of Senders about the action after message m_1 exhibit a similar patter as in treatment $N1$, with beliefs consistently smaller than the actual frequencies. In any case, it seems that the strength of the initial prior probabilities is not as large as in treatment $N1$, and beliefs mostly adjust with experience.

In the other treatments beliefs also adjust towards the frequencies of realization. Summarizing,

Result 3 *Beliefs adjust towards observed realization of the relevant events. There is some evidence that beliefs about types (objective uncertainty) adjust slower than beliefs about strategies of other players.*

A more detailed model of belief adjustment explaining the observed patters of beliefs is left for future research.

5 Conclusion

In a situation where probabilistic information is not provided, subjects learn about it from experience. The results reported in this paper show that belief adjustment starts from a uniform prior and adjusts towards experienced outcomes.

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 albeit slowly. For strategies the influence of prior is weaker.

There are some issues that were 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 B-participant: **E**	**F**
Decision of the A-participant	**C**	**15**, 10	**80**, 80
	D	**25**, 10	**50**, 50

Payoff table for Type 2 of the A-participant:

		Decision of the B-participant: **E**	**F**
Decision of the A-participant	**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":

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

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

		Decision of the B-participant: **E**	**F**
Type of the A-participant	**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

$$50 \cdot (1 - (1 - E/100)^2) \quad \text{if the B-participant actually chooses "E"}$$
$$50 \cdot (1 - (1 - F/100)^2) \quad \text{if the B-participant actually chooses "F"}$$

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

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

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. (2009) “Signaling without Common Prior: An Experiment”, CeDEx discussion paper 2009-08, University of Nottingham.

[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. (2007) "Stated Beliefs versus Inferred Beliefs: A Methodological Inquiry and Experimental Test", working paper, University of Central Florida and University of Houston.

[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.